

QoS ENHANCEMENTS AND THE NEW TRANSPORT SERVICES

André Danthine, Olivier Bonaventure¹, Yves Baguette²,
Guy Leduc³ and Luc Léonard²

¹ Research Assistant of the Université de Liège

² Research Assistant of the F.N.R.S. (National Fund for Scientific
Research, Belgium)

³ Research Associate of the F.N.R.S.

Institut d'Electricité Montefiore, B28

Université de Liège,

B-4000, LIEGE, Belgium

E-mail: danthine@vm1.ulg.ac.be

INTRODUCTION

The transport protocols TCP and TP4 have been designed in the late seventies at a time when the network environment was essentially based on the switched telephone network and on leased-lines and when the packet switching was an emerging concept.

In ISO, the transport service was seen at the beginning as a unique service based on connection mode and, in this framework, one protocol class, TP4, was designed for the poor network service environment [Dan 90]. TP4, as well as TCP, were designed to be able to recover from the worst situation in terms of packet losses and packet disorders.

In the last ten years, we have seen a tremendous change in the communication environment. The LANs have drastically changed the scene and the high-speed LANs have accelerated the trend. More recently, the MANs have opened new possibilities not to mention the broadband ISDN which is expected to be deployed starting in the middle of this decade.

Not only the communication environment has drastically changed during the eighties, but we have also faced a drastic change in application requirements of which we will mention only the client-server paradigm and the multimedia-based applications.

In this paper we will first justify the need of new transport services and protocols based on the analysis of the consequences of the changes we have been facing

- in network performance
- in network services
- in application requirements.

We will, in the second part, present an enhancement of the QoS semantics in order to be able to specify the new transport service in a way compatible to the requirements of the transport service users. In our proposed semantics, a QoS parameter is seen as a structure of three values, respectively called "compulsory", "threshold" and "maximal quality", all these values being optional. Each one expresses a contract between the service users and the service provider. This means, and this is the most fundamental difference with the present situation, that the service provider is now submitted to some well defined duties.

The last part of the paper will be a brief presentation of examples coming from the OSI 95 transport services.

CHANGE IN NETWORK PERFORMANCE

Let us first point out the robustness of the OSI Reference Model which has been able to integrate the LANs in its layered architecture [Dan 89] and is very likely to integrate the B-ISDN.

The basic contributions of the LANs have been a drastic improvement of the network performance in terms of data rate as well as in terms of bit error rate (BER) and packet error rate. The error recovery was not any more a basic goal of the DL layer and this explained the tremendous success of the connectionless-mode.

The LANs have not only introduced an increase of the global communication capability due to the high value of their data transmission rate but, more important, they have offered the user an access data rate several orders of magnitude greater than what was available before.

If, as usual, the first implementations of the LAN controllers were not always able to offer the maximum access rate theoretically available at the MAC level [DHH 88a], the market, in about five years, was able to offer well-designed and well-implemented LAN interfaces.

It did not take long to realise that an access rate available at the MAC or LLC levels was far from being available above the transport layer [DHH 88a][Svo 89a, 89b][HeS 89][CCC 91]. This is not surprising, because opening a highway does not by itself change drastically the speed of the bicycles.

How to Improve ?

From the observations made in the second half of the eighties about the transport performance on top of LANs, detailed studies of mechanisms and implementations have been done [CJR 89], [GKW 89], [Mei 91], [JSB 90], [Zitt 91] and ended up in two different schools of thought.

The first school advocated the design of new transport protocols such as NETBLT [CLZ 87], VMTP [Che 88],[ChW 89], SNR [NRS 90] and XTP [Ches 89], [Wha 89], [PE 91].

The second school is putting the burden of the poor transport performance to the bad implementation quality and is claiming the capability of existing or slightly modified transport protocols to better handle the performance available at the MAC level [CJR 89].

Both schools have arguments and we would like here to discuss them, keeping in mind that these LANs do not only offer a better access rate, they are also characterised by a very low probability of corrupted packets by transmission errors.

Mechanisms and Implementations

We claim that there is a basic interest to adapt or to modify the protocol mechanisms when the network service is drastically changing. The following example will support our claim.

