LANCASTER UNIVERSITY

Computing Department

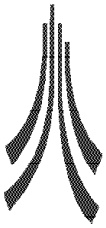

The OSI 95 Transport Service and the New Environment : a Contribution to the Discussion on the Enhanced Communications Functions and Facilities.

Author(s) : Andrew Campbell, David Hutchison and Helmut Leopold

Supervisor :

Reviewer(s) :

Contents

1 Introduction 1
1.1 Scope 1
2 Distributed Multimedia Communication Requirements 2
Introduction 2
General Multimedia Communication System Requirements 2
Continuous Media 2
Error Control 3
Timeliness 3
Error Recovery and Application Requirements 3
End-to-End Flow Control 4
End-to-End Argument 4
Profile and Class of Service Selection 4
Other Issues 4
Quality of Service 5
Transport Level QOS Parameters 5
Guaranteed QOS 6
Dynamic QOS 7
Group and Multicast Support 7
The Concept of Groups 8
Multicast Services 8
Orchestration 8
Orchestration Functionality 8
Enhanced Transport Services 10
Continuous Media Service 10
Transaction Oriented Service 10
Bulk Data Transfer 11
Baseline Requirements 11
3 The OSI 95 Transport Service 11
Introduction 12
TPX Features 12
Flexible Error Control 12
TPX QOS Parameters 12
Compulsory QOS 13
QOS Negotiation 13
QOS Degradation 14
Continuous Media Support 14
Group and Multicast Support 14
Out-of-band Data Transfer Facility 14
Graceful Release Mechanism 14
The OSI 95 Transport Service 15
The OSI 95 Enhanced Connection-mode Transport Service 15
Fast Connection -mode Transport Service 16
Unacknowledged Connectionless-mode Transport Service 17
Acknowledged Connectionless-mode Transport Service 18
Request/Response Transport Service 19
4 Evaluation 20
Error Detection and Recovery 21
Timeliness 22
End-to-end Flow Control 23
Profile and Class Selection of Service Specification 23
Security 24

Heterogeneity 25
Quality of Service 25
Group and Multicast Support 30
Enhanced Transport Services 32
Recommendation 35
5 Integrated QOS for Multimedia Communications 37
Introduction 37
A New QOS Architecture for Distributed Multimedia Communications 37
The Need for a QOS-A 37
QOS Requirement for Continuous Media 37
QOS-A Requirements 37
Current Notions of QOS 37
OSI Perspective 37
CCITT Perspective 37
CCITT I-Series Perspective 37
Evaluation 38
An Extended View of QOS 39
Baseline Architecture 40
QOS Management Functions 41
Monitoring and Policing 41
Layer Specific QOS Considerations 42
6 Conclusion 43
Appendix 46
A: A List of Requirements 46
List of Baseline Requirements 46
TPX Service Requirements 49
Non TPX Service Requirements 52
B: Abbreviations 54
References 55

1. Introduction

The evolution of distributed computing is being simultaneously influenced by the new multimedia applications requirements [Leopold,92a] and by the increasing capability of high-speed multiservice networks [Boerjan,91]. The performance of the communications system will be critical in meeting the needs of the emerging distributed multimedia applications; in particular new enhanced transport facilities and functions should offer comprehensive *Quality of Service (QOS)* configurability for a wide variety of media types and specific user requirements. It is generally accepted that existing transport protocols such as TCP and TP4, designed for unreliable lower speed networks, will be unable to best exploit the potential bandwidth offered by ATM networks and meet the high performance requirements of new applications [Danthine,92a]. Therefore, it is expected that new high performance transports, such as TPX[1] [Danthine,92b][Baguette,92a] will provide enhanced functions [Baguette,92b] and facilities [Leonard,92] to operate efficiently in the new environment [Boerjan,91] [Greaves,92]. The primary objective of such a transport service is to provide a high performance interface which can both guarantee QOS to a wide variety of distributed multimedia applications and exploit the underlying multiservice networks.

Integration of distributed multimedia systems support into a communications architecture encompassing the new multiservice networks poses significant challenges. A key observation about the new environment is that QOS provides a unifying theme around which most of the new communications requirements can be grouped [Leopold,92b]. For applications relying on the transfer of multimedia information, and in particular continuous media, it is essential that QOS is guaranteed system-wide, including the distributed system platform, the transport protocol and the multiservice network. Enhanced protocol support such as end-to-end QOS negotiation, re-negotiation, indication of QOS degradation and co-ordination over multiple related connections are also required. However, little attention has so far been paid to the definition of a coherent framework that incorporates QOS interfaces, management and mechanisms across all the layers. An important first step in meeting such a requirement is the specification of a *Quality of Service Architecture* (QOS-A) which provides a framework to specify and implement the required performance properties of new multimedia applications [Blair,91] over enhanced communication functions and facilities such as TPX [Danthine,92b] and multiservice ATM-based networks [Boerjan,91].

1.1 Scope

The aim of this report is to study to what extent OSI 95 Transport Services (TS) support distributed multimedia applications, and at the same time preserve the capabilities offered by B-ISDN based on ATM technology. We build on the results of BELL-1[Boerjan,91] and ELIN-1[Leopold,92a] by evaluating OSI 95 transport service specification [Leonard,92] as the critical component of an integrated QOS-A for the new environment.

This report is structured as follows. We first presents, in section two, a summary of the communication requirements of new distributed multimedia applications. This is followed, in section three, by an overview of the enhanced OSI 95 transport service; this provides a suitable background for the evaluation of OSI 95 transport services in section four. In this section we take the novel approach of a quantitative requirements analysis of OSI 95 transport against the communication service requirements; each requirement is given a number, these are listed in appendix A and enumerated in section two. In section four we evaluate the OSI 95 transport service against these key transport service requirements, this results in quantitative conclusion to transport suitability. Finally in section five we introduces the notion of an integrated QOS architecture for the new environment.

---

[1] The enhanced transport service specification ULg-4 [Leonard,92] and the TPX specification ULg-5 [Baguette,92b] focus on one of the five classes of TS initially described in [Baguette,92a] We refer to all five classes collectively as “the OSI 95 Transport Service” (OSI 95 TS) and the Peer-to-Peer Enhanced Connection-mode Transport Service (PP-ECOTS) as the Enhanced Transport Service (ETS). Broadly speaking OSI 95 TS refers to the enhanced transport facilities and TPX to the underlying transport functions.

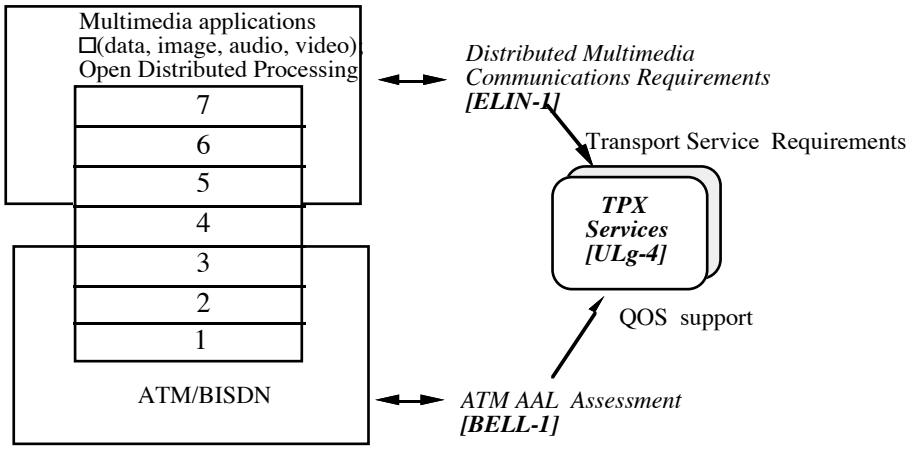

Figure 1: Evaluation of the OSI 95 Transport Service

As Figure 1 indicates, the focus of our evaluation will be the OSI 95 Transport Service (OSI 95 TS) in the new environment.

2. Distributed Multimedia Communication System Requirements

2.1 Introduction

This section presents a summary of the major finding of the ELIN1 deliverable[1] on distributed multimedia communication system requirements. The new requirements are categorised as *(i)* general multimedia communication issues; *(ii)* group and multicast support; *(iii)* QOS; and *(iv)* orchestration. In addition, this section also identifies the types of enhanced transport services which emerged during the analysis of distributed multimedia applications [Williams,91].

This summary includes a comprehensive list of the baseline communication requirements. We adopted a convention [Rn] to number each requirements in this section. For example the [R001] references an error control baseline requirement in section 2.2.1.3 :-

"One function of a transport protocol is to provide the required end-to-end error rate desired by the application[R001]"

By adopting such an approach, we attempt to provide the basis for a more quantitative requirement analysis of OSI 95 TS in section four.

2.1.1. Caveat

Many of our requirements are very high level and broad in nature and other are much more specific. This reflects the fact that a summary is presented, and not the detailed requirement specification [Leopold,92a]. Furthermore, some requirements have gained wider acceptance in the field, and in this case we have detailed the requirement in the preceding text. Because some of requirements have not mature as yet, we felt a broader description of the issue would suffice.

2.1.2 General Multimedia Communications Requirements

2.1.2.2. Continuous Media

Multimedia is characterised by continuous media such as voice, video, high quality audio, and graphical animation, which place much greater demands on communications than still media, such as, text, images and non animated graphics. The former consist of continuous

[1]Results from ELIN-1 [Campbell,92b] were introduced into the ISO as a contribution to the New Work Item on Enhanced Communications Functions and Facilities (ECFF); our work was used as primary input for the guidelines document [ISO,92d].

sequences of samples that have strict timing dependencies [Nicolous,91]. These are often referred to as *streams* [Coulson,90]. Continuous media requirements are quite demanding, for example, minimum latency, guaranteed bandwidth, real-time transmission, bounded jitter, QOS guarantees, *orchestration* [Berrocal,92] within single streams and between multiple streams, media and application specific protocols and error recovery strategies etc. These various requirements are discussed in greater detail below.

2.1.2.3. Error Control

Any network will produce errors and there are two types of expected errors in high-speed networks: bit corruption due to noise and packet loss due to congestion. One function of a transport protocol is to provide the required end-to-end error rate desired by the application [R001]. Correct data delivery encompasses two issues: *(i)* correctness in the value domain; and *(ii)* correctness in the time domain. For the correctness in the value domain, we have two approaches to error control:

> *Natural Redundancy in Multimedia:]*High reliability is not always of pressing importance since media which includes natural redundancy such as voice and video are inherently tolerant of small amounts of error. Thus, instead of high reliability, error announcement (i.e. error detection and indication) at the receiver [R002] side gives the user the possibility to determine the grade of error rate and degradation in the requested service.
>
> *Error Sensitive Information:* In contrast to the requirement of information types with natural redundancy, information which is highly error sensitive, such as file transfer of control data, X-Ray pictures, etc., needs an error control mechanism for higher reliability [R003].

2.1.2.4. Timeliness

For the correctness in the time domain, the communication system should support timeliness (i.e. bounded time limits) of media delivery [R004]. This can be achieved by providing a validity time to each message with an associated deadline. Thus, a global time, or a common reference time that provides system-wide time synchronisation is a basic issue in distributed real-time systems [R005].

Distributed real-time application must be informed when the communication system fails to transmit a message within its agreed time frame [R006]. Notification of such a failure allows the application to perform exception handling if necessary, we describe this as bounded invocation.

Taking into account the timeliness requirement we get a third information type, which is sensitive to errors in both domains, the value and time domain.

> *Real-Time Error Sensitive Information:* Real-time applications which are sensitive to errors must transmit redundant information (i.e. Forward Error Correction (FEC)) in order to recover from loss without retransmission [R007], which may cause unacceptable delays.

2.1.2.5. Error Recovery and Application Requirements

To allow distributed multimedia applications to make a trade off between reliability and timeliness on the one hand, and to give to the application the possibility to perform error recovery at the application level on the other hand, the new transport mechanism must support flexible error control [R008]. For example, video can sustain a certain degradation in quality by simply redisplaying the previous frame, that is, continuity, which would be meaningless for speech. With regard to applications, some may prefer or tolerate slower delivery rates in order to improve quality, thus it is important that applications have a means of controlling and selecting error recovery strategies as *(i)* go-back-N[R009]; *(ii)* selective retransmission[R010]; *(iii)* FEC[R011]; and *(iv)* no error correction [R012]. Furthermore the transport service should allow flexible error detection and recovery mechanisms[R013], for example the transport

service user should be able to select just error indication or possibly only error correction.

2.1.2.6. End-to-End Flow Control

The most critical requirement of a high speed transport is to sustain isochronous transmission and reception rates. To do this the transport protocol must be designed to support continuous transmission of data [R014]. This requires an open window flow control mechanism where the sender does not wait to send (rate-based flow control)[R015]. End-to-end flow control techniques have a major impact on the effective throughput of the transport protocol. The receiver must process the received data stream in real-time and provide the buffering implicitly agreed upon during connection establishment[R016]

2.1.2.7. End-to-End Argument: End-System Responsibilities

In multimedia systems it is not sufficient to guarantee the transmission of information between two endpoints. It is also important that the operating system provides appropriate response[R017]. Hence, QOS issues such as throughput, latency and real-time responsiveness are end-system concerns as well as concerns of the communications infrastructure. This makes the relationship between, and integration of, the operating system and the communications infrastructure particularly important in distributed multimedia systems.

2.1.2.8. Profile and Class of Service Selection

The transport service and protocol should be flexible enough, to support the full range of multimedia applications and the wide variety of communication characteristics, and work over different high performance networks[R018]. The idea of protocol selection [R019](i.e. profile selection) is well established within the OSI RM, and therefore will also be supported in our model. Although we advocate the use of QOS parameterisation of connections [R020](see section on QOS), we do not envisage a single fully generic transport protocol that can cater for all types of traffic equally well. Instead, due to the diverse nature of distributed multimedia applications, different protocols will be required for different types of traffic and control data. We thus envisage that communication sub-systems of the future will be composed of horizontal and vertical subdivisions in a *protocol matrix*. The user will be required to select a protocol profile that will provide a suitable service for each traffic type without duplication of functionality at the various protocol layers. Class of service selection will implicitly select between reliable and unreliable communications, and thus use various error control mechanism. In addition, to the provision of profile selection, we see the need for an extension of the more traditional notion of class of service selection employed in current OSI standards. In the current OSI framework, classes of transport service are selected according to the network service over which connections are to operate; that is whether the network is connectionless, connection oriented, reliable, unreliable etc. For future multimedia applications the notion of class of service selection should be extended to encompass more user oriented functions [R021].

2.1.2.9. Security and Heterogeneity

There remain other issues such as heterogeneity [R022] and security [R023] which were considered outside of the scope of ELIN-1. However, these issues are an integral part of the new protocol environment within modern high performance networks.

2.1.2.10. Implementation Philosophy

It has been suggested [Chesson,89] that the protocol overhead is less than 20% of all processing time in communications and that most of the overhead is expended on supporting timers and copying data. The reduction in the copying can be achieved by adopting a high performance buffer architecture which utilises a shared buffer mechanism between the ULA and the lower levels [R024]. This may contravene the OSI style of data encapsulation but is seen to be a necessary component of a transport service built specifically to exploit high throughput for continuous media traffic.

2.2.1.3 Quality of Service

Quality of Service parameters [R025] in current communications infrastructures permit the specification of user requirements which may or may not be supported by underlying networks. Usually, these values are agreed between the service provider and the customer at the time the customer subscribes to a particular network service. Another function of these parameters is to form a basis for charging customers for pre-specified services.

With the emergence of digital video, audio, and other continuous media types, increased requirements are placed upon QOS support. For example, to support video connections high throughput is required and therefore high bandwidth guarantees will have to be made. Audio, on the other hand, will not require such a high bandwidth.

End-to-end delay and delay jitter (i.e. variance in delay) are new factors which must be taken into account for continuous media transfers. In particular, interactive continuous media connections impose stringent delay constraints, derived from human perceptual thresholds, which must not be violated. Delay jitter must also be kept within rigours bounds to preserve the intelligibility of audio and voice information.

It is therefore important to support a notion of a Quality of Service Architecture [Campbell,92c], whereby application requirements can be easily mapped down onto the communications sub-system, that is, high level requirements can be mapped onto communication channels and resources. This mapping should be clean, simple and efficient, protecting application programmers from communication details, that is, they should be stating what they require rather than how it is to be achieved.

Particular applications may require certain minimum standards, below which transmission may as well terminate, on the other hand other applications may require only an indication that the QOS has degraded [R026], in this case the application is responsible for taking any action regarding the connection. It is up to the communications sub-system to guarantee that the required QOS is maintained throughout transmission [R027]. Users must have some means of defining the QOS they require at an appropriate level, these may include specific quality statements, don't care, best fit options or guaranteed [R028].

2.1.3.1. Transport level QOS Parameters

A set of suitable QOS parameters [Leopold,92a] which are meaningful to the transport level and the levels below, and which may be employed in characterising individual continuous media types, are *throughput* [R029], *end-to-end delay* [R030], *jitter* [R031], *error rate* [R032] and *priority* [R033]

At connection establishment time it should be possible to quantify and express preferred, acceptable and unacceptable tolerance levels [R034] for each of these parameters. The requested parameters should then undergo full end-to-end negotiation [R035]. The finally agreed QOS should then be guaranteed for the duration of the connection[R036] or at least a QOS degradation indication [R037] should be provided if the contracted values are violated. The background to these parameters is as follows:

Throughput and End-to-End Delay

Multimedia applications require high bandwidth (from 12 Kbytes for a page of text to 150 Mbits/s for HDTV) and high performance delivery mechanisms with bounded end-to-end latency. These bandwidth requirements may be reduced with the use of appropriate compression techniques [R038] at the expense of latency. Based on these figures, new multiservice networks, will be able to handle the transmission of raw data. This basic scenario seems to be acceptable until one considers real-time requirements, multiple transmissions, orchestration and the need for a guaranteed QOS, all of which impose significant demands on the communication stack. Also, late data in continuous media streams is as serious a problem as lost data, and therefore the new service must attempt to provide the required level of error control [R39]without resorting to retransmission, due to the timeliness requirement.

Jitter

There is never complete synchronisation between a sender and a receiver since variable delay (i.e. jitter) is introduced by a real network. Therefore, the sender and receiver may drift in and out of synchronisation. Continuous media streams require that jitter delay be kept within a small upper and lower bound. Jitter introduced by queuing in intermediate nodes may be reduced by the use of buffering at the receiver [R040]. The amount of jitter allowed must be traded off against the end-to-end delay requirements and the receiver buffer availability. It is important to note that the jitter detection may is done within the network [R041] and jitter correction at the receiver [R042].

Priority

The communication system should support multiple levels of priority for multimedia traffic and the pre-emption of network resources for high priority applications [R043]. The transport service must be flexible enough to allow the transport service user to attach relative priorities across end system connections [R044]. It is feasible that this could be supported during connection set-up using an associated QOS parameter. High priority connections will get serviced before lower ones. Furthermore, it should be possible that lower priority connection packets will be dropped first before high priority packets [R045], should the network become congested, which is probable. Real-time systems require guaranteed delivery of high priority control messages such as a control message[R046] even under peak-load condition.

Error Rates

Packet error rates are dependent upon the size of the Transport Protocol Data Unit (TDPU) at the transport layer and in some way the Transport Service Data Unit (TSDU) at the Transport Services Access Point (TSAP). Bit error rates are associated with reliability of the network. The user should just have to deal with packet errors rate at the TSAP [R047]. Further research is necessary to work out an acceptable specification from the packet error rates.

The following section summarises additional QOS issues required to meet the needs of the new distributed multimedia applications.

2.1.3.2. Guaranteed QOS

Multimedia applications which can accurately forecast their resource requirements may prefer to reserve transport and network resources in advance via a connection-mode service [R048]. Transactions [R049] which require fast response often cannot wait for the connection time associated with resource allocation which is greater than twice the round trip delay. Applications that require any QOS higher than best-effort [R050] (offered by current networks) must request resources be allocated at the end system and in the network.

End-to-end Resource Allocation and Reservation

An important requirement of continuous media QOS management at the end systems and network level is a suitable reservation scheme [R051] which is able to set up and guarantee network resources such as bandwidth and end-to-end delay for high performance continuous media communications between network subscribers. This could be achieved using resource reservation protocols based on such as SRP (Session Reservation Protocol) [Anderson,91], ST-II (Stream Protocol) [Topolcic,90] or directly provided by the network services, as proposed in ATM networks with signalling[1] protocol using Q93b.

Currently, the literature discusses two methods for resource allocation and reservation in distributed communication infrastructures. In the first approach, the transport service and protocol takes the responsibility for the reservation of communication resources. This includes reservation of communication system resources, such as transport and network resources, but excludes other end system resources, such as application and operating system resources. An alternative to this approach [Campbell,92d], and perhaps more suitable for the new

[1]Appendix D describes QOS in ATM-based B-ISDN, this includes Q93b and the BAP proposal.

environment, is to consider resource management as orthogonal issue to the transport service and protocol. Ideally the applications should be able to reserve end systems and communications resources in advance of data transfer or at the time of data transfer. In addition, the reservation should be truly end-to-end in nature, including all resources associated with distributed communications. Decoupling resources management and communications in this way would require the provision of a new reservation mechanism. We propose that such a QOS reservation protocol could be responsible for interacting with resource managers at the end systems and in the network. The division between resource management and data transfer management would result in the realisation of an integrated end-to-end resource reservation model which could support diverse QOS needs of distributed applications.

While peak bandwidth network resource reservation strategies lead to lower utilisation of network bandwidth, guaranteeing continuous media QOS is more important than maximising utilisation of the link bandwidth.

2.1.3.3. Dynamic QOS Control

In distributed multimedia systems it is not sufficient to specify a QOS level, protocol profile and service class at connection time which will statically remain in force for the life-time of the connection. A more flexible service is needed whereby users can dynamically alter the QOS of a connection [R052] while it is active. This is required because there are applications such as continuous media which have extremely diverse QOS demands which may vary during the course of a connection. Such dynamicity is illustrated in the following examples:

> If *best effort* guarantee of QOS level is selected at connection time, the QOS level may degrade, and this will result in the transport user being informed of the degradation (via a QOS degradation indication [R053]). Rather than simply closing the connection down or accepting the degradation, the user may want to re-assess his priorities and, for example, close down other connections to save resources and then upgrade the first connection or retain the connection with the currently degraded QOS.
>
> The user may use the same connection successively for different purposes at different times. Examples are the transmission of full motion video interspersed with intervals of slow motion, or upgrading from monochrome to colour video, or telephone quality to CD quality audio. Another possibility is the in-service insertion of a compression module which may reduce the bandwidth requirement but increase the susceptibility to errors.

Of course, to permit QOS negotiation and re-negotiation [R054] at the TSAP, the necessary support must be provided either within or below the transport protocol itself. Thus mechanisms are required to alter link-level bandwidths and/or processing and buffering resources at intermediate nodes. One possibility is the use of re-configuration at the network level in the context of protocols such as the real-time channel administration protocol [Banerjea,91] where resource reservations are made at intermediate nodes. Even when an underlying connection must be torn down and re-established in order to achieve an alteration in QOS level, there are strong arguments for doing this transparently behind the transport service interface. For example, it allows the maintenance of buffers and protocol state over the successive connections which may minimise the delay before data flow may resume; it eases the management burden on the transport service user.

2.1.4 Group and Multicast Support

The applications study [Williams,91] highlights the importance of group [R055] and multicast [R056] services as characteristics of distributed multimedia communications. For example, many CSCW applications require sophisticated patterns of connectivity between different sites. Note that multicast is important for hard real time data transfer [R057], bulk data transfer [R058], transaction [R059] and continuous media services [R060]. Consequently, the underlying transport mechanism [R061] and the network provider [R062] must support group connection management [R063] and multicast addressing respectively [R064]. Connection management must be flexible enough to support a wide range of data transfer, from a simplex

connection [R066] between a camera and a remote display to more complex configurations that support multiple remote sources connected to a common local or remote sink [R067].

2.1.4.1. The Concept of Groups

A group is a concept to form an association between a set of users, peers or object entities. Grouping is motivated by the need to *(i)* encapsulate the internal state and hide interactions among the group members; *(ii)* provide fault tolerance by means of functional redundancy; *(iii)* reduce network and sender resources by delivering a single message to multiple destinations without the overhead of multiple discrete addressing; *(iv)* provide a high level communication interface that simplifies applications software complexity for simultaneously communications with multiple receivers.

2.1.4.2. Multicast Services

Multicast is the simultaneous transmission of a single message to multiple destinations (i.e. all destinations will receive the same message). A multicast service must provide the same QOS [R068] associated as with the more traditional point-to-point network access services, in terms of selection of end-to-end delay, jitter, throughput and reliability. A multicast service should be configurable to the communication need of distributed multimedia applications; that encompasses:

- *communication topology: (i)* multicast (one-to-many) communications [R069]; *(ii)* broadcast (one- to-all) communications [R070]; *(iii)* multi-collect (many-to- one) communications [R071]; and finally *(vi)* conference [R072] (typically many-to-many in video conferencing configuration) communications.

- *communication semantics:* That means either the sender will not have any knowledge about the number of recipients that actually received the message (un-acknowledged mode [R073]); or each message is acknowledged by each recipient (acknowledged mode [R074]); or the receivers retain the delivery of the message to the user, until a notification arrives from the sender [R075]

- *collection and distribution functionality:* is important to have a form of *naming transparency* i.e. group names [R076] to address the receivers on the one side, and to process the replies properly on the other side (reply-handling transparency).

- *collection and distribution policy:* It should be possible to specify any set of receivers or senders and the timing relation of the communication [R077]. Examples are at least, at most, exactly; one, some, all; etc.

2.1.5 Orchestration

The development of communications systems to support distributed multimedia application which exploit continuous media introduces new synchronisation requirements [Berrocal,92]. In particular, the upper layers require support for the following two varieties of synchronisation:

- *event-driven synchronisation* [R078] This is the act of notifying that a relevant event or set of events has taken place, and then causing an associated real-time action or actions to take place. For example, a user clicking on the stop button relating to a video play-out could cause the play-out to stop instantaneously.

- *continuous synchronisation* [R079] This is where two or more continuous streams need to be kept in step by an on going commitment to a repetitive pattern of event driven synchronisation relationships. An example scenario where it is required to form a continuous synchronisation relationship is a language laboratory where separate audio tracks in different languages are stored on a single server but are to be distributed to different workstations in a real-time interactive language lesson. Another example arises where it is required to associate captions from a text file with an on-going video play-out. In both of these scenarios there is a need to maintain a tight on-going linkage between real-time continuous media information streams.

We refer to these sorts of linkage as *orchestration*, and have developed an architecture to model and realise such relationships[Campbell,92a]. We employ the term orchestration rather than *synchronisation* for two reasons. Firstly, cross-stream relationships may encompass more than just temporal co-ordination, and secondly the term synchronisation is already overloaded and given a different emphasis in current OSI usage (the OSI concept pertains to checkpointing and state synchronisation between peer entities).

2.1.5.1. Orchestration Functionality

Focusing on the continuous synchronisation problem, the following functionality must be provided by the infrastructure to ensure that media flowing in separate (non-multiplexed) [R080]but related connections can be correctly co-ordinated in time:-

- the ability to start [R081]related continuous media data flows precisely together. If the relationship is not correctly initiated, there is no possibility of maintaining a correct temporal relationship.

- the ability to start and stop the flow of continuous media information on sets of related connections together in real-time.

- the ability to create [R082] related connections with compatible QOS. This is so that the connections will maintain a compatible temporal transmission rate in the required ratio.

- the ability to monitor [R083] the on-going temporal relationship between related connections, and to regulate the connections to perform fine grained corrections if synchronisation is being lost. It is almost inevitable that related connections will eventually drift out of synchronisation due to such factors as the potentially long duration of continuous media connections in typical applications, the inevitable discrepancies between remote clock rates, and temporary "glitches" occurring in individual connections.

As shown in Figure 2, we attempt to place the orchestration services in perspective by illustrating their use in a microscope application scenario, developed at Lancaster University [Campbell,92c].

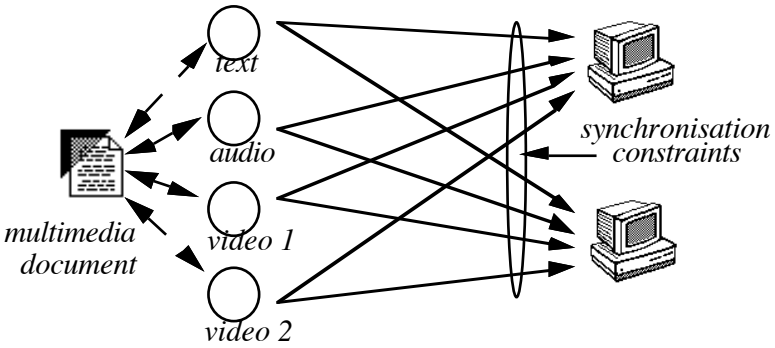

Figure 2: Orchestration Requirement .

Consider an application configuration where a multimedia document needs to be played back to two separately located scientists who are conferring over a video telephone. The document contains two video clips which view the same sample at different levels of magnification, together with occasional voice annotation and text captions. The various components of the document are separately stored but must nevertheless be synchronised on playout both together and across the two scientist's workstations. The details of the required synchronisation (i.e. timing of text captions, starting time of annotation, degree of permitted slack etc.) are contained in a 'script' encapsulated within the multimedia document structure. Figure 2 also illustrates the logical topology associated with this configuration.

2.1.6 Enhanced Transport Services

Traditional transport services, such as OSI Transport Class 4, have mainly serviced the communication needs of file transfer and low-bandwidth interactive applications. It is generally accepted that these protocols are not suitable for multimedia and real-time communications, and that dedicated transport services and supporting protocols are needed to exploit the performance of the new high speed network environment such as the ATM based multiservice networks. The following enhanced transfer services have been identified: *(i)* continuous media service[R084]; *(ii)* transaction oriented service [R085]; and *(iii)* bulk data transfer service[R086]. These transport services should operate over connection-oriented [R087] and connectionless services[R088] and the entire communication topology of group communication. The error control mechanisms must provide both assured [R089] and unacknowledged [R090] capabilities. The following section describes each of the three types of transport services mentioned above.

2.1.6.1. Continuous Media Services

Continuous media is the essence of multimedia and has by far the most stringent and novel requirements in terms of guaranteed QOS and orchestration. Continuous media traffic differs from traffic generated by traditional distributed applications. Instead of the bursty traffic generated by, e.g., remote procedure calls, continuous media packets are transmitted by an application at a fixed rate (e.g. a video frame every 33 ms) and received at a fixed rate. This kind of traffic is often referred to as isochronous and is characterised by its strict adherence to periodicity, in contrast to the asynchronous nature of bulk data transfer. Continuous media transfer requires new transport mechanisms that support the following functionality [R091].

- high bandwidth [R092];
- a flexible approach to error recovery for continuous media [R093];
- end-to-end rate based flow control mechanisms [R094];
- mapping application requirements onto the communications sub-system [R095];
- QOS degradation indication [R096];
- QOS negotiation and re-negotiation [R097];
- transport and network resource allocation to offer guaranteed end-to-end QOS [R098];
- single and multiple streams[R099];
- orchestration [R099];
- group/multicast [R100];
- isochronous delivery mechanisms[R102];
- non multiplexed simplex connections [R103].

All the above requirements need to be guaranteed on an *end-to-end basis,* whether it be from kernel to kernel, application to application or application to remote device etc. Additionally, the above continuous media requirements indicate a definite *need for flexibility* from the selection of protocols, QOS, synchronisation and error recovery etc. Flexibility should be achieved without incurring excessive performance and complexity overheads.

2.1.6.2. Transaction Oriented Services

Transactions occur in distributed systems such as distributed operating systems, databases or reservation systems. Although a *Request/Response Communication* (a client issues a request, a server responds to that request) is the basis for a transaction, it is important to note there are a lot of differences. The term transaction will imply the characteristics of *failure atomicity, failure isolation, permanence and serialisability*. The Transaction paradigm is an application oriented term and should not be misunderstood as a communication method for

lower layers of the OSI model. A similar communication issue is the Remote Procedure Call (RPC). [Dark,91] considers how the RPC, a widely used and accepted abstraction for distributed programming, may be utilised to support distributed multimedia computing. RPC is based on the premise that distributed programming will be more readily adopted if inter-machine communications can be handled in a familiar and simple way. The procedure call is the primary means of invoking and transferring control and data between local procedures. In essence the RPC simply extends the procedure call concept to allow the control and transfer of data across communication networks. Rather than sending and receiving messages a program calls a remote operation and receives return values. This hides distribution complexities from the programmer and provides a familiar interface, that is, the procedure call. In an ideal world one would not be able to distinguish between remote and local procedure calls, distribution would be totally transparent.

The following section describes some transactional oriented issues.

• *Reliable request/response* [R104] *interaction versus unreliable* [R105] *messaging communication:* In the case of a reliable request/response service both the request and the response have to be delivered reliably. Connection oriented mode, which supports reliable communications, requires a connection set-up which adds additional overhead to the interaction. This overhead could be too high in relation to the duration of the whole interaction. If this is the case then connectionless services or fast connection set-up and release could be more appropriate to real-time multimedia operations.

• *Maintaining Connections with Long Lifetimes*: [R106] If connections are used to support a reliable kind of communication, the duration requirements for connections may vary from milliseconds to long periods of time. The communication system should maintain connections that operate over long periods of time by polling the connection over its life-time. In addition, network resources required for long term connections should be managed efficiently; this is dependent on the underlying network provider but could include that dynamic allocation and deallocation of resources or the opening and closing of a connections when the communicating entities are active or not.

2.1.6.3. Bulk Data Transfer

Bulk data transfer is a service that carries a large amount of data with relatively loose latency constraints. Generally, all data must be reliably received before the receiving application can proceed. Bulk data transfer should be efficient over large bandwidth-delay product channels. As transmission must be reliable, retransmission of corrupted data is required [R107]. High throughput is the major performance issue associated with the transfer of bulk data between network subscribers. This required throughput is often limited by the transport protocol, particularly due its flow control mechanism.

2.1.7 Baseline Requirements

In total, 107 *baseline* communication system requirements are identified, [R001] through [R107]. The baseline requirements broadly related to service and protocol issues of the ULA and the transport sub-system. Because our evaluation concentrates on the transport service access point, we partition the baseline requirements into *OSI 95 transport service* and *non OSI 95 transport service* requirements; restricting the analysis of TPX to service requirements only. A summary of the baseline requirements, OSI 95 transport service requirements and OSI 95 transport non-service requirements can be found in appendix A2 for quick reference.

3. Summary of the OSI 95 Transport Service

3.1 Introduction

In this section we provide a summary of [Baguette,92a] [Baguette,92b] [Danthine,92b] [Leonard,92] the OSI 95 Transport Service. Before describing the services and their respective time sequence diagrams, we outline the rationale behind the new transport service, by briefly describing some new features which underpin the TPX philosophy [Danthine,92a] [Baguette,92a]

3.1.1 OSI 95 Transport Service Features

The design objective of new OSI 95 transport services, specified by the University of Liege, was to take into account *(i)* existing connection-mode [ISO 8072,86] and connectionless-mode [ISO 8072 ADD1,86] services; *(ii)* XTP [PE, 92] ; *(iii)* new network performance and services [Boerjan,91]; and finally *(iv)* the new distributed multimedia communications requirements [Leopold,92a].

The new OSI 95 TS reflect the new requirements for both high performance communications and enhanced QOS support, these new transport service features include:-

- flexible error control for transport connections;
- additional QOS parameters;
- a new QOS semantic: compulsory QOS values;
- full end-to-end QOS negotiation and QOS re-negotiation;
- support for continuous media communications;
- user indications for QOS degradation;
- group and multicast support;
- an out-of-band data transfer facility;
- and graceful release mechanism for transport connections.

Most of the new features are associated with the ETS [Leonard,92]. In addition, the literature [Baguette, 92a] states that these features are for further study with respect to all other OSI 95 transport services. In section three each service is described and open issues are identified. ETS is covered in more detail in [Leonard,92] and therefore we focus our assessment on this service. The remaining services are still under consideration with important issues still for further study. The following section addresses each of the new features in more detail.

3.1.1.1. Flexible Error Control

The OSI 95 TS will provide a selectable error recovery [Danthine,92b], depending on the application. For example data transfer and continuous media transfer have very different error control requirements. For the connection-mode transport service, TPX will permit tuning (which may lead to suppression) of the error recovery mechanisms. Multimedia applications may require error detection without error recovery, take for example a live video stream transfer, where the TS user only requires notification of an error in a TSDU. In this case TPX will deliver an indication primitive with no video data, with corrupted video data or with dummy video data replacing the corrupted or missing video data. The OSI 95 TS will supply a flexible error management interface allowing the TS user or acceptor to select the appropriate error control mechanisms.

In [Leonard,92] error control is tightly bound to the residual TSDU error rate and is defined as the maximal number of errors given that may occur in a window. In addition [Danthine,92b] states that the transport connection (TC) error control can be independently selected and negotiated for each direction of the connection. The TC error control is encapsulated in a error policy and divides into the corrupted TSDU and lost TSDU error rates. Because the TS user may require to be informed if an error exception occurs, a guarded error indication selection is proposed [Baguette,92a].

3.1.1.2. The OSI 95 Transport Service QOS Parameters

TPX will provide the following new QOS parameters

Performance QOS parameters

- throughput
- traffic type (isochronous or non-isochronous);
- transit delay;
- delay jitter;

Information transfer QOS parameters

- protection;
- priority;
- TSDU_Size;
- error_control_selection;

3.1.1.3. Compulsory QOS

Because multimedia communications require hard guarantees for throughput, delay and jitter, TPX takes a new approach towards QOS support called *compulsory QOS*. Today most transport services only support *best effort* communications, this is because most networks can only operate at this level of services. With the advent of new multiservice networks, such as ATM-based networks, it will be feasible to offer a wide variety services and most importantly, better service with hard guarantees. This will allow transport providers, such as TPX, to guarantee compulsory levels of QOS parameters. Best effort also means a qualitative semantic without any actual commitment to a given quantitative value for say throughput or delay. Evidently this type of best effort QOS model is not suitable for the new environment and QOS requirements of distributed multimedia applications such as video conferencing or hard real time applications.

Compulsory QOS means that the TS provider will monitor any parameter which falls in this category, if the level of service of this parameter falls below a contracted value then it is up to the TS provider to release the TC. This would happen in the case where a compulsory QOS could not be maintained by the provider.

TPX allows the TC user to supply both a minimum and maximum value for compulsory QOS parameters. In addition TPX presents the TS user with a flexible QOS parameter interface; users can mix compulsory and best effort parameters on a single transport connection. For example, a TS user may need hard guarantees (compulsory) on throughput and have less stringent requirements (best effort) on jitter. Section five builds on this new QOS perspective by defining a *Quality of Service Architecture (QOS-A)* for multimedia communications.

3.1.1.4. QOS Negotiation

TPX defines three *performance QOS parameters (i)* throughput; *(ii)* transit delay; and*(iii)* and delay jitter which can be negotiated and re-negotiated separately for both directions of transmission on a TC. Other parameters, such as the *information QOS parameter* (that defines the TSDU size) and the *traffic type* component (isochronous traffic or not) of the throughput parameter can not be negotiated but are fixed by the calling TS user. Only the called TS user can modify the proposed compulsory values of performance parameters at connection time. However, the TS provider may make some limited changes, such as to decrease the maximum throughput (but not below the minimum) and increase the minimum (but not above the maximum) throughput but within bounds.

The called TS user may modify the performance QOS parameters on reception of a T-CONNECT.indication informing the calling TS user in the T-CONNECT.confirm. Again the called TS user can decrease the maximum throughput and increase the minimum but within the bounds of the initial call. By restricting the allowed modifications as above, the compulsory character of the QOS proposed by the calling TS user is only strengthened by the negotiation.

3.1.1.5. QOS Re-Negotiation

TPX takes on board recommendations made in [Leopold,92a] for the re-negotiation of QOS parameters during the lifetime of a transport connection between two applications. In [Baguette,92a] the authors identify the type of the problems which have emerged in trying to meet such a new application requirement. As yet no solution has been chosen; however more work has been under taken in [Leonard,92] to overcome contention between the called and calling TS users.

3.1.1.6. QOS Degradation Indication

In addition to providing a flexible error control interface, the ETS supports QOS degradation indication when throughput degrades close to the contracted compulsory minimum value. This service is important for continuous media applications which may tolerate diminished throughput for short periods; alternatively the application may take some appropriate action such as dropping its bandwidth requirement to achieve a more stable service.

3.1.1.7. Continuous Media Support

Continuous media is the essence of multimedia and has by far the most stringent and novel requirements in terms of guaranteed QOS, support for isochronicity, flexible error control, sustained high throughput, jitter and multimedia synchronisation. The ETS provides an enhanced connection-mode service with extended QOS features for a continuous media communications. Important new ETS features, such as QOS semantics, flexible error control, isochronous support at the TSAP via the traffic classification parameter and QOS degradation indication, all go to meet a number of important novel requirements which collectively meet new continuous media communications needs.

3.1.1.8. Group and Multicast Support

TPX will provide an extended addressing capability, this will probably be multicast for data transfer. Future TPX service specification will take into account the work done in ISO JTC 1.21.09.01 on Multipeer Data Transmission (MPDT). [Danthine,92b] states that TPX multicast should be extended to cover all TPX services. The literature [Danthine,92b][Danthine,92b] also introduce the concept of Active Group Integrity (AGI) for the connection-mode service. This allows the TS user to request that a connection be set-up between at least one other destination, a number of destinations or a list of destinations. AGI provides QOS support for the connection duration, this is similar to the new QOS semantic promoted in the TPX specification. If AGI is can not meet a QOS requirement of the TS user then the service is not accepted or terminated. Members of a group can be added and removed from an AGI session. When a user leaves a multicast connection, the AGI should be re-evaluated.

3.1.1.9. Out-of-Band Data Transfer Facility

The concept of out-of-band data transfer appears to be rather close to the concept of expedited data transfer exhibited by ISO connection-mode transport service. [Baguette,92b] argues that TPX out-of-band data transfer facility is more suitable for the new environment, and therefore proposes that expedited data transfer should be suppressed.

The out-of-band data transfer will be used in conjunction with the ETS, and will be based on either the TPX request/response TS or the TPX acknowledged connectionless-mode TS. The motivation behind such a feature is described in [Danthine,92b][Baguette,92a]. Briefly, it is envisaged that an out-of-band data transfer capability will allow some TSDUs to be transmitted on the request/response TS or acknowledged connectionless-mode TS while maintaining a relationship between these TSDUs and the open transport connections. Whilst the out-of-band capability is seen as a key feature of TPX it is still not clear how the relationship of the out-of-band TSDUs will be associated with the connection-mode TS, [Baguette,92a] states that it is for further study.

3.1.1.10. Graceful release mechanism for transport connections

Graceful release of a transport connection, at the transport level, supports the orderly

release of the connection and guarantees that all data in transit is delivered before the connection is closed. This scheme is based on the orderly session release, defined in the current ISO Session Service standard [ISO 8326]

TPX advocates the migration of graceful release from the session layer to the transport layer. The arguments in favour of this are as follows:-

- the transport layer deals with most aspects of reliable data transfer other than reliable closure of a connection;
- the introduction of the graceful release is seen as providing the TS user with a new service, with the same motivations as providing the new TPX acknowledge connectionless-mode and request/response services. The authors are attempting to offer the community a richer set of services;
- because the transport layer is now responsible for graceful release, the upper layers are relieved of the burden of supporting this facility;

Graceful TC release requires a mutual agreement between all TS users of a connection. The request to release the connection can be initiated by either or both TS users and may be accepted or rejected by either one. By adopting graceful release TPX aligns itself with XTP and TCP on this issue. [Danthine,92b][Baguette,92a] states that the introduction of this feature in TPX would ease the migration of applications from the TCP world.

The following section described each of the new services in more detail.

3.1.2 OSI 95 Transport Services[1]

Briefly, the design objective of new OSI 95 TS, specified by the University of Liege, was to take into account *(i)* existing connection-mode [ISO 8072] and connectionless-mode [ISO 8072 ADD1] services; *(ii)* XTP [PEI 92] ; *(iii)* new network performance and services; and finally *(iv)* the new distributed multimedia communications requirements [ELIN1].

The five transport classes identified in [Danthine,92b][Baguette,92b][Leonard,92] and collectively described as the OSI 95 Transport Service are:-

- an Enhanced connection-mode Transport Service [Leonard,92] (ETS);
- a fast connection-mode service;
- basic (unacknowledged) connectionless mode service;
- an acknowledged connectionless-mode service
- and a request/response connectionless-mode service.

These services are described in more detail below.

3.1.2.1. The OSI95 Enhanced Connection-mode Transport Service

The ETS [Leonard,92] [Danthine,92b] is an enhancement of the existing ISO connection-mode service and will conform to the classical four service primitives for connection control This service will provide a full end to end QOS negotiation between peer transport entities before the data transfer phase can begin.

The proposed modifications to the ISO connection-mode service are as follows :-

- the introduction of a *graceful TC release facility* , which will guarantee that all data in transit on a connection is delivered correctly in the event of the connection being terminated. Data integrity of this nature is the responsibility of the session layer in the

[1]Currently the OSI 95 Transport Service only supports a peer-to-peer communications services. This is a limitiation of the existing CLNS; however the authors plan to add a multipeer service to to the existing peer-to-peer services in due course.

OSI RM. However, [Danthine,92b] argues that the transport-layer deals with most aspects relating to reliable data transfer, and the graceful connection release is concerned with reliability and therefore the transport layer is the natural place for this mechanism;

• the suppression of *expedited data transfer facility*, which is made redundant by the addition of an out of band data transfer capability. Out of band data transfer provides more flexibility and is perceived to be more suitable [Danthine,92b] than expedited data transfer capability in the new environment;

• to add *QOS support* , which will include *(i)* a new QOS semantic which will guarantee a compulsory level of service for a TC; *(ii)* monitoring QOS of a TC; flexible selection of error control; and *(iv)* the addition of a new jitter QOS parameter;

• the addition of a *multicast addressing* capability.

Time Sequence Diagram

As shown in Figure 3, the time sequence diagram for the ETS is built using T-Connect.

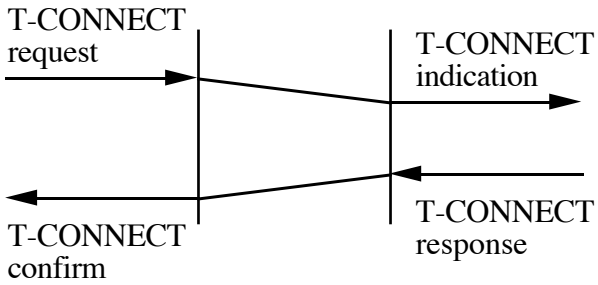

Figure 3: Successful TC establishment for the ETS

3.1.2.2. Fast Connection-mode Transport Service

High-speed networks, such as B-ISDN, exhibit large 'throughput * round-trip delay' products which are in the order of several Mbits. It has been argued [Danthine,92b][Baguette,92a] that the TPX connection-mode TS connection time can be prohibitive for some applications which can not afford to wait for at least one round-trip delay before data transfer can take place. Because some applications are sensitive to the connection time overhead, TPX offers a fast connect service.

The fast connect idea has been introduced in XTP for example. The basic idea, transposed from ISO 8072 TS, allows the calling TS user to submit T-DATA.requests before receiving the T-CONNECT.confirm. TPX introduces the idea of a T-LISTEN.request as a pending request for the incoming T-FastCONNECT.request; the T-LISTEN.request must correctly match the QOS requirements of the incoming T-FastCONNECT, this therefore implies a priori knowledge of the connections QOS characteristics.

Time Sequence Diagram

The time sequence diagram for the fast connect TS is shown in Figure 4. The called TSAP user must append T-LISTEN.request prior to the T-FastCONNECT.request being received at the called TSAP. If the addressing and QOS characteristics associated with the T-FastCONNECT.request are in line with the T-LISTEN.request then the fast connection is successfully established; if this is the case then the called TS provider responds by issuing a T-FastCONNECT.confirm. Once issued, the called TS provider is now ready to process T-DATA.indications. If the matching of characteristics fails then a T-DISCONNECT is issued by the called TS and any subsequent T-DATA.requests are discarded.

TPX fast connect provides low latency data transfer with limited QOS negotiation. This service is associated with the service concept of "take it or leave it". [Danthine,92b] recommends that this service would be suitable for connections which have well-known QOS

characteristics.

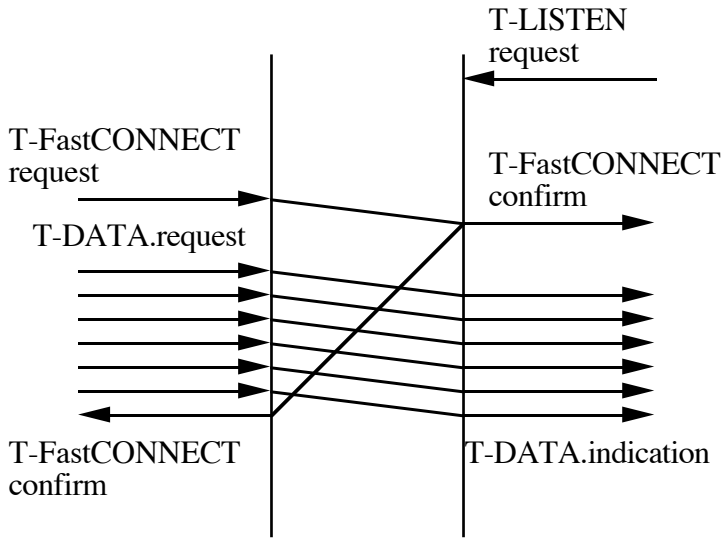

Figure 4 Time Sequence Diagram for the Fact Connect

3.1.2.3. Unacknowledged Connectionless-mode Transport Service

The TPX unacknowledged connectionless-mode TS is expected to be an enhancement of the current ISO transport service [ISO 8072 ADD1,86]. Figure 5 shows the normal time sequence diagram for the service. The impact of QOS and multicast addressing are for further study [Danthine,92b]

Time Sequence Diagram

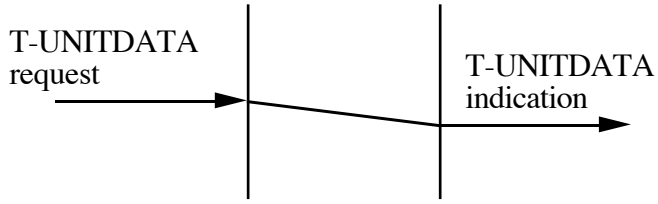

Figure 5: The Unacknowledged CL Transport Service

3.1.2.4. Acknowledged Connectionless-mode Transport Service

The unacknowledged connectionless-mode service is also an extension of the ISO connectionless-mode transport service [ISO 8072 AAD1,86]. The difference between this service and the unacknowledged connectionless-mode service is that the TS user will receive a positive or negative confirmation at the calling TSAP. The acknowledged connectionless-mode service may be suitable for applications which require the transfer of a single TSDU but can not afford the overhead associated with the connection time of the connection-mode transport service. In this case the application achieves a level of confidence that the TSDU was received or not; in the case of negative acknowledgement the TS user may choose to re-transmit the failed TSDU.

Time Sequence Diagram

Figure 6 shows normal time sequence diagram for this service. The impact of QOS and multicast addressing are for further study [Danthine,92b]

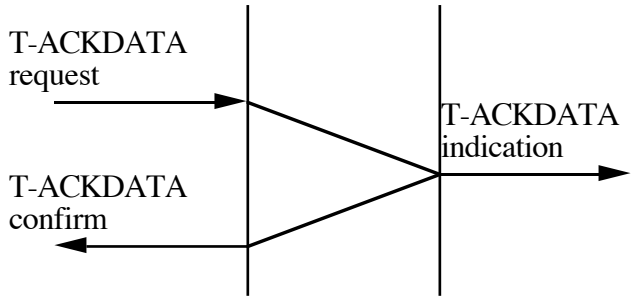

Figure 6: The Acknowledged CL Transport Service

3.1.2.5. Request/Response Transport Service

The TPX request/response connectionless-mode TS is a further extension of the acknowledged connectionless TS.

The motivation behind this services is described in [Baguette,92a] and is summarised as :-

- client/server applications may find it useful a request/response connectionless-mode TS at their disposal;
- the request/response connectionless-mode TS may be used as a short 2 way data transfer which can be achieved by means of a single 2 TSDU exchange; in this case the application would not be required to open a connection to preserve the sequence of TSDUs in both directions;
- compared with a 2 TSDU exchange over a TC this service would exhibit a reduced overhead, reducing latency from the TS user's point of view;
- finally, the request/response service could be implemented using the acknowledged connectionless mode, however tin the case of the request/response service the TS provider would automatically associate the request and the response TSDUs ;

Time Sequence Diagram

The request/response service is illustrated in Figure 7; a successful scenario consist of four primitives. These are (i) T-REQRESP.request which is invoked at the calling TSAP; (ii) T-REQRESP.indication at the called TSAP; (iii) T-REQRESP.response is issued by the called TSAP; and finally (iv) T-REQRESP.confirm is received by the calling TSAP. Exception may occur and are detailed in [Baguette,92a] One possible scenario is when a negative T-REJECT-REQRESP.indication primitive is received by the calling or called TSAP. In this case the T-REJECT-REQRESP.indication would include an indication why the service failed in the form of a reason code.

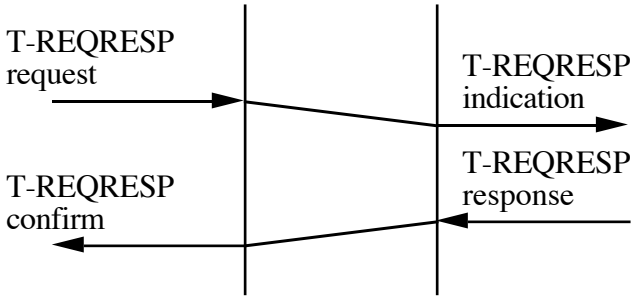

Figure 7 TPX Request/Response Transport Service

In this section we have outlined the new features and services of TPX. The next section offers some architectural remarks surrounding TPX and the ATM.

3.2 Conclusion

A baseline architecture for continuous media support, shown in Figure 8, is a first attempt

to place new OSI95 facilities such as orchestration and ETS into a provisional framework based on ATM networks. We concentrate on ETS, orchestration and AAL support because these facilities have already been well specified by ULg, Lancaster/ELIN, and CCITT/ETSI respectively.

In order to provide architectural support for continuous media, it is necessary to look at the problems associated with mapping orchestration on the ETS at the TSAP. In addition, we need to determine whether the ETS can support continuos media and how ETS can operate in the ATM environment. These issues are addressed in section 4 on evaluation.

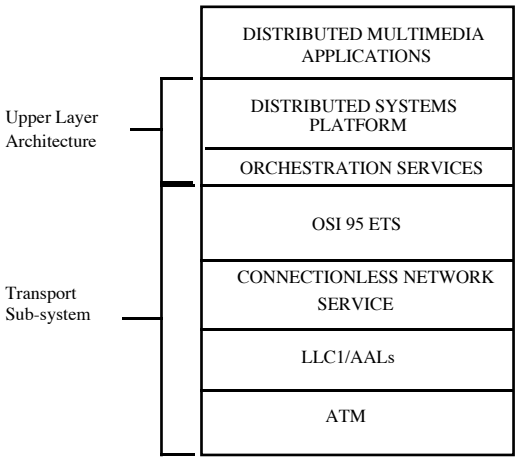

Figure 8: OSI95 Architectural Support for Continuous Media Communications

The stack is divided up into two regions for simplicity; the upper layer architecture and the transport sub-system. Each region is briefly described below:-

- *upper layer architecture* is concerned with the support of distributed multimedia applications. It incorporates data presentation, (specifically compression and decompression), and multimedia profiles selection. The upper layer architecture also interacts directly with the orchestration services in the case of continuous media and directly with TPX for applications which do not require synchronisation services.
- *orchestration* is clearly seen as operating above the TSAP. The main task of the orchestration protocol and services is to provide multimedia synchronisation and QOS management for continuous media communications across multiple related streams. Orchestration services include support for event based and continuous synchronisation and required enhanced QOS support from the underlying transport sub-system.
- *transport sub-system* should included the *(i)* ETS; *(ii)* Connectionless-mode Network Services and Protocol (CLNS/CLNP) [Baguette,92b]; and finally *(iii)* Unacknowledged connectionless-mode services at the data link (LLC1) [OSI95 TA,91]. Further clarification of the harmonisation of OSI and B-ISDN [Boerjan,91] is required before the lower layers of the transport sub-system clearly mapped.

This chapter has presented summaries of the OSI95 deliverables : BELL-1, ELIN-1 and ULg-4. We then briefly described the type of OSI95 architectural support needed to accommodate continuous media applications. In the next section we evaluate the transport sub-system against the service requirements identified in section two and listed in Appendix A.

4. Evaluation

4.1 Introduction

Fundamental to the success of TPX, as a suitable transport service and protocol for the new environment, is its ability to meet the new demanding needs of both distributed multimedia applications as well as the more traditional data communications requirements. The main object of this report is to assess the suitability of the proposed new OSI 95 transport services in meeting the requirements summaries in section two.

4.2 OSI 95 Transport Service Evaluation

The full range of OSI 95 transport services as described in [Danthine,92b] [Baguette,92a] [Leonard,92] and summarised in section two are evaluated in this section. The object of this evaluation is to determine to what extent the OSI 95 TS meet the needs of distributed multimedia communication systems as formulated in [Leopold,92a].

The TPX services identified in [Danthine,92b][Baguette,92a][Leonard,92] are as follows:-

- an Enhanced Connection-mode Transport Service (ETS) ;
- a fast connection-mode service;
- basic (unacknowledged) connectionless mode service;
- an acknowledged connectionless-mode service;
- and a request/response connectionless-mode service.

It should be noted that the authors of TPX have focused on the specification of the ETS [Leonard,92][Baguette,92b]. Consequently, the remaining services are described in much less detail [Danthine,92b] [Baguette,92a].

The evaluation provided in this section is structured as follows. For each functional area, the following information is provided. First, a list of specific requirements in the area is listed. Second, an evaluation of the extent to which TPX meets these requirements is offered. Finally, some conclusions are drawn.

The specific requirements are derived from section two and each requirement is assigned a requirement number. For example in section 2.2.1.3 on error control, [R001] is used to reference the following requirement :-

"One function of a transport protocol is to provide the required end-to-end error rate desired by the application[R001]"

By adopting such an approach, we attempt to provide a more quantitative requirement analysis of TPX. In total, 107 *baseline* communication system requirements are identified, [R001] through [R107]. The baseline requirements broadly related to service and protocol issues of the ULA and the transport sub-system. Because our evaluation concentrates on the transport service access point, we partition the baseline requirements into *TPX service* and *TPX non service* requirements, restricting the analysis of TPX to service requirements only. A summary of the baseline requirements, TPX service requirements and TPX non-service requirements can be found in appendix A2 for quick reference.

4.2.1 Error Detection and Recovery

Distributed multimedia applications have varying sensitivity to errors. In some of the identified applications, for example continuous media applications, uncorrected errors can be tolerated and an elaborate error control scheme is not required. Rather, these applications require a light weight error control which includes notification when errored TSDUs arrive, in order that the application can take some error recovery action. Other applications such as bulk data transfer do however require error free communications. These types of application need the transport service to support error detection and correction to ensure reliable operations.

The error control and error recovery service requirement are listed below:-

Error Control Requirements

[R001] Provide the required end-to-end error rate desired by the application.
[R002] Error announcement (i.e. error detection and indication) at the receiver.
[R003] Error control mechanism for higher reliability i.e. data transfer or real-time.

Error Recovery and Application Requirements

[R008] The new transport mechanism must support flexible error control .

[R012] No error correction.
[R013] Flexible error detection and recovery mechanisms.

4.2.1.1. Requirements Covered by TPX

TPX will provide a selectable error recovery facility [R008], which is selectable by the application [R001]. For example, data transfer and continuous media transfer have very different error control requirements. ETS will permit tuning (that may lead to suppression [R012]) of the error recovery mechanisms [R003]. Multimedia applications may require error detection without error recovery [R013]. An example of an application with such a requirement is a live video stream transfer, where the TS user only requires notification of an error in a TSDU. In this case TPX will deliver an indication primitive [R002] together with the corrupted video data (if available) and then the application can choose to either use the corrupted data or replace it with 'dummy' data. TPX will supply a flexible error management interface allowing the TS user or acceptor to select the appropriate error control mechanisms. The ETS primitives T-Connect and T-Renegotiate also help in accommodating the TS user error control requirement. This is achieved by specifying a *error control policy* at the TSAP. This policy is configurable for each direction of the connection [Baguette,92a] and includes parameters for :-

- *corrupted TSDU error rate* - the TS user can also signify whether or not a corrupted TSDUs should be signalled and/or delivered;
- *lost TSDUs error rate* - in the face loss data, the TS user can indicate to the service provider to replace lost data with prescribed dummy data or not, or just to signal the occurrence of loss.

The transport service interface for the indication of corrupted or lost data is part of the T-Data primitive. [Baguette,92] implies that this *error type indicator* can signal loss and/or corruption of the TSDU, and in some cases point to the region of corruption in a TSDU. This type of error control is well suited to continuous media communications.

4.2.1.2. Comment

The provision of comprehensive error control is comprehensively catered for in the ETS. One of the key qualities in support of error control for the new environment is flexibility. In this area TPX succeeds well and clearly covers all the services requirements in this area.

4.2.2 Timeliness

Multimedia and other real-time applications have data exchange requirements where the validity of the data is time dependent. For example, video TSDUs which arrive after their associated playout times are as useful as lost TSDUs to the receiving application. Therefore, the requirement for timeliness of a TSDU is based on the need to meet a deadline at the receiver. Transport and network providers must be developed to enforce the concept of timeliness.

The TPX service requirements associated with timeliness are as follows:-

Timeliness

[R004] The transport should support bounded delivery of data to the receiver.
[R005] Global time, or a common reference time.
[R006] User indication for failure to transmit a message within its agreed time frame.
[R007] Real-time error sensitive information (real time error control for high reliability).

4.2.2.1. Requirements Covered by TPX

TPX supports the idea of timeliness through the introduction of a new *compulsory QOS* semantic which guarantees [R004] QOS parameters at the TSAP. TPX also assumes bounded throughput, jitter and end-to-end delay at the network services access point. Unfortunately, the network service (CLNS) over which TPX operates can not provide such assurances at the present time.

TPX requires the TS user to characterise its traffic throughput requirement in TSDU size and throughput. TPX requires this to support the compulsory QOS semantic. TPX defines throughput as follows:-

> *"the ratio of the size of the submitted TSDU to the time elapsed until the occurrence of the next t-DATA.request on the same TC"*

Throughput refers to the *instantaneous* throughput rather than the *average* throughput. The latter would be calculated over a fixed period. The TS user specifies the maximum and minimum throughput and the TSDU size. TPX then supports this range by determining the minimum and maximum times between any consecutive TSDU transmissions. TPX allows the TS user to deliver T-DATA.requests under the compulsory minimum throughput; this corresponds to the maximum transmission time. A connection is *not* disconnected in the case of a late submission of a TSDU on the part of the TS user [Danthine,92a]. However, TPX will shut down a connection if the network can not support the minimum compulsory throughput.

4.2.2.2. Requirements Not Covered by TPX

The requirements which have not been covered in the TPX service specification are as follows:-

Timeliness

[R005] Global time, or a common reference time.
[R006] User indication for failure to transmit a message within its agreed time frame.
[R007] Real-time error sensitive information (real time error control for high reliability).

For correctness in the time domain, the communication system should support timeliness (i.e. bounded time limits) of media delivery. This can be achieved by associating a *validity time* with each message together with an associated deadline. This requires the availability of an abstraction of global time, or a common reference time that provides system-wide time synchronisation. The provision of a global time abstraction is a basic issue in distributed real-time systems [R005].

Distributed real-time applications must be informed when the communication system fails to transmit a message within its agreed time frame [R006]. Notification of such a failure allows the application to perform exception handling if necessary; we describe this as *bounded invocation.*

Real-time applications which are sensitive to errors must transmit redundant information in order to recover from loss without retransmission [R007], which if adopted may cause unacceptable delays.

4.2.2.3. Comment

TPX does not meet the baseline requirements [R005][R006][R007]. Clearly the idea of global time [R005] as a common reference is important in distributed continuous media communications. What is not clear, at this moment, is whether this is an issue for the transport function or the higher layers. The relationship between the transport provider and the user with regard to correctness in the time domain is for further study.

Bounded invocation [R006] requires the transport to inform the user if a TSDU has missed its designated transmission opportunity. From the current TPX specification it is not currently clear how this could be engineered.

It is widely accepted by the data communications community that error control should be performed by the transport function, and therefore only at the end points. Furthermore, error control must take into account real-time error sensitive data transfer [R007]. It is not clear from the TPX service specification [Leonard,92] how the TS user specifies error recovery with an accompanying error control constraint. For example continuous media cannot rely on any sort of recovery mechanism which uses retransmission as a means of providing higher reliability because of the unacceptable latency overhead involved.

4.2.3 Profile and Class of Service Selection

The TS should be flexible enough to accommodate a full range of applications with a wide variety of communication characteristics and should exploit the features of the various high performance networks.

Service requirements related to profile and class of service selection are as follows:-

Profile and Class of Service Selection

[R018] The transport service and protocol should be flexible enough, to support the full range of multimedia applications and the wide variety of communication characteristics, and work over different high performance networks.
[R019] The idea of protocol selection.
[R021] User semantics for selection of protocol matrix, rather than network one.

4.2.3.1. Requirements Covered by TPX

TPX supports five transport services which can accommodate a diverse group of application needs [R018]; the inter networking requirements is implicit in the selection of the Connectionless Network Service defined in ISO/IEC 8348. In essence, TPX supports the concept of service selection [R091], and indirectly protocol selection through the T-Connect and T-Renegotiate TS primitives. Furthermore, TPX reinforces protocol selection by incorporating error control information in the T-DATA.indication. If compulsory throughput drops below a pre-defined threshold TPX will inform the TS user of this event. This is referred to as QOS degradation indication in the requirements document.

4.2.3.2. Comment

One of the requirements [R021] is outside the scope of TPX but we mention it here for completeness. This is the consideration that application programmers may not wish to be concerned with the precise details of a possibly complex set of transport service QOS parameters and services. Thus there should arguably be a shorthand way of specifying commonly used 'channel types' by name such as HifiAudio or StandardVideo. These standard channel types could also be parameterised with arbitrary semantics: e.g. Monochrome-Video(MPEG, 512, 256) would mean that the data stream will be compressed according to the MPEG standard and the throughput can be deduced from the required window size of 512 x 256 pixels. These channel types would have to mapped onto specific TPX QOS parameters at the TSAP. Table 1 illustrates an example set of user level channel types and the mapping to TPX QOS parameters.

Channel Type	*Throughput*	*Transit delay jitter*	*Transit delay*	*Traffic type*	*Corrupted TSDU error rate*
StandardVideo	25 Mbps	10 ms	250 ms	isochronous	10^{-3}
SlowScanVideo	1 Mbps	10 ms	250 ms	isochronous	10^{-2}
MPEGVideo	10 Mbps	1 ms	250 ms	isochronous	10^{-9}
VoiceAudio	64 Kbps	10 ms	250 ms	isochronous	10^{-1}
HiFiAudio	2 Mbps	5 ms	500 ms	isochronous	10^{-5}

Table 1 Standard Channel Types with TPX QOS mapping

4.2.4 Security

Security is not covered in any great detail in the requirements document [Leopold,92a] or in the TPX literature. However, secure communications is perceived to be an area of importance within open distributed systems. It is envisioned that both unclassified and classified multimedia applications will utilise the new network environment.

The only requirement mentioned in [Leopold,92a] is a very general one:-

Security Requirements

[R023] Transport support for security.

4.2.4.1. Comment

TPX provides a protection [R023] parameter at the TSAP, as part of its information QOS parameter set, which forms the basis for future security enhancement. In this area TPX should support authentication, integrity, confidentiality and access control which are essential for secure communications.

4.2.5 Heterogeneity

The area of inter-networking is crucial today with the co-existence of several distinct network environments requiring connectivity. Communication support for heterogeneity is therefore an important requirement when designing future transport protocols. We feel the introduction of TPX into the ISO SC6 new work item on Enhanced Communications Functions and Facilities [ISO,92d] presents an opportunity to harmonise TPX with existing networks. As stated in 3.3.4.1, TPX assumes a CLNS/CLNP. Future enhanced network functions and facilities must include multicast services and enhanced QOS, as well as inter-networking protocols.

4.2.6 Quality of Service Parameterisation

With the emergence of digital video, audio, and other continuous media types, increased requirements are placed upon QOS support. For example, to support video connections high throughput and bounded jitter are required and therefore high bandwidth guarantees will have to be made.

The QOS requirements listed in appendix A includes:-

Quality of Service

[R025] Quality of Service parameters.
[R026] Indication that the QOS has degraded.
[R027] Communications sub-system to guarantee QOS is maintained throughout transmission.
[R028] Users QOS semantics incorporating mixed QOS on parameters.

Transport level QOS Parameters

[R029] Throughput.
[R030] End-to-end delay.
[R031] Jitter.
[R032] Error rate.
[R033] Priority.
[R034] QOS tolerance levels for each QOS parameter.
[R035] The requested parameters should then undergo full end-to-end negotiation.
[R036] The finally agreed QOS should then be guaranteed for the duration of the connection.
[R037] QOS degradation indication.

Jitter

[R040] Jitter correction must be offered at the receiver.

Priority

[R043] Multiple levels of priority for multimedia traffic and the pre-emption of network resources for high priority applications.
[R044] The transport service must be flexible enough to allow the transport service user to attach relative priorities across end system connections.
[R045] It should be possible that lower priority connection packets will be dropped first before high priority packets.
[R046] Real-time systems require guaranteed delivery of high priority control messages such as a control message, high priority out of band service required.
[R047] The user should just have to deal with packet errors rate at the TSAP.

3.3.8.1 Requirements Covered by TPX

Because multimedia communications require hard guarantees for throughput, delay and jitter, TPX takes a new approach towards QOS support with its *compulsory QOS* semantic. Today most transport services only offer *effort* communications because most networks providers can only operate at this level of service. With the advent of new multiservice networks, such as ATM-based networks, it will be feasible to offer a wide variety of services and services with hard guarantees. This will allow transport providers, such as TPX, to guarantee compulsory levels of QOS parameters [025][027][R036]. 'Best effort' traditionally means a qualitative semantic without any actual commitment to a given quantitative QOS for, say, throughput or delay. Evidently this type of best effort QOS is not appropriate for the new environment.

Compulsory QOS means that the TS provider must be able monitor [Danthine,92b] any parameter which falls in this category; if the level of service of a compulsory QOS parameter falls below a contracted value then it is up to the TS provider to release the TC.

The ETS allows the user to supply both a minimum and maximum range for compulsory QOS parameters [R034]. In addition TPX presents the TS user with a flexible QOS parameter interface; users can mix compulsory and best effort parameters [R028] on a single transport connection. For example, a TS user may need hard guarantees (compulsory) on throughput and have less stringent requirements (best effort) on jitter. However, it is not clear from the TPX specification [Baguette,92] how this mix might be achieved. Note that section five builds on this new QOS perspective by exploring the notion of a *Quality of Service Architecture (QOS-A)* for multimedia communications.

As described in [Leonard,92][Baguette,92] ETS supports a sets of QOS parameters. Here we divide parameters into two groups, performance QOS parameters and information QOS:-

Performance QOS parameters

- throughput [R029];
- traffic type: identifying the transfer as isochronous or not [R102];
- transit delay[R030];
- delay jitter [R031][R040];

Information transfer QOS parameters

- protection [R023];
- priority [R044];
- TSDU size range [R029];
- error_control_selection [Danthine,92a] [R008]: this parameter provides the TS user with a powerful *error policy* [Baguette,92] individually configurable for each direction of the connection. This error control facility includes a *(i)* corrupted TSDU error rate [R032] with indication ; *(ii)* Lost TSDU error rate [R032] with indication and *(iii)* Dummy replacement indicator and dummy character.

TPX defines three performance QOS parameters *(i)* throughput; *(ii)* transit delay; and *(iii)* delay jitter which can be negotiated (via the T-Connect primitive) and re-negotiated (via the T-Renegotiate primitive) separately for each direction of transmission. Other parameters, such as the *information QOS parameter* (that defines the TSDU size) and the *traffic type* component (isochronous traffic or not) of the throughput parameter can not be negotiated but are fixed by the calling TS user. This is also true for the dummy replacement indicator and the dummy character QOS parameter [Baguette,92]. Only the called TS user can modify the proposed compulsory values of performance parameters at connection time. However, the TS provider may make some limited changes, such as to decrease the maximum throughput (but not below

the minimum) and increase the minimum (but not above the maximum) throughput[1].

The called TS user may modify the performance QOS parameters on reception of a T-Connect indication informing the calling TS user in the T-Connect confirm. Again the called TS user may decrease the maximum throughput and increase the minimum throughput but within the bounds of the initial call. By restricting the allowed modifications, the compulsory character of the QOS proposed by the calling TS user is only strengthened by the negotiation.

TPX takes on board recommendations made in [Leopold,92a] for the re-negotiation of QOS parameters during the lifetime of a transport connection between two applications [R035].

In addition to providing a flexible error control interface, TPX supports a QOS degradation indication [R037] when throughput degrades close to the contracted compulsory minimum value. This service is important for continuous media applications which may tolerate diminished throughput for short periods. What is not clear from the TPX specification is the extent to which the indication of service degradation is extensible to other performance related QOS parameters such as jitter and end-to-end delay.

4.2.6.1. Comment

The ETS accommodates a priority parameter in the T-Connect request. The transport provider and the called TS user may modify this parameter [Danthine,92b][R043][R044]. However, the condition for such a modification is not clear in the literature. [R046] states that real-time systems require guaranteed delivery of high priority control messages such as a control messages even under peak-load condition. ETS logically divides control and data of the TC [Baguette,92]. However, it is unclear whether TPX can differentiate between control and data in terms of transmission priority as control. Moving back to the transport service level, the out of band facility may well suit the transmission of high priority user control information [R042]. Again the TPX service specification [Leonard,92] does not provide much detail about this service.

The need for enhanced QOS parameterisation at the TSAP has been clearly identified in [Leopold,92a]. New services, such as a continuous media service, require a rich set of QOS parameters as a tool to characterise diverse traffic profiles. In addition error control selection and QOS negotiation are key attributes of any new transport facility.

Negotiation of QOS is clearly an important feature embraced by ETS. With the exception of the comments concerning clarification of service, we consider that all the QOS parametisation requirements can be met by the current TPX service specification [Baguette,92].

4.2.7 Guaranteed QOS

Multimedia applications which can accurately forecast their resource requirements may prefer to reserve transport and network resources in advance via a connection-mode service. Applications that require any QOS higher than best-effort (offered by current networks) must request certain guarantees at and above the traditional best effort service. In essence, the transport protocol must provide a QOS environment where rigid application QOS requirements can be met.

The TPX service requirements for guaranteed QOS include:-

Guaranteed QOS Requirements

[R048] Reservation of transport and network resources via a connection-mode service.
[R049] Transactions QOS support.
[R050] Applications that require any QOS higher than best-effort (offered by current networks) must request resources be allocated at the end system and in the network.

[1] The ETS T-Connect and T-Renegotiate primitives offer an additional parameter called “direction to open”. This gives the user the option of stating the direction of flow i.e from the called to the calling TSAP or vis-a-versa.

4.2.7.1. Requirements Covered by TPX

The ETS is a connection oriented service with full end to end QOS negotiation at the transport layer [R048]. Currently TPX relies on a CLNS which provides no guarantees. As pointed out in [Baguette,92b] TPX requires enhancement of the current ISO/IEC 8348 specification, providing an enriched CLNS with QOS reservation and multicast at the NSAP. Work is progressing in the ISO with ECFF. ECFF encompasses the transport, network and data link layers. Furthermore, [Baguette,92] considers that resource reservation of network resources using protocols such as ST-II [Topolcic,90] constitute a noteworthy approach. Such an addition to the network service would allow TPX to offer compulsory QOS at the TSAP.

4.2.7.2. Comment

By extending the traditional best effort model, to include compulsory QOS, TPX provides the type of rich QOS interface required by applications which demand stringent QOS guarantees over the duration of a connection, such as continuous media or hard real time applications. Because continuous media requires true end-to-end guarantees, which include the host operating system and other associated application platforms, it will of course be necessary to build on the QOS offered by the network and transport layer, in order that a true end-to-end service is supported.

As stated earlier [Baguette,92] points out the limitations of the current CLNS, which TPX is based upon. TPX facilities, offered to the TS user, are dependent upon suitable high performance network services and as yet no such network services exist in the standards although work is under way. [Baguette,92] suggests that *"a network service with QOS guarantees under TPX, associated with an additional mechanism for resources assessment and reservation in the Transport Layer, would allow an extension of the current semantic of enhanced QOS in the ETS definition. This extension would aim at introducing the concept of guaranteed QOS at the ETS interface."*

Like any transport protocol, TPX relies on the QOS offered by the underlying network. One function of TPX is to build on the reliability and services supplied by the network. However as stated in [Baguette,92] *"no work is currently going on at the ISO regarding the provision of network service with guaranteed QOS"*. One issue which arise in support of QOS is end-to-end resource reservation. TPX is solely responsible for transport resource allocation and interacting with future network resource reservation mechanisms to support compulsory QOS. However, there is a need to be able to reserve all resources needed by an association, including bounding QOS within the end system. Evidently, this is not within the scope of a TPX alone but requires the provision of an integrated QOS architecture for all layers of the architecture. This need is discussed in detail is section five.

Finally, compulsory QOS has been specified for the ETS. It is, however, still not clear from the specification if TPX will extended this model to the other services such as request/response service; this is for further study.

4.2.8 Dynamic QOS Control

In distributed multimedia systems it is not sufficient to specify a QOS level, protocol profile and service class at connection time which will statically remain in force for the life-time of the connection. A more flexible service is needed whereby users can dynamically alter the QOS of a connection while it is active. This is required because there are applications such as continuous media which have extremely diverse QOS demands which may vary during the course of a connection.

4.2.8.1. Requirements Covered by TPX

There are three classes of negotiation and renegotiation of QOS parameters in ETS [Leonard,92][Baguette,92]. QOS parameters are negotiated and renegotiated using the T-Connect and T-Renegotiate phases respectively.

Class 1

The following QOS parameters are negotiated [R054] between the TS users and the TS provider :-

- throughput ;
- transit delay;
- delay jitter; and
- priority.

The values associated with throughput, transit delay and delay jitter are negotiated for each direction of transfer of the TC, whereas the negotiated priority applies to both directions of transfer.

Class 2

Other parameters, such as protection, error control policy and lost TSDU error rate, corrupted TSDU error rate are negotiated between TS users only. The error control policy and values for TSDU error rates are negotiated for each direction of transmission. The exception here is that protection parameters like priority apply to both directions of the transport connection.

Class 3

Finally in class 3, the remaining QOS parameters include:-

- traffic type indicator;
- dummy replacement indicator; and
- dummy character;

The parameters are exchanged without negotiation; the calling TS user imposes values in one direction and the called TS user in the other. A TS user fixes the traffic type (i.e. isochronous or not isochronous traffic) for its output direction of data transfer. In addition, if a TS user requests an indication when data loss occurs at the receiver it can also stipulate that lost data is replaced by dummy data; in this case the TS user must request the replacement explicitly by supplying dummy data in the T-Connect or T-Renegotiate primitives.

Generally, only the called TS user can modify the proposed compulsory values of performance parameters at connection time. However, the TS provider may make some limited changes, such as to decrease the maximum throughput (but not below the minimum) and increase the minimum (but not above the maximum) throughput but within bounds.

TPX takes on board recommendations made in [Leopold,92a] for the re-negotiation [R054] of QOS parameters during the lifetime of an active transport connection [R052] between two TSAPs. A number of techniques for achieving dynamic re-negotiation of QOS at the TSAP are described in the TPX literature. Problems associated with re-negotiation of performance parameters on an active connection are *(i)* at what point will the new QOS become active; *(ii)* who can re-negotiate; and *(iii)* how are conflicts handled in the protocol. [Leonard,92] and [Baguette,92a] clarify these problems and provide rules to overcome possible conflicts.

Traditionally, transport protocols only provide best effort service with abrupt disconnects. TPX enhances the QOS model by providing (i) QOS negotiation on performance based parameters between end points; (ii) QOS re-negotiation on performance related parameters; compulsory QOS semantic; and finally QOS indication [R053] on errored TSDUs and degrading throughput.

The binding service level requirements are:-

Dynamic QOS Control Requirements

[R052] Dynamically alter the QOS of a connection while it is active.
[R053] QOS degradation indication.
[R054] QOS negotiation and re-negotiation at the TSAP.

4.2.8.2. Comment

The requirements document called for a QOS.degradation.indication when any QOS parameters drop below contracted values. Because continuous media application require flexibility via notification of service degradation, the provision of such an QOS degradation indication would allow the application to take some compensatory action. Clearly, TPX allows negotiation and re-negotiation on throughput, delay and delay jitter, thus meeting the requirements described down in ELIN1. In addition, TPX supports the idea of informing the TS user when a compulsory QOS parameter degrades below a threshold.

TPX incorporates an error type indicator in the ETS T-Data indication primitive. The selectable error control policy is related to the connection set up and renegotiation phases. This is exactly the type of error control indication requested in the requirements document [Leopold,92a] for multimedia applications. Furthermore, ETS [Danthine,92b] supports a complex throughput semantic, which allows the TS user to receive a QOS indication if the throughput is approaching the minimum compulsory level (if selected by the TS user in the error control selection of the connect primitive). Here, again TPX supplies a most useful extension to the QOS semantic, allowing applications to take appropriate action as service degrades; this feature is especially useful for continuous media applications.

One concern is that [Danthine,92b] describes a threshold throughput mechanism which would allow the TS provider to inform the TS user when the contracted throughput has degraded. However, neither the TPX service specification [Leonard,92] or the TPX protocol specification [Baguette,92] reflect this idea. The ability of the service provider to monitor errors and provide indications to the TS user should be extended to cover throughput [Danthine,92b], jitter and transit delay. In section five we expand on the idea of comprehensive monitoring of the performance with respect to the QOS contract between the TS user and TS provider. In addition we also investigate the need for resource reservation, allocation, monitoring and policing at the ULA and transport sub-system; these are seen as key components in a future QOS Architecture for high performance networking [Campbell,92d].

4.2.9 Group and Multicast Support

The requirements document highlights the importance of group and multicast services as characteristics of distributed multimedia communications, especially in the case of emerging CSCW applications. Many distributed multimedia applications require sophisticated patterns of connectivity between different sites which need to be supported at the end systems. Groups and multicast support are arguably the most fundamental capabilities which will lead to the rapid development of open distributed system development. Group and multicast support are two areas which impact the complete architectural framework, from the network to the distributed platform. The addressing capability and semantic provided by TPX is now compared against the list of group and multicast requirements, requirements number [R055] through [R077]:-

Group and Multicast Support Requirements

[R055] The applications study [Williams:91] highlights the importance of group and multicast [R056] Services as characteristics of distributed multimedia communications.
[R057] Note that multicast is important for hard real time data transfer.
[R058] Bulk data transfer.
[R059] Transaction.
[R060] Continuous media services.
[R061] Consequently, the underlying transport mechanism.
[R063] Group connection management at the TSAP.
[R065] Connection management flexibility for diverse data transfer requirements.
[R066] From a simplex connection.

[R067] Complex connection configurations support.
[R068] A multicast service must provide the same QOS associated as with the more traditional point-to-point network access services, in terms of selection of end-to-end delay, jitter, throughput and reliability.

Multicast Services Requirements

[R069] Multicast (one-to-many) communications.
[R070] Broadcast (one- to-all) communications.
[R071] Multi-collect (many-to- one) communications.
[R072] Conference (typically many-to-many in video conferencing configuration).
[R073] Un-acknowledged multicast mode.
[R074] Acknowledged mode.
[R075] Synchronised mode - This mode may be used to implement atomic multicast.
[R076] Group names support for transparency.
[R077] Collection and distribution policy.

4.2.9.1. Requirements Covered by TPX

TPX will provide an extended addressing capability, this will probably be multicast [R56] for data transfer. Future TPX service specification will take into account the work done in ISO JTC 1.21.09.01 on Multipeer Data Transmission (MPDT). MPDT, or multicast, includes a multipeer concept which has also been adopted by the ANSI High Speed Transport Service (HSTS). As stated TPX will provide support for the type of addressing modes described in the Multipeer Addendum [ISO,92c]

Briefly, [Baguette,92a] states that multipeer data transmission will provide peer-to-peer exchange of information between members of a group. Conceptually, information may flow unrestricted between all members of the group. Any group member may send information to group members at any time [R069][R070][R071][R072][R067]. There are several possibilities as to the semantic of multicast. Two of those possibilities are *unacknowledged transmission* [R073] and *acknowledged protocol multicast* [R074][R075]. Along with the semantics of the multicast transmission, MPDT and therefore implicitly TPX specifies a transmitter and receiver regime [R077]. The first regime is where there is only one fixed multicast transmitter. All return messages are unicast to the originator of the multicast. In this regime, it is not necessary for the recipients to know that they are members of a group, or who the other members of the group are, or even that the transmission is multicast. The second regime is similar to the first with the ability to pass the transmitting role on to other participants [R077]. This requires that all group members know about the existence of the group and the definition of a token passing mechanism. The third and final regime specified in MPDT is the unrestricted [R077] case where all group members may send messages to all other members.

A central concept of multipeer data transmission is that of groups [R076] which participate in multipeer communications. Groups can have several distinct properties, these are as follow:-

- Groups are identified by a set of rules;
- Groups are either active or inactive;
- Active group membership is dynamic;
- Groups may be determinate or indeterminate;
- Groups may have transmit and receive regimes;

[Danthine,92b] states that multicast should be extended to cover all TPX services [R057][R058][R059][R060]. The literature [Danthine,92b] also introduces the concept of Active Group Integrity (AGI) for the ETS [R063]. This allows the TS user to request that a connection be set up to one other destination, a number of destinations or a list of destinations [R065]. AGI provides QOS support for the connection duration; this is similar to the new QOS [R068] semantic promoted in the TPX specification. If AGI is can not meet a QOS requirement of the TS user then the service is not accepted or is terminated. Members of a group can be added and removed from an AGI session; when a user leaves a multicast connection the AGI has to be re-evaluated. Finally, ETS is currently specified for peer-to-peer communications

only; this is a product of the lack of multicast support at the network layer[1]. It is planned to add multipeer services to the existing peer-to-peer services of ETS in due course [Baguette,92].

3.3.11.3 Comment

TPX recommends that multicast [MPDT] and group [MPDT][AGI] support be included in the transport services but this is clearly an area which is outside the scope of the current specifications [Leonard,92]. [Danthine,92a] states that multicast should be extended for all services but this is an area for further study; the requirement relating to group support at the TSAP [R061] is not currently met. However, it is recognised as an important requirement which will be met in the future [Danthine,92a][Baguette,92] when an enhanced CLNS becomes available. What is evident is that the authors of TPX consider extending the addressing space to be an important feature of the new transport services and fundamental to meeting the requirements of the new environment

Most of the group and multicast requirements [R055] through [R077] have been covered either implicitly, by incorporating the work carried out on MPDT and AGI into the extended addressing models, or directly in the TPX specification documentation. Group and multicast support is an area of critical importance to the success of TPX. Therefore, we would have liked to have seen more details of the impact of this extended addressing provision on TPX transport service specifications. Even though most of the requirements covered are in a very early stages of service realisation, we are happy that this work is in hand, and that TPX service will support a multicast service and group semantics in due course.

It is not clear from the TPX literature [Danthine,92b] how the multicast addressing capability at the transport layer maps to the network [R062][R064]. It is also unclear what assumptions about the underlying networks addressing TPX makes (other than an enhanced CLNS).

TPX only supports full duplex connection via the enhanced connection-mode service. It allows the user to fully specify the QOS for each on direction of the full duplex, which is desirable and flexible. However, it is preferable with CM transport systems to provide an option of simplex configuration of virtual circuits [R066] at the TSAP rather than the more conventional full duplex connections. This follows from the inherently unidirectional nature of many CM transmissions; for example, a remote camera to local video monitor connection is typical in a multimedia conferencing environment. The Stream abstraction at the platform level is also unidirectional. As resources must be explicitly reserved for CM connections, it is wasteful of network capacity to support a full duplex connection if only unidirectional transfer is required. The reviewers understand that TPX is not solely interested in CM and it goes along way to meeting the simplex requirement [R066]. The ETS must provide general service support across a spectrum of applications. On reflection, however, the ETS offers QOS configurability in each direction of transmission which is acceptable for CM. Because some TS users may only require a simplex connection, we suggest that TPX accepts a null QOS parameter, possibly throughput, as an indication to suppress one of the simplex connections of its current full duplex.

4.2.10 Enhanced Transport Services

Traditional transport services, such as OSI Transport Class 4, have mainly serviced the communication needs of file transfer and low-bandwidth interactive applications. It is generally accepted that these protocols are not suitable for multimedia and real-time communications, and that dedicated transport services and supporting protocols are needed to exploit the performance of the new high speed network environment such as the ATM based multiservice networks.

The following new service requirements have been identified:-

---

[1] Currently TPX [Baguette,92b] assumes CLNS which has no support for a multicast service. ULg plan to extent ETS to support multicast but assume that CLNS will be enhanced with a multicast NSDU transfer capability

Enhanced Transport Services Requirements

[R084] continuous media service.
[R085] transaction oriented service.
[R086] Bulk data transfer service.
[R087] These transport services should operate over connection-oriented.
[R088] Connectionless services.
[R089] Error control for assured communications
[R090] Unacknowledged capabilities.

4.2.10.1. Requirements Covered by TPX

TPX supports many of the requirements essential to provide a continuous (isochronous) media service [R084]. Continuous media services will operate over the ETS [R087] and will depend on the new compulsory QOS parameters, flexible error control capabilities and traffic type selection. A transaction oriented service [R085] is supported by the TPX request/response transport service or the acknowledged connectionless-mode service [R088]. The traditional bulk data transfer services [R086] can be provided by the ETS. However, it is not clear from [Leonard,92] how one would select assured mode for reliable delivery of unsequenced data Finally, inter-networking and RPC type service are supported by the TPX connectionless-mode service [R090].

Comment

TPX provides a rich set of transport services which meet most of the enhanced service needs for the new environment.

4.2.10.2. Continuous Media Service

Continuous media is the essence of multimedia and has by far the most stringent and novel requirements in terms of guaranteed QOS and orchestration. Continuous media traffic differs from traffic generated by traditional distributed applications. Instead of the bursty traffic generated by, e.g., remote procedure calls, continuous media packets are transmitted by an application at a fixed rate (e.g. a video frame every 33 ms) and received at a fixed rate. This kind of traffic is often referred to as isochronous and is characterised by its strict adherence to periodicity, in contrast to the asynchronous nature of bulk data transfer.

Traditionally, the end-to-end flow control mechanism has been transparent to the TS user, and therefore the requirement [R094] is legitimately outside of the scope of a transport service requirement, in fact it is a requirement associated with the protocol mechanism itself. We therefore do not consider this a binding requirement on TPX services but we include it as a strong recommendation which TPX could adopt in future. Currently, the TPX specification [Baguette,92] does not discuss the internals of its protocol transfer mechanism. Because of the introduction of isochronicity at the TSAP, we believe TPX will naturally be rate-based flow control oriented. Rate base flow control [Cherition,86][Chesson] [Clark,88] is more suited for transporting continuous media. Attractive characteristics include the de-coupling of flow control from error control mechanisms, and the natural correspondence between the notion of continuous data flow and rate controlled mechanism. In section five, we look at the architectural considerations surrounding the need to support continuous media at all layers of the stack.

A continuous media service is based around a connection-mode service where resources are reserved in advance of transmission. It is recognised that some of these requirement listed below may be duplicates of previously listed requirements. However, for a transport to offer a continuous media service it must collectively [R091] support the following requirements [R091]-[R107]:-

Continuous Media Transport Service Requirements

[R091] Continuous media requires new transport mechanisms [R92] to [R103] inclusive.
[R092] High bandwidth.
[R093] A flexible approach to error recovery; including real time sensitive error control.

[R096] QOS degradation indication.
[R097] QOS negotiation and re-negotiation.
[R099] Single and multiple streams.
[R101] Group/multicast.
[R102] Isochronous delivery mechanism
[R103] Non multiplexed simplex connections.
[R104] Reliable request/response.
[R105] Interaction versus unreliable.
[R106] Maintaining Connections with Long Lifetimes.
[R107] Reliable, retransmission of corrupted data is required.

Requirements Covered by TPX

TPX offers support for continuous media (isochronous) communications at the TSAP. The TS user is supplied with a traffic type indicator parameter in the T-Connect request The traffic type can be either set to non-isochronous or isochronous. If the TS user specifies isochronous then the TS provider guarantees to provide a transmission opportunity to the TS user at a fixed rate [R102], described in the literature as the *isochronous time*. So as long as the TS user supplies T-Data request at a fixed rate at the TSAP, TPX will ensure that this behaviour is also guaranteed by TPX protocol itself.

As stated earlier in this assessment TPX meets the following requirements:-

[R092] High bandwidth.
[R093] A flexible approach to error recovery.
[R096] QOS degradation indication.
[R097] QOS negotiation and re-negotiation.
[R099] Single and multiple streams.
[R101] Group/multicast.
[R103] Non multiplexed simplex connections.
[R104] Reliable request/response.
[R105] Interaction versus unreliable.
[R106] Maintaining Connections with Long Lifetimes.
[R107] Reliable, retransmission of corrupted data is required.

Comment

The ETS provides the required service interface for continuous media communications. The next stage in building continuous media support is to look at the mapping of orchestration services on the ETS[1].

[1] Orchestration is a ULA activity. It monitors the temporal relationship between multiple related streams or connections. After conducting this analysis we see no reasons why orchetration could not be implemented on top of the ETS. This is for further study by Lancaster.

4.3 Recommendation

The preceding analysis of TPX services and features against the distributed multimedia communications system requirements has shown that TPX provides a rich base of services which covers many of these new requirement. The case of non compliance of TPX services to the new requirements has been described in four. In section four we took an initial look at the problems associated with mapping TPX services to AAL services. Even though mapping was not comprehensively covered in this section, we see no reason why TPX should not be able to efficiently exploit the capabilities of ATM technology. A comprehensive mapping of TPX to AAL is outside the scope of this assessment.

The integration of distributed multimedia systems support into a communications architecture encompassing the new multiservice networks poses significant challenges. In this section we have presented an evaluation of the emerging TPX and AAL services against the new requirements. A key observation is that QOS provides a unifying theme around which most of the application communication service needs, TPX services and AAL services can be grouped.

For applications relying on the transfer of multimedia information, and in particular continuous media services, it is essential that QOS is guaranteed system-wide, including the distributed system platform, the TPX protocol and the ATM network. Enhanced protocol support such as end-to-end QOS negotiation, re-negotiation, indication of QOS degradation's and co-ordination over multiple related connections are also required. However, little attention has so far been paid to the definition of a coherent framework that incorporates QOS interfaces, management and mechanisms across all the layers. The next section describes the first stage in the development of an integrated QOS-A, as shown in Figure 10, which provides a framework to specify and implement the required performance properties of new multimedia applications using the OSI 95 TS over ATM-based networks.

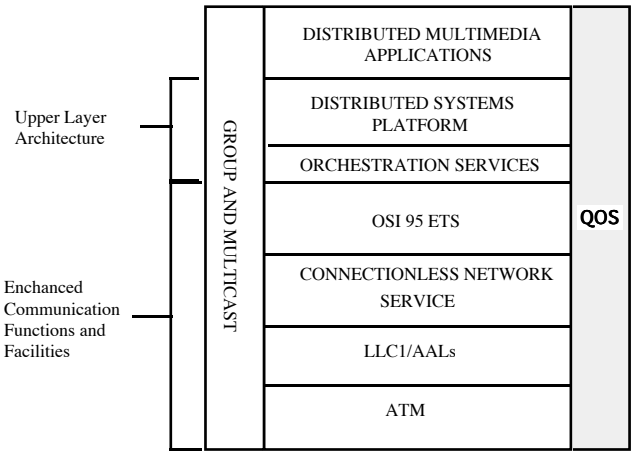

Figure 10 An Integrated QOS-A for Distributed Multimedia Communications

Like any transport protocol, TPX relies on the QOS offered by the underlying network. One function of TPX is to build on the reliability and services supplied by the network. One issue which arise in support of QOS is end-to-end resource reservation[1]. TPX covers transport and network resource reservation only, however there is a need to be able to reserve all resources needed by an association, including bounding QOS with in the operating systems. Evidently, this is not within the scope of a TPX alone but requires the provision of an

[1][Baguette,92b] adds that resources reservation mechanisms are important to the OSI 95 TS compulsory QOS semantic being realised.

integrated QOS architecture for all layers of the architecture. This is discussed in detail is section five

5. Integrated Quality of Service for Multimedia Communications

5.1 Introduction

New communications infrastructures should offer comprehensive QOS configurability for a wide variety of media types and specific user requirements. An important first step in meeting such a requirement is the specification of a *Quality of Service Architecture* (QOS-A) which provides an integrated framework for QOS specification which incorporates TPX's compulsory QOS semantic and resource control over all architectural layers from the distributed application platform to the ATM network.

The approach of this section is to present a set of key QOS requirements and map these requirements onto a provisional QOS-A which has emerged from [Leopold,92a] [Danthine,92b][Boerjan,91][Campbell,92d]. Because of the likely complexity of a fully general QOS-A, we limit the scope of our discussion in this report to aspects of a QOS-A for the support of continuous media communications. We also concentrate on ATM at the network layer rather than consider the full range of multiservice networks. However, a generalised QOS-A should eventually be extensible to incorporate other areas of QOS provision such as real-time control systems, file transfer and real-time transaction processing.

5.1.1 The Need for a QOS-A

5.1.1.1. QOS-A Requirements

Because of the increased range and complexity of QOS provision required by the emerging distributed applications, it becomes essential that the necessary extensions to QOS provision are not done on a piecemeal basis. Instead, we advocate the notion of a comprehensive QOS-A, whereby application requirements can be mapped through all the levels of the system. Thus, communications abstractions at the application platform level should provide QOS abstractions which can be mapped down through all intermediate layers to the multiservice network access point in a coherent and integrated way. This mapping should be clean, simple and efficient, protecting application programmers from communication details; that is, they should be stating what they require rather than how it is to be achieved.

In addition to mapping functions, the QOS-A should provide a framework for the support of QOS throughout the layers. In particular, there is a need for enhanced protocol support which is lacking in current protocol specifications. The framework will include *management functions* and selectable QOS support *mechanisms*. Examples of management functions are:-

- end to end QOS negotiation including admission control for new connections;
- policing to ensure that users are not violating negotiated QOS parameters;
- monitoring to ensure that negotiated QOS levels are being maintained by the service provider;

QOS support mechanisms take the form of a pool of available procedures which can be configured and inserted into a protocol stack. Depending on the QOS contract established at negotiation time, the QOS-A would be responsible for building a suitable stack profile by selecting from the available set of mechanisms. Examples of such mechanisms are error control modules, parameterisable scheduling modules and jitter smoothing modules. Once again, many of these mechanisms will be applicable at multiple system layers. For example, scheduling appears in the context of end user threads and also in network switches.

5.1.2 Current Notions of QOS

Traditionally, the term 'quality of service' in the communications context referred to certain characteristics of network services as observed by TPX users. These characteristics were not controllable by users, and described only those aspects of services attributable to the network provider. QOS parameters in current communications infrastructures, like OSI and CCITT do

permit the specification of some user requirements but these are almost never supported by the underlying network. For example, the current OSI standards threat QOS in a layer specific way, and QOS definition has been looked at by separate committees (i.e. the presentation, session, TPX, network and data link committees) working in isolation. Thus the relationship between QOS layers is not clearly defined and there is no consistent, integrated notion of QOS which relates user requirements to the network provider services.

The following sub-sections review the degree of QOS provision in sample architectures at the session, transport and network layers. The session and transport layer specifications are taken from the ISO's Reference Model for Open Systems Interconnection (OSI-RM) [Henshall,88] and the network specifications are taken from the CCITT's series I recommendations for ATM cell switching.

5.1.2.1. Evaluation

The clearest point to emerge is that QOS is currently looked on largely as a service provider issue whereas the requirements identified in section 2 imply that QOS should also be a user level issue. It is also clear that the OSI standards are currently incomplete and inconsistent for specifying QOS properties. In particular, the protocol specifications and service definitions do not include any notion of QOS management and the semantics of responsibilities and guarantees are not clear. Furthermore, those functions which are defined are almost never supported by protocols and networks.

Another important point is that the OSI upper layers have no notion of QOS: QOS parameters are simply mapped through to the TPX layer. This is also true in alternative upper layer architectures such as the object-based ODP architecture [ISO,89]. If users want to specify QOS they are forced to drop below the level of abstraction provided by these architectures and interact with layers that are supposed to be hidden. Furthermore, there are very limited facilities for QOS negotiation at the user level. A user must simply specify the parameter values required and let the lower layers either accept or reject the proposal.

The CCITT's ATM recommendations are more comprehensive in scope with a fairly detailed traffic characterisation model. Here it is the mapping between the higher layers and the ATM adaptation layer, and also the mechanisms required to support particular QOS specifications which are lacking in substance. These are precisely the concern of a generalised QOS-A. An important step in our work will be resolving the present inconsistencies in the relationship traffic characterisation parameters of ATM and the OSI-RM. Other requirements are the development of protocol support for QOS in terms of the various QOS management functions and support mechanisms, examples of which were given in section two above.

Finally, a major limitation of all current notions of QOS is that the value of a QOS parameter, whether negotiated or not, remains the same through the lifetime of a connection: i.e. once negotiated a QOS parameter is never re-negotiated. Another implication of this is that the service-provider is committed to provide the QOS over the lifetime of the connection. There is, however, no guarantee that the service-provider will be able to maintain the originally specified values: in fact maintaining end-to-end service levels in the face of variable load is an unsolved problem involving resource scheduling at multiple levels. Even when the QOS of a connection does deteriorate the service provider is under no obligation to signal such a change in QOS to the users of the connection. The provider may, however, disconnect the connection unilaterally.

The essential characteristics of the current state of QOS provision may therefore be summarised as follows:-

- *lack of overall framework*: the framework for QOS must extend from the distributed application platform through the transport sub-system and the network. It must also encompass QOS considerations in areas such as orchestration and groups;
- *inconsistency*: the framework must build on and reconcile the existing notions of QOS particularly in the OSI-RM and the ATM series I recommendations;

- *incompleteness*: the framework should include extensions to current QOS provision as detailed in the following section;
- *lack of mechanisms to support QOS guarantees*: research is needed in basic mechanisms such as scheduling so that contracted QOS levels can, in fact, be maintained.

5.1.3 An Extended View of QOS

5.1.4 Baseline Architecture

Architectural Layers

The baseline for the development of the QOS-A is the layered architecture depicted at the left hand side of figure 11. This has been derived from our experimentation to date with distributed multimedia applications. A detailed description of our current infrastructure and its implementation is given in [Davies,91]. The remainder of figure 12 illustrates the aspects of the QOS-A to be described in this section.

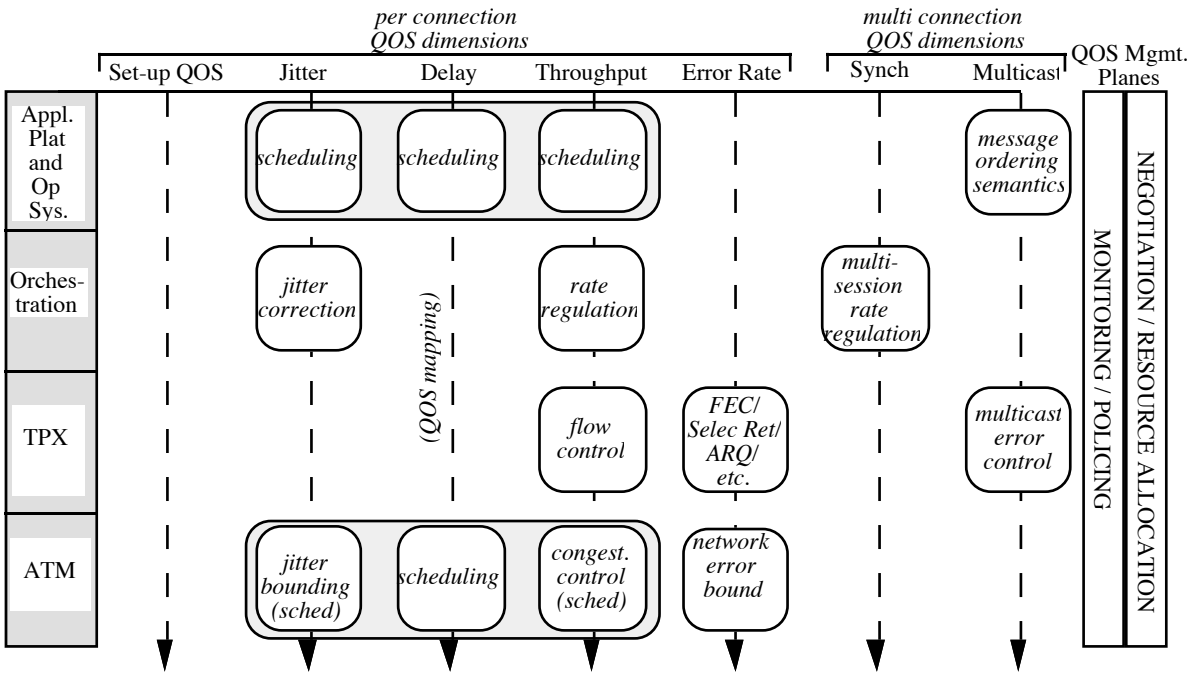

Figure 11: QOS-A

The upper layer in the layered architecture consists of distributed applications platform which is provided by an ODP compatible distributed systems platform augmented with services to provide multimedia communications, QOS configuration and synchronisation [Coulson,90].

Below the platform level is a layer of services used to add value to the functionality provided by the lower transport layer. Specifically, these services control jitter and rate regulation for continuous media streams. They also provide these services, together with cross-stream synchronisation, across multiple application related connections. Because these services are concerned with co-ordinating multiple sources and sinks we refer to them as *orchestration* services; a full description of these services can be found in [Campbell,92a].

Below the orchestration services is TPX which includes an enhanced connection-mode service [Danthine,92b] and a new QOS semantic which are suitable for continuous media communications. It is highly configurable in terms of QOS and offers full end-to-end QOS negotiation and re-negotiation.

5.1.4.1. QOS Dimensions

In figure 11 we have attempted to extract a canonical and orthogonal set of *dimensions* within which traffic can be characterised in our chosen domain of continuous media communications. The chosen set of dimensions are: *set-up QOS* (i.e. the OSI establishment and release parameters), *jitter* (i.e. variation in delay), *delay*, *throughput* and *error characteristics*. In addition to these fundamental dimensions, two additional dimensions are included: *synchronisation* between media streams, and aspects of *multicast* quality of service. The essence of these latter two dimensions is that they are applicable over multiple connections whereas the others apply to single connections. In fact, the multi-connection dimensions also subsume the fundamental dimensions but additional quality of service characteristics arise as emergent properties.

Later sections describe how traffic may be specifically characterised at the various layers, and how levels of service along the canonical QOS dimensions are maintained through profile selection at the different layers.

5.1.4.2. QOS Management Functions

QOS Negotiation

The most fundamental aspect of the QOS-A is the interface at which desired levels of QOS can be requested, negotiated and contracted. In a layered architecture such as figure 11, there are multiple instances of this interface; each instance has a user above the interface and a provider below. The function of the QOS-A here is to permit end-to-end QOS negotiation from the top user level down to the network layer and up again at the remote site. A successful negotiation at each interface level results in a contract with two major clauses:-

- an agreed *level of service* which the provider level must undertake to maintain, and
- an agreed *level of traffic* which the user level must undertake not to exceed.

Both the level of service and the level of traffic will be expressed in terms of a common, layer specific, *traffic characterisation* language based on the fundamental QOS dimensions. Examples of layer specific traffic characterisations appear in section 4.3.

A further aspect to the contract is the *degree of commitment* in the above clauses: e.g. is the provider committed to maintaining the level of service in all conditions or are there circumstances in which the level of service may be relaxed? A related question is what sanctions will be imposed by the provider if the contracted traffic level commitment is exceeded by the user? Both of these points can also be related to the *cost* of the service; presumably a higher commitment by the provider for the same nominal service will cost more, as will the option of a lower commitment from the user.

To express degrees of commitment either a relative measure such as priority levels can be used, or absolute measure such as a percentage. Absolute measures can also be expressed as a step function such as {*deterministic, probabilistic, best-effort*}. An absolute scale is also appropriate for cost measures. Even if only a step function was ultimately available at the bottom level, percentages measures of commitment at higher levels could be used to express trade-offs. For example, in a videophone connection consisting of separate video and audio channels, a slightly lower commitment probability for the video channel would be appropriate Even if the two probabilities chosen mapped to the same step function value at the lower layer (e.g. probabilistic), the higher commitment of the audio channel would ensure that if one channel had to be taken down due to lack of resources, the more important audio channel would be preserved.

As pointed out in section 2 the same QOS interface is not necessarily appropriate for all layers in a layered architecture. However, in abstract terms, the interface at each layer will consist of some subset of the following general components:-

- a notation for quality characterisation along the QOS dimensions,

- a notation for commitment specification,
- a notation for cost specification, and
- a protocol for QOS negotiation.

There are other possibilities in the negotiation process which may be of use at different layers. For example, rather than simply proposing a level of service, upper and lower bounds on acceptability could be proposed by the user and the provider could return a contract as close to the upper bounds as possible. Also, re-negotiation of the QOS on a live connection could be permitted. Finally, the user could specify alternative *degradation paths* to be taken when commitments are not met by the provider. Possible alternatives are: ignore the situation, simply inform the user, or inform the user and reconfigure according to a user specified degradation path. As an example of the latter a hifi audio channel could be degraded to a 64Kbps voice audio channel. At the same time, the system would inform the user who would adjust the source and sink codecs as appropriate.

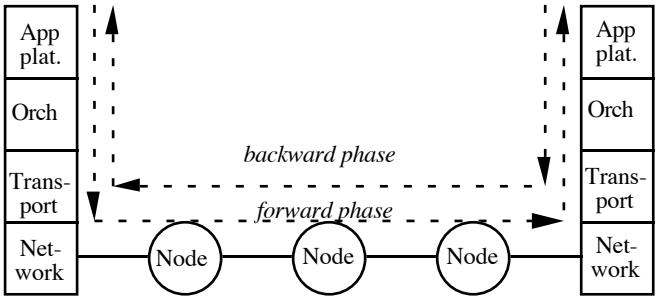

Figure 12: End-to-end QOS Negotiation

The mechanism of negotiation is illustrated in figure 12. In broad terms, negotiation is a two phase process. On the forward phase, from the source towards the sink, each intermediate resource holding node or end-system layer attempts to allocate resources to the request. Each layer in the system contains an *admission control* module which determines whether or not the request should be considered. For example, the operating system will decide if it is able to create a new real-time thread and the network layer will determine whether or not it can allocate a new connection. If the request is accepted by admission control, as many resources as are available are dedicated to the request. Eventually, at the sink end, the allocated QOS is compared against the requested requirement along each QOS dimension and the amount of over-commitment, if any, is calculated. On the second phase, from sink to source, any over-commitment is divided among the intermediate nodes and the over committed resources are released. Examples of negotiation mechanisms from the literature include [Anderson,91] and [Topolcic,90] which address negotiation in the network. Full end-to-end QOS negotiation is discussed in [Nicolaou,91].

Monitoring and Policing

QOS must be monitored at all layers of the QOS-A to ensure that the negotiated levels of QOS are being maintained. Each layer will collect statistical information associated with the on-going connection performance and make an assessment of the QOS measured against QOS requested. Monitors may then either attempt to take action to restore QOS levels themselves or they may choose simply to inform the upper layer that there is a problem.

Policing, on the other hand, is primarily seen as a network (and perhaps transport) function and may not be carried out at all layers. In particular, the application platform and orchestration layers are simply not able to violate QOS agreements by the nature of QOS support at those layers. Examples of monitoring and policing at the various levels in figure 11 are as follows:-

Application platform

An important consideration at the platform level is the monitoring of end-to-end communications and will thus take into account possible QOS degradation's arising from the end-system in addition to those arising from the communications sub-system. Application

platform level monitoring will be mainly concerned with monitoring the performance of the operating system thread scheduling mechanism to ensure that data which arrives correctly is also delivered correctly to its final destination.

Orchestration

A fundamental part of orchestration is the ability to monitor the on-going temporal relationship between of connections, and to regulate the connections to perform fine grained corrections if synchronisation is being lost. It is almost inevitable that related connections will eventually drift out of synchronisation due to such factors as discrepancies between remote clock rates and network congestion caused by temporary blocking at intermediate switches. The degree of multimedia synchronisation accuracy required by the user is viewed as a monitoring issue in the QOS-A. The monitor may choose to attempt to correct the drift (perhaps by renegotiating transport QOS), or if the drift is too great it may simply inform the upper level.

TPX

The transport level monitoring entity captures statistical information related to the connection performance between two TSAPs. In the case of guaranteed QOS, the QOS state for each parameter is periodically measured to determine the *(i)* typical response time; *(ii)* average throughput over an interval *(iii)* amount of time buffering/blocking in the network per packet; *(iv)* worst case responses time; and *(v)* best case response time. If any of the negotiated QOS parameters degrade below the specified minimum tolerance, then a *QOS.Indication* [Campbell,92a] is raised by the monitoring mechanism detailing which QOS parameter has degraded and its current measured value, the indication also includes a statistical profile of all monitored parameters. On receiving a QOS.Indication, the application is free to make a judgement based on the current connections performance; for example the application may accept a poorer QOS in the light of network congestion, on the other hand it may initiate a full end to end re-negotiation and also upgrade the level of commitment of the connections QOS.

Network

Once a connection has been successfully negotiated between two end points, the source must be policed at the edge of the network to ensure that it does not exceed the traffic profile declared at ATM call-setup time. This is particularly important if a statistical multiplexing approach is assumed. Should the user exceed the agreed throughput levels, the policing function will intervene and discard cells. In addition a network level QOS.Indication will be generated indicating a *bandwidth violation* has been detected on the connection. On reception of the indication, the application can either regulate its traffic or re-negotiate a higher throughput to meet its end to end throughput requirement.

In a logical sense, monitoring is also required at the network level to ensure that the user is receiving the negotiated level of service. However, we envisage that this will mainly be the responsibility of the transport. Broadly speaking, the transport will monitor and the network will police.

6. Conclusion

6.1 The OSI 95 TS and the New Environment

The OSI 95 TS[1] was evaluated against a total of 83 communication service requirements out of 107 baseline requirements listed in Appendix A. Of these 85 service realted requirements, the OSI 95 TS covered 59 explicitly, 21 implicitly and 3 not at all. The major area of non-compliance was group and multicast support .

TPX recommends that multicast [MPDT] and group [MPDT][AGI] support be included in the OSI 95 but this is clearly an area which is outside the scope of the current specifications. [Leonard,92] [Baguette,92a] states that multicast should be extended for all services but this is an area for further study; the requirement relating to group support at the TSAP [R061] is not currently met. However, it is recognised as an important requirement which will be met in the future [Danthine,92a][Baguette,92b] when an enhanced CLNS becomes available. What is evident is that the authors of TPX consider extending the addressing space to be an important feature of the OSI 95 TS and fundamental to meeting the requirements of the new environment. For this reason we do not argue that TPX fails to meet the 22 group and multicast communication service requirements.

Most of the group and multicast requirements [R055] through [R077] have been covered either implicitly, by incorporating the work carried out on MPDT and AGI into the extended addressing models, or directly in the TPX specification documentation. Group and multicast support is an area of critical importance to the success of TPX. Therefore, we would have liked to have seen more details of the impact of this extended addressing provision on TPX transport service specifications. Even though most of the requirements covered are in a very early stages of service realisation, we are happy that this work is in hand, and that TPX service will support a multicast service and group semantics in due course.

The provision of comprehensive error control is comprehensively catered for in the ETS. One of the key qualities in support of error control for the new environment is flexibility. In this area the ETS succeeds well and clearly covers all the services requirements in this area. TPX provides a protection [R023] parameter at the TSAP, as part of its information QOS parameter set, which forms the basis for future security enhancement. In this area TPX should support authentication, integrity, confidentiality and access control which are essential for secure communications. The ETS accommodates a priority parameter in the T-Connect request. The transport provider and the called TS user may modify this parameter [Danthine,92b][R043][R044]. The OSI 95 TS out of band facility may well suit the transmission of high priority user control information [R042]. However TPX service specification [Leonard,92] does not provide detail about this service.

The need for enhanced QOS parameterisation at the TSAP has been clearly identified in [Leopold,92a]. New services, such as a ETS, provide a rich set of QOS parameters as a tool to characterise diverse traffic. In addition error control selection and QOS negotiation are two key attributes of the OSI 95 TS. Negotiation of QOS is clearly an important feature embraced by ETS; we consider that all the QOS parameterisation requirements can be met by the current TPX service specification [Leonard.92].

By extending the traditional best effort model, to include compulsory QOS, ETS provides the type of rich QOS interface required by applications which demand stringent QOS guarantees over the duration of a connection, such as continuous media or hard real time applications. Because continuous media requires true end-to-end guarantees, which include the host operating system and other associated application platforms, it will of course be necessary to build on the QOS offered by the network and ETS, in order that a true end-to-end service is supported.

---

[1]Most of our comments here are related to the ETS of the OSI 95 TS

As stated earlier [Baguette,92] points out the limitations of the current CLNS, which TPX is based upon. TPX facilities, offered to the TS user, are dependent upon suitable high performance network services and as yet no such network services exist in the standards although work is under way. [Baguette,92] suggests that *"a network service with QOS guarantees under TPX, associated with an additional mechanism for resources assessment and reservation in the Transport Layer, would allow an extension of the current semantic of enhanced QOS in the ETS definition. This extension would aim at introducing the concept of guaranteed QOS at the ETS interface."* The compulsory QOS has been specified for the ETS. It is, however, still not clear from the specification if TPX will extended this model to the other services such as request/response service; this is for further study.

The requirement document [Leopold,92a] called for a QOS.degradation.indication when any QOS parameters drop below contracted values. Because continuous media application require flexibility via notification of service degradation, the provision of such an QOS degradation indication would allow the application to take some compensatory action. Clearly, ETS allows negotiation and re-negotiation on throughput, delay and delay jitter, thus meeting the requirements described down in ELIN1. In addition, TPX supports the idea of informing the TS user when a compulsory QOS parameter degrades below a threshold.

The OSI 95 TS incorporates an error type indicator in the ETS T-Data indication primitive. The selectable error control policy is related to the connection set up and renegotiation phases. This is exactly the type of error control indication requested in the requirements document [Leopold,92a] for multimedia applications. Furthermore, ETS [Danthine,92b] supports a complex throughput semantic, which allows the TS user to receive a QOS indication if the throughput is approaching the minimum compulsory level (if selected by the TS user in the error control selection of the connect primitive). Here, again TPX supplies a most useful extension to the QOS semantic, allowing applications to take appropriate action as service degrades; this feature is especially useful for continuous media applications.

One concern is that [Danthine,92b] describes a threshold throughput mechanisms which would allow the TS provider to inform the TS user when the contracted throughput has degraded. However, neither the TPX service specification [Leonard,92] or the TPX protocol specification [Baguette,92] reflect this idea. The ability of the service provider to monitor errors and provide indications to the TS user should be extended to cover throughput [Danthine,92b], jitter and transit delay. In section 4.0 we expand on the idea of comprehensive monitoring of the performance with respect to the QOS contract between the TS user and TS provider. In addition we also investigate the need for resource reservation, allocation, monitoring and policing at the ULA and transport sub-system; these are seen as key components in a future QOS Architecture for high performance networking [Campbell,92d].

We have presented a draft QOS-A based on our experimental work and have discussed requirements and possible mechanisms for QOS support. In addition AAL services evaluated in section 3.2 , also provide the range of services required to meet the transmission needs of the new environment. In section 3.4 we took an initial look at the problems associated with mapping TPX services to AAL services. Even though mapping was not comprehensively covered, we see no reason why the ETS should not be able to efficiently exploit the capabilities of ATM technology. A comprehensive mapping of TPX to AAL is outside the scope of this assessment

[Danthine,92b][Baguette,92a][Leonard,92] present a rich set of transport services which meet most of the enhanced service needs for the new environment. Finally, the ETS provides a continuous media communications service which is a real strength of the new facility. The next stage in building continuous media support is to look at the mapping of orchestration services on the ETS[1].

[1] Orchestration is a ULA activity. It monitors the temporal relationship between multiple related streams or connections. After conducting this analysis we see no reasons why orchetration could not be implemented on top of the TPX ETS. This is for further study by Lancaster.

6.1.1 QOS-A

This deliverable has argued for a comprehensive QOS architectural framework in the light of evaluating the OSI 95 TS against the new applications QOS requirements in the new environment. Standards have an important role to play in development of a QOS-A for high performance distributed computing. The current OSI and CCITT standards do not take into account the particular requirements of an integrated QOS model. It can therefore be anticipated that the OSI RM will have to be extended to support the new QOS requirements [Campbell,92d].

In ISO, a New Work Items (NWI) on QOS and ECFF has recently been initiated (ISO/IEC JTC1/SC21 and SC6 respectively). The NWI on QOS aims to address QOS in a consistent way. This activity will cover QOS very broadly, but has begun by investigating user requirements for QOS [ISO,92a] and some architectural issues [ISO,92b]. We have submitted a contribution to the ISO [Campbell,92f] describing some early results of work on QOS. It is Lancaster's intention to participate actively in it at both these NWIs; we see a need for close co-operation between SC21 and SC6 on the issue of a QOS-A for high performance communications.

One potential difficulty in achieving a unified QOS-A is that the ISO and CCITT have contrasting views on QOS. The CCITT perspective is network oriented and the ISO approach is geared towards the user. Another important view point, which is just being taken into account, is Open Distributed Processing QOS requirements. This has been recently initiated in the UK with a liaison between ODP[1] and IST/21/-/1/5 panels. We see the work done in OSI/ULA -ODP architectural mapping presented in ELIN-3 [Singer,92] as being important to the development of a QOS-A which provides high performance communications in an ODP environment

[1]Appendix C investigates the type of QOS support required by ODP.

Appendix A

In Appendix A1 we presents an abbreviated list of the distributed multimedia communication system requirements, distilled from section 2.1. The remaining part of Appendix A is structured as follows. Because the complete list of requirements includes non TPX service requirements, such as orchestration, we define the *TPX Service Requirements* from the complete list in A1. The baseline requirements, presented in A2, and associated requirement numbers are used in the evaluation section of this report. Each requirement is tested against the current TPX specification. Any discrepancies between TPX and the requirements are listed in section 2.1.

Appendix A3 lists any requirements not included in the evaluation of TPX. These *Non TPX Service Requirements* include requirements on other parts of the stack, such as the network layer or the Upper Layer Architecture. In addition, they include protocol aspects of the communication reqirements; these are not judged to be relevant at the Transport Service Access point, and a considered to be internal requirements on the TPX protocol not service

A1 List of Baseline Requirements

Error Control

[R001] Provide the required end-to-end error rate desired by the application.
[R002] Error announcement (i.e. error detection and indication) at the receiver.
[R003] Error control mechanism for higher reliability i.e. data transfer or real-time

Timeliness

[R004] The transport should support bounded delivery of data to the reciever
[R005] Global time, or a common reference time
[R006] User indication for failure to transmit a message within its agreed time frame.
[R007] Real-time error sensitive information

Error Recovery and Application Requirements

[R008] The new transport mechanism must support flexible error control .
[R009] Go-back-N.
[R010] Selective retransmission.
[R011] FEC
[R012] No error correction.
[R013] Flexible error detection and recovery mechanisms,

End-to-end Flow Control

[R014] Support continuous transmission of data .
[R015] This requires an open window flow control mechanism where the sender does not wait to send (rate-based flow control).
[R016] The receiver must process the received data stream in real-time and provide the buffering implicitly agreed upon during connection establishment
[R017] It is also important that the operating system provides appropriate response.

Profile and Class of Service Selection

[R018] The transport service and protocol should be flexible enough, to support the full range of multimedia applications and the wide variety of communication characteristics, and work over different high performance networks.
[R019] The idea of protocol selection .
[R020] protocol matrix
[R021] User semantics for selection of protocol matrix, rather than network one.

Other Issues

[R022] Transport support for heterogeneity .
[R023] Transport support for security .

Implementation Philosophy

[R024] The reduction in the copying can be achieved by adopting a high performance buffer architecture which utilises a shared buffer mechanism between the ULA and the lower levels.

Quality of Service

[R025] Quality of Service parameters
[R026] Indication that the QOS has degraded
[R027] Communications sub-system to guarantee QOS is maintained throughout transmission.
[R028] Users QOS semantics incorporating mixed QOS on parameters.

Transport level QOS Parameters

[R029] Throughput .
[R030] End-to-end delay .
[R031] Jitter
[R032] Error rate.
[R033] Priority.
[R034] QOS tolerance levels for each QOS parameter.
[R035] The requested parameters should then undergo full end-to-end negotiation.
[R036] The finally agreed QOS should then be guaranteed for the duration of the connection
[R037] QOS degradation indication.

Throughput and End-to-End Delay

[R38] These bandwidth requirements may be reduced with the use of appropriate compression techniques at the expense of latency.
[R39]Also, late data in continuous media streams is as serious a problem as lost data, and therefore the new service must attempt to provide the required level of error control [without resorting to retransmission, due to the timeliness requirement.

Jitter

[R040] Jitter correction must be offered at the receiver.
[R041] Jitter detection is done within the network.

Priority

[R043] Multiple levels of priority for multimedia traffic and the pre-emption of network resources for high priority applications.
[R044] The transport service must be flexible enough to allow the transport service user to attach relative priorities across end system connections .
[R045] It should be possible that lower priority connection packets will be dropped first before high priority packets.
[R046] Real-time systems require guaranteed delivery of high priority control messages such as a control message, high priority out of band service required.
[R047] The user should just have to deal with packet errors rate at the TSAP.

Guaranteed QOS

[R048] Reservation of transport and network resources via a connection-mode service .
[R049] Transactions QOS support
[R050] Applications that require any QOS higher than best-effort [R050] (offered by current networks) must request resources be allocated at the end system and in the network
[R051] suitable reservation scheme which is able to set up and guarantee network resources such as bandwidth and end-to-end delay between network subscribers.

Dynamic QOS Control

[R052] Dynamically alter the QOS of a connection while it is active.
[R053] QOS degradation indication.
[R054] QOS negotiation and renegotiation at the TSAP.

Group and Multicast Support

[R055] The applications study [Williams:91] highlights the importance of group and multicast
[R056] Services as characteristics of distributed multimedia communications.

[R057] Note that multicast is important for hard real time data transfer.
[R058] Bulk data transfer .
[R059] Transaction .
[R060] Continuous media services .
[R061] Consequently, the underlying transport mechanism .
[R062] Network provider support for multicast.
[R063] Group connection management at the TSAP
[R064] Multicast addressing respectively at the NSAP.
[R065] Connection management flexibility for diverse data transfer requirements
[R066] From a simplex connection.
[R067] Complex connection configurations support
[R068] A multicast service must provide the same QOS associated as with the more traditional point-to-point network access services, in terms of selection of end-to-end delay, jitter, throughput and reliability.

Multicast Services

[R069] Multicast (one-to-many) communications.
[R070] Broadcast (one- to-all) communications.
[R071] Multi-collect (many-to- one) communications.
[R072] Conference (typically many-to-many in video conferencing configuration)
[R073] Un-acknowledged multicast mode.
[R074] Acknowledged mode.
[R075] Synchronised mode - This mode may be used to implement atomic multicast.
[R076] Group names support for transparancy
[R077] Collection and distribution policy

Orchestration

[R078] Event-driven synchronisation.
[R079] Continuous synchronisation.
[R080] The transport must support non-multiplex connections at the TSAP
[R081] The ability to start related continuous media data flows precisely together.
[R082] The ability to create [R82] related connections with compatible QOS.
[R083] The ability to monitor the on-going temporal relationship between related connections

Enhanced Transport Services

[R084] continuous media service.
[R085] transaction oriented service.
[R086] Bulk data transfer service.
[R087] These transport services should operate over connection-oriented.
[R088] Connectionless services.
[R089] Error control for assured communications
[R090] Unacknowledge capabilities.
[R091] Continuous media requires new transport mechanisms [R92] to [R103] inclusive
[R092] High bandwidth.
[R093] A flexible approach to error recovery; including a form of FEC.
[R094] End-to-end rate based flow control mechanisms.
[R095] Mapping application requirements onto the communications sub-system.
[R096] QOS degradation indication.
[R097] QOS negotiation and re-negotiation.
[R098] Transport and network resource allocation to offer guaranteed end-to-end QOS.
[R099] Single and multiple streams.
[R100] Orchestration.
[R101] Group/multicast.
[R102] Isochronous delivery mechanisms.
[R103] Non multipled Simplex connections.
[R104] Reliable request/response
[R105] Interaction versus unreliable .
[R106] Maintaining Connections with Long Lifetimes:.
[R107] Reliable, retransmission of corrupted data is required.

A2 OSI 95 Transport Service Requirements

In this section we list the TPX baseline service requirements which are used to evaluate the TPX TS specification [Danthine,92b] against. This list does not include any non-service (see A3) requirements such as protocol related requirements or ULA issues.

Error Control

[R001] Provide the required end-to-end error rate desired by the application.
[R002] Error announcement (i.e. error detection and indication) at the receiver.
[R003] Error control mechanism for higher reliability i.e. data transfer or real-time

Timeliness

[R004] The transport should support bounded delivery of data to the reciever[R005] Global time, or a common reference time
[R005] Global time, or a common reference time
[R006] User indication for failure to transmit a message within its agreed time frame.
[R007] Real-time error sensitive information

Error Recovery and Application Requirements

[R008] The new transport mechanism must support flexible error control .
[R012] No error correction.
[R013] Flexible error detection and recovery mechanisms,

End-to-end Flow Control

[R014] Support continuous transmission of data .

Profile and Class of Service Selection

[R018] The transport service and protocol should be flexible enough, to support the full range of multimedia applications and the wide variety of communication characteristics, and work over different high performance networks.
[R019] The idea of protocol selection .
[R021] User semantics for selection of protocol matrix, rather than network one.

Other Issues

[R022] Transport support for heterogeneity .
[R023] Transport support for security .

Quality of Service

[R025] Quality of Service parameters
[R026] Indication that the QOS has degraded
[R027] Communications sub-system to guarantee QOS is maintained throughout transmission.
[R028] Users QOS semantics incorporating mixed QOS on parameters.

Transport level QOS Parameters

[R029] Throughput .
[R030] End-to-end delay .
[R031] Jitter
[R032] Error rate.
[R033] Priority.
[R034] QOS tolerance levels for each QOS parameter.
[R035] The requested parameters should then undergo full end-to-end negotiation.
[R036] The finally agreed QOS should then be guaranteed for the duration of the connection
[R037] QOS degradation indication.

Jitter

[R040] Jitter correction must be offered at the receiver.

Priority

[R043] Multiple levels of priority for multimedia traffic and the preemption of network

resources for high priority applications.
[R044] The transport service must be flexible enough to allow the transport service user to attach relative priorities across end system connections .
[R045] It should be possible that lower priority connection packets will be dropped first before high priority packets.
[R046] Real-time systems require guaranteed delivery of high priority control messages such as a control message, high priority out of band service required.
[R047] The user should just have to deal with packet errors rate at the TSAP.

Guaranteed QOS

[R048] Reservation of transport and network resources via a connection-mode service .
[R049] Transactions QOS support
[R050] Applications that require any QOS higher than best-effort [R050] (offered by current networks) must request resources be allocated at the end system and in the network

Dynamic QOS Control

[R052] Dynamically alter the QOS of a connection while it is active.
[R053] QOS degradation indication.
[R054] QOS negotiation and renegotiation at the TSAP.

Group and Multicast Support

[R055] The applications study [Williams:91] highlights the importance of group and multicast
[R056] Services as characteristics of distributed multimedia communications.
[R057] Note that multicast is important for real time data transfer.
[R058] Bulk data transfer .
[R059] Transaction .
[R060] Continuous media services .
[R061] Consequently, the underlying transport mechanism .
[R063] Group connection management at the TSAP
[R065] Connection management flexibility for diverse data transfer requirements
[R066] From a simplex connection.
[R067] Complex connection configurations support
[R068] A multicast service must provide the same QOS associated as with the more traditional point-to-point network access services, in terms of selection of end-to-end delay, jitter, throughput and reliability.

Multicast Services

[R069] Multicast (one-to-many) communications.
[R070] Broadcast (one- to-all) communications.
[R071] Multi-collect (many-to- one) communications.
[R072] Conference (typically many-to-many in video conferencing configuration)
[R073] Un-acknowledged multicast mode.
[R074] Acknowledged mode.
[R075] Synchronised mode - This mode may be used to implement atomic multicast.
[R076] Group names support for transparancy
[R077] Collection and distribution policy

Enhanced Transport Services

[R084] continuous media service.
[R085] transaction oriented service.
[R086] Bulk data transfer service.
[R087] These transport services should operate over connection-oriented.
[R088] Connectionless services.
[R089] Error control for assured communications
[R090] Unacknowledge capabilities.
[R091] Continuous media requires new transport mechanisms [R92] to [R103] inclusive
[R092] High bandwidth.
[R093] A flexible approach to error recovery; including a form of FEC.
[R096] QOS degradation indication.

[R097] QOS negotiation and re-negotiation.
[R099] Single and multiple streams.
[R101] Group/multicast.
[R102] Isochronous delivery mechanisms.
[R103] Non multiplexed Simplex connections.
[R104] Reliable request/response
[R105] Interaction versus unreliable .
[R106] Maintaining Connections with Long Lifetimes:.
[R107] Reliable, retransmission of corrupted data is required.

A3 Non OSI 95 Transport Service Requirements

In this section includes the communication requirements which were not considered in the evaluation section of this deliverable. These non TPX service requirements are only listed here for completeness. Most requirements are associated with ULA functionality such as orchestration, the network layer or transport protocol issues.

Error Recovery and Application Requirements

[R009] Go-back-N.
[R010] Selective retransmission.
[R011] FEC

End-to-end Flow Control

[R015] This requires an open window flow control mechanism where the sender does not wait to send (rate-based flow control).
[R016] The receiver must process the received data stream in real-time and provide the buffering implicitly agreed upon during connection establishment
[R017] It is also important that the operating system provides appropriate response.

Profile and Class of Service Selection

[R020] protocol matrix

Implementation Philosophy

[R024] The reduction in the copying can be achieved by adopting a high performance buffer architecture which utilises a shared buffer mechanism between the ULA and the lower levels.

Throughput and End-to-End Delay

[R38] These bandwidth requirements may be reduced with the use of appropriate compression techniques at the expense of latency.
[R39]Also, late data in continuous media streams is as serious a problem as lost data, and therefore the new service must attempt to provide the required level of error control [without resorting to retransmission, due to the timeliness requirement.

Jitter

[R041] Jitter detection is done within the network.

Guaranteed QOS

[R048] Reservation of transport and network resources via a connection-mode service .
[R051] suitable reservation scheme which is able to set up and guarantee network resources such as bandwidth and end-to-end delay between network subscribers.

Group and Multicast Support

[R062] Network provider support for multicast.
[R064] Multicast addressing respectively at the NSAP.

Orchestration

[R078] Event-driven synchronisation.
[R079] Continuous synchronisation.
[R080] The transport must support non-multiplex connections at the TSAP
[R081] The ability to start related continuous media data flows precisely together.
[R082] The ability to create [R82] related connections with compatible QOS.
[R083] The ability to monitor the on-going temporal relationship between related connections

Enhanced Transport Services

[R093] A flexible approach to error recovery; including a form of FEC.
[R094] End-to-end rate based flow control mechanisms.
[R095] Mapping application requirements onto the communications sub-system.

[R098] Transport and network resource allocation to offer guaranteed end-to-end QOS.
[R100] Orchestration.
[R101] Group/multicast.

Appendix B

Abbreviations

AAL	ATM Adaptation Layer
ALS	Application Layer Structure
ANSA	Advanced Networked Systems Architecture
ASE	Application Specific Element
ATM	Ayschronous Transfer Mode
BAP	Broadband Application Part
B-ISDN	Broadband ISDN
BCOB	Broadband Connection Oriented Bearer service
BCLB	Broadband Connectionless Bearer service
CBDS	Connectioless Broadband Data Service
CO	Connection Oriented
CL	Connectionless
CBR	Constant Bit Rate
CLP	Cell Loss Priority
CPCS	Common Part Convergence Sublayer
DQDB	Distributed Queue Dual Bus
ETS	The TPX Peer to Peer Enhanced Connection-mode Transport Service
ETSI	European Telecommunications Standards Institute
ISCP	ISDN Signalling Control Part
JPEG	Joint Picture Expert Group
NNI	Network Node Interface
ODP	Open Distributed Processing
OS	Operating System
OSI	Open System Interconnection
QOS	Quality of Service
QOS-A	Quality of Service Architecture
SMDS	Switched Multimegabit Data Service
SNI	Subscriber Network Interface
SSCS	Service Specific Convergence Sublayer
ST-II	STream Protocol (TCP/IP)\
SRP	Session Reservation Protocol
TC	Transport Connection
TS	Transport Service
TSAP	Transport Service Access Point
NSAP	Network Service Acess Point
ULA	Upper Layer Architecture
UNI	User Network Interface
UPC	User Parameter Control
VBR	Variable Bit Rate
VCC	Virtual Channel Connection
VPC	Virtual Path Connection
XTP	eXpress Transfer Protocol

References

[Anagnostou,91] Anagnostou, M.E. et al, "Quality of Service Requirements in ATM-Based B-ISDNs", *Computer Networks and ISDN Systems*, Vol 14, No 4, May 1991, pages 197-204.

[Anderson,91] Anderson, D.P., R.G. Herrtwich, and C. Schaefer. "SRP: A Resource Reservation Protocol for Guaranteed Performance Communication in the Internet", *Internal Report* University of California at Berkeley, 1991.

[ANSI,91] ANSI, "High Speed Transport Protocol (HSTP) Specification", *X3S3.3/91-264*, American National Standards Institute, September 1991.

[APM,89] APM Ltd. "The ANSA Reference Manual Release 01.01", Architecture Projects Management Ltd., Poseidon House, Castle Park, Cambridge, UK, July 1989.

[Baguette,92a] Y.Baguette "Enhanced Transport Service Specification TPX", Technical Report OSI95/ULg/A/22/TR/R, ISO/IEC JTC1/SC6 N7323, University of Liege, Liege, Belgium, May 1992.

[Baguette,92b] Y.Baguette. "TPX Specification", OSI95/ Deliverable ULg-5/R/V1, University of Liege, Liege, Belgium , October 1992.

[Ball,90] Ball, F., D. Hutchison, A.C. Scott, and W.D. Shepherd. "A Multimedia Network Interface." *3rd IEEE COMSOC International Multimedia Workshop (Multimedia '90),* Bordeaux, France, Nov 1990.

[Boerjan,91] Deliverable Bell-1: ATM LLC Assessment, 1991 Joost Boerjan Esprit project OSI95

[Berrocal,91] Berrocal J, A. Campbell , G. Coulson, F. Garica, D. Hutchison and A Linares. "Assessment of Synchronisation Functions", OSI95/LANC/C2/05/TR/P/V1, MPG-92-33, Computing Department, Lancaster University, Bailrigg, Lancaster LA1 4YR, UK. January,1992.

[Blair,91] Blair, G., D. Hutchison, W.D. Shepherd, H. Leopold. *"State of the Art in Distributed Multimedia Computing"*, OSI95/Deliverable LANC-1/P MPG-91-26, Computing Department, Lancaster University, Bailrigg, Lancaster LA1 4YR, UK. August 1992.

[Birman,82] Birman, K.P., and T.A. Joseph. "Exploiting Virtual Synchrony in Distributed Systems." *Operating System Review Op. Sys. Review Vol. 21 No. 5,* pp 123-138.

[Campbell,92a] Campbell, A., G. Coulson, F. Garcia and D. Hutchison. "A Continuous Media Transport and Orchestration Service", proceedings *ACM SIGCOMM '92,*, Baltimore, USA, October 1992.

[Campbell,92b] Campbell, A., Coulson G., Garcia F., D. Hutchison and H. Leopold "Contribution to the new ISO Work Item: Key Issues in Distributed Multimedia Communications", ISO/IEC JTC1/SC21/WG1 San Diego, *International Standards Organisation*, UK, November 92

[Campbell,92c] Campbell, A., G. Coulson, F. Garcia and D. Hutchison. "Orchestration Services for Distributed Multimedia Synchronisation", to appear: A. Danthine, O. Spaniol, eds, *High Perfomance Networking,* Elsevier Science Publishers (North Holland), Amsterdam, MPG-90-39 Computing Department, Lancaster University, Bailrigg, Lancaster LA1 4YR, UK. August 1992.

[Campbell,92d] Campbell, A., G. Coulson, F. Garcia, D. Hutchison and H. Leopold. "Integrated Quality of Service for Multimedia Communications", t*o be presented at IEEE INFOCOM'93, San Francisco, USA*, April 93, MPG-90-32, Computing Department, Lancaster University, Bailrigg, Lancaster LA1 4YR, UK. June, 1992.

[Campbell,92e] Campbell, A., G. Coulson, F. Garcia and D. Hutchison. "Resource Management in Multimedia Communication Stacks", *to be presented at the 4 th IEE Conference on Telecomunication, Manchester, UK,* April 93, MPG-90-40, Computing

Department, Lancaster University, Bailrigg, Lancaster LA1 4YR, UK. November, 1992.

[Campbell,92f] Campbell, A., Coulson G., Garcia F., D. Hutchison and H. Leopold. "A Suggested Quality of Service Architecture for Multimedia Communications", ISO/IEC JTC1/SC21/WG1 N1201*, International Standards Organisation*, UK, October 1992

[Clark,92] Clark, D.D., Shenker S., and L. Zhang, "Supporting Real-Time Applications in an Integrated Services Packet Network: Architecture and Mechanism", *Internal Report, Laboratory for Computer Science, MIT, USA,* January 1992.

[Coulson,90] Coulson, G., G.S. Blair, N. Davies, and A. Macartney. "Extensions to ANSA for Multimedia Computing", *To appear in Computer Networks and ISDN Systems* MPG-90-11, Computing Department, Lancaster University, Bailrigg, Lancaster LA1 4YR, UK. Oct 25 1990.

[Dabbous,92] Dabbous, W. and J. Bolot. "*Study of congestion avoidance mechanisms*", OSI95/Deliverable INRIA-5/R/V, /OSI95/A11/v1, INRIA, Unite de Recherche de Sophia-Antipolis, 2004 Route des Lucioles - BP109, F-06561 Valbonne Cedex, France, June, 1992.

[Dark,91] Dark, P., D. Hutchison and Shepherd, W.D., "Remote Procedure Calls and Distributed Multimedia Systems", OSI95/Deliverable LANC-2/P , MPG-91-26, Computing Department, Lancaster University, Bailrigg, Lancaster LA1 4YR, UK. October 1991.

[Danthine,92a] Danthine, A. "A New Transport Protocol for the Broadband Environment", OSI95/ULg/A/14/TS/P, presented at *IFIP Workshop on Broadband Communication*, University of Liege, Liege, Belgium, January 1992.

[Danthine,92b] Danthine, A., Y.Baguette, G. Leduc and L. Leonard. "The Enhanced Connection-mode Transport Service of OSI95", Technical Report OSI95/ULg/A/24/TR, University of Liege, Liege, Belgium June 1992.

[Davies,91] Davies, N., G. Coulson, N. Williams, and G.S. Blair. "Experiences of Handling Multimedia in Distributed Open Systems", *To be presented at SEDMS '92* Computing Department, Lancaster University, Bailrigg, Lancaster LA1 4YR, UK. November 1991.

[Deloddere,92] Deloddere, D., P. Reynders and P. Verbeeck,"Architecture and Implementation of Connectionless Server for B- ISDN", ISS 1992, Yokohama.

[de Prycker,91] de Prycker, M., "Asynchronous Transfer Mode: Solution for Broadband ISDN", ISBN 0-13-053513-3, Ellis Horwood, New York, 1991.

[Feldmeier,92] Feldmeier, D. "Architectural Concepts for High Speed Communications Systems" *Internal Report,* Bellcore, Morristown, NJ, USA, February 1992.

[Ferrari,90] Ferrari, D., and D.C. Verma, "A Scheme for Real-Time Channel Establishment in Wide-Area Networks", *IEEE JSAC*, Vol 8, No 3, April 1990, pages 368-377.

[Govindan,91] Govindan, R., and D.P. Anderson. "Scheduling and IPC Mechanisms for Continuous Media." *Thirteenth ACM Symposium on Operating Systems Principles,* Asilomar Conference Center, Pacific Grove, California, USA, SIGOPS, Vol 25, Pages 68-80.

[Greaves,91] Greaves, D.J, L.J. French and D.R. Milway. “Specification of a High Performance Worstation ATM Interface” OSI9/ORL/TN/B3/12/P/V1, Olivetti Research Limited, Old Addenbrookes Site, 24a Trumpington Street, Cambridge, May 1992.

[Hehmann,91] Hehmann, D.B., R.G. Herrtwich, W. Schulz, T. Schuett, and R. Steinmetz. "Implementing HeiTS: Architecture and Implementation Strategy of the Heidelberg High Speed Transport System" *Second International Workshop on Network and Operating System Support for Digital Audio and Video,* IBM ENC, Heidelberg, Germany, 1991.

[Henshall,88] Henshall, J., and S. Shaw, "OSI Explained: End-to-end Computer Communication Standards", ISBN 07458-0253-2, Ellis Horwood, New York, 1988.

[I.371] Traffic Control and Congestion Control in B-ISDN CCITT Rec. I.371 Geneva 1992

[I.37Y] Network capabilities to support multimedia services. CCITT Draft Rec. I.37Y Dec'91

[I.350] General aspects of Quality of Service and network performance in digital networks,

including ISDNs CCITT Rec. I.350

[ISO 8072,86] ISO/EC JTC1 Information Processing Systems - Open Distributed Processing *"Trnsport Service Definition" ,* International Standards Organisation, 1986.

[ISO 8072 ADD1,86] ISO/EC JTC1 Information Processing Systems - Open Distributed Processing *"Trnsport Service Definition - Addendum 1: Connectionless-Mode Transmission"* , International Standards Organisation, 1986.

[ISO,89] ISO. "Basic Reference Model of Open Distributed Processing", *Working document on structures and functions,* ISO/IEC JTC1/SC21 N4022, 10 November 1989.

[ISO,92a] ISO/EC JTC1, "User Requirements for Quality of Service", ISO/IEC JTC1/SC21/WG1 N1146*, International Standards Organisation*, UK, May 1992.

[ISO,92b] ISO/EC JTC1, "Quality of Service Framework - Outline", ISO/IEC JTC1/SC21/WG1 N1145*, International Standards Organisation*, UK, March 1992.

[ISO,92c] ISO/EC JTC1 SC21 N3287, Information processing systems -- Open System Interconnection -- Draft Addendum to ISO 7498, " Multipeer Data Transmission" *International Standards Organisation*, 1992.

[ISO,92d] ISO/EC JTC1/SC6 N7309, Draft Guidelines for Enhanced Communications Functions and Facilities" *International Standards Organisation*, 1992.

[ISO 8348,92] ISO/EC JTC1, "Network Service Definition"*, International Standards Organisation*, 1992.

[Landegem,91] Van Landegem and R. Peschi, "Managing a Connectionless Virtual Overlay Network on Top of an ATM Network", ICC'91.

[Leonard,92] L.Leonard "Enhanced Transport Service Specification", Technical Report OSI95/Deliverable /ULg-4/R/V3, University of Liege, Liege, Belgium October 1992.

[Leopold,92a] Leopold, H., Blair, G., Campbell, A., Coulson, G., Dark, P., Garcia, F., Hutchison, D., Singer, N., and N. Williams, "Distributed Multimedia CommunicationsSystem Requirements", *OSI95/Deliverable ELIN-1/C/V3,* Alcatel ELIN Research, A-1210 Vienna, Ruthnergasse 1-7, Austria, April 1992.

[Leopold,92b] Leopold H., A. Campbell , D. Hutchison and N. Singer, N. "Towards an Integrated Quality of Service Architecture (QOS-A) for Distributed Multimedia Communications", to appear A. Danthine, O. Spaniol, eds, *High Perfomance Networking,* Elsevier Science Publishers (North Holland), Amsterdam, OSI95/ELIN/D3/02/TR/P/V1 Alcatel Austria - ELIN, Forschungszentrum G.m.b.H, Ruthnergasse 1-7, A-1210 Wien, Austria.

[Leopold,92c] Leopold H., A. Campbell , D. Hutchison and N. Singer, N. "Will B-ISDN Services Meet the Needs of Distributed Multimedia Communications", *to be presented at the 4 th IEE Conference on Telecomunication, Manchester, UK,* April 93, OSI95/ELIN/D3/02/TR/P/V1 Alcatel Austria - ELIN, Forschungszentrum G.m.b.H, Ruthnergasse 1-7, A-1210 Wien, Austria.

[Leopold,92d] Leopold, H., "Overview of B-ISDN AAL Protocols and Services" Summary of ETSI Standardisation work, OSI95/ELIN/D3/04/TN/R/V2, Alcatel Austria - ELIN, Forschungszentrum G.m.b.H, Ruthnergasse 1-7, A-1210 Wien, Austria, October 1992

[Nicolaou,91] C. Nicolaou, "A Distributed Architecture for Multimedia Communication Systems", *Ph.D Thesis*, University of Cambridge, Computer Laboratory, April 1991.

[PE,92] Protocol Engines, Inc., "XTP Protocol Definition - Revision 3.6", January 1992

[Race D510] General aspects of Quality of Service and System Performance in IBC RACE IBC Common Functional Specification D510 (based on RACE QOSMIC project)

[Race 2044] Signalling Requirements Baseline Document RACE 2044 Magic, WP3 September 1992

[Rozenblad,92] Interconnection of LANs/802.6 Customer Premises Equipments (CPEs) via SMDS on Top of ATM: a case description.", W. Rozenblad, B. Li, R. Peschi, IFIP 92, Liege.

[Shepherd,91] Shepherd, W.D., D. Hutchison, F. Garcia and G. Coulson. "Protocol Support for Distributed Multimedia Applications." *Second International Workshop on Network and Operating System Support for Digital Audio and Video,* IBM ENC, Heidelberg, Germany, Nov 18-19 1991.

[Singer,92] Singer, N., G. Blair, A. Campbell, G. Coulson, K. Frimpong-Ansah, F. Garcia, D. Hutchison and H. Leopold."OSI/ULA - ODP Architecture", *OSI95/Deliverable ELIN-3/P/V1,* Alcatel ELIN Research, A-1210 Vienna, Ruthnergasse 1-7, Austria, April 1992.

[Sluman,92] Sluman, C., "Quality of Service in Distributed Systems", BSI/IST21/-/1/5:33, *British Standards Institute*, UK, October 1991.

[Topolcic,90] Topolcic, C. "Experimental Internet Stream Protocol, Version 2 (ST-II)", *Internet Request for Comments No. 1190* RFC-1190, October 1990.

[Van Achter,92] D. Van Achter. "Routing approaches for remote LAN bridging over SMDS using the spanning tree", Interworking '92, Berlin.

[Williams,91] Williams, N., G.S. Blair. "Distributed Multimedia Application Survey", *Internal Report N. MPG-91-11.* Computing Department, Lancaster University, Bailrigg, Lancaster LA1 4YR, UK. April 1991.

[Wolfinger,91] Wolfinger, B., and M. Moran. "A Continuous Media Data Transport Service and Protocol for Real-time Communication in High Speed Networks." *Second International Workshop on Network and Operating System Support for Digital Audio and Video,* IBM ENC, Heidelberg, Germany, 1991.