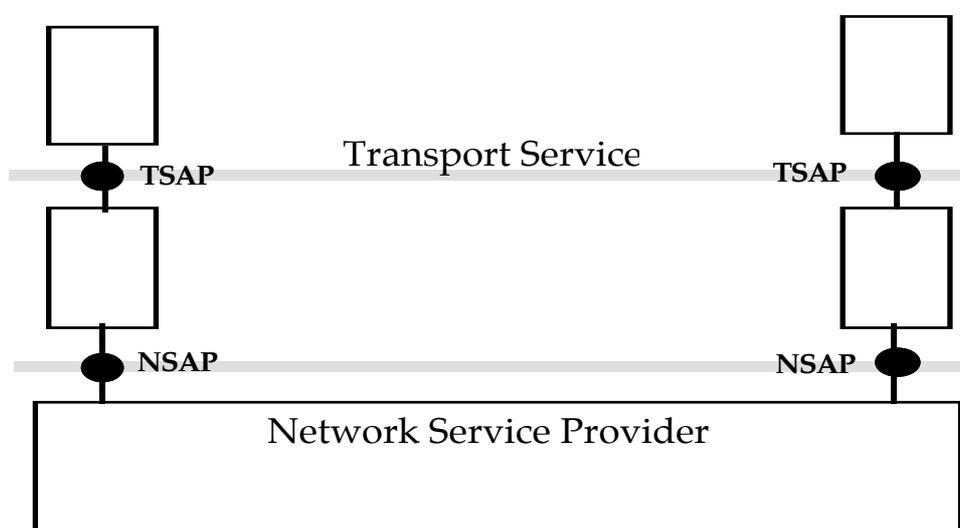


Figure 1. The basic OSI reference model

The figure 1 represents a simplified view of the basic OSI reference model and we

assume, on the NSAP-to-NSAP association, a guaranteed throughput and a fixed round trip time. On top of such an well-defined N-service, we assume an ideal implementation of the transport entities which does not introduce any latency to process TSDUs into NSDUs. This does not mean however that the throughput and the round-trip delay on the TSAP-to-TSAP association, say a transport connection, will be equivalent to the values at the network level.

The product of “throughput x round-trip delay” - also called “bandwidth x delay” - is becoming, with high speed networks, an important characteristic to be dealt with. The value of such a product has evolved from 1 Kb (2 Kbps x 500 msec) to 600 Kb (10 Mbps x 60 msec), not to mention the very high speed environment. If the window associated with a transport connection does not allow the pipelining of enough megabits, the transport entity will be obliged to stop transmitting by lack of acknowledgements and of flow control credits. Therefore the performance at the transport level will be much lower than the performance at the network level **even with an ideal implementation of the transport layer**.

On the other hand, if protocol mechanisms may have such an influence on the performance even with an ideal implementation, it is also true that some mechanisms may induce very poor implementations. For instance a time-out per packet is not the best mechanism from an implementation point of view.

With a transport entity trying to send as many packets as possible, a lot of ACKs to handle represents a burden and the classical ARQ mechanism has therefore to be reviewed to reduce the penalties on performance.

Last but not least, there remain implementation problems which are independent of any protocol mechanism such as buffer handling and interaction with the operating system [CJR 89], [MaB 91].

There exist of course many choices of implementation from VLSI to pure software, more or less integrated in the operating system, using additional dedicated computing power or not [GKW 89], [JSB 90], [Zitt 91].

The syntaxes of the PDU are not neutral with respect to the implementation results. Variable fields may save bandwidth but are more difficult to “siliconize”. The place of the CRC field in the PDU is also important if hardware implementation with insertion and removal “on the fly” is envisaged.

With the availability of microprocessors of more than 10 Mips, with the possibility to use parallelism in the implementation of the transport entities, with the possible use of dedicated chips, we have a wide variety of possible implementations of the transport entities, all of them attempting to offer a transport service with a minimum drain on the host capabilities.

The conclusion to this section on the link between mechanisms and implementations is that even if it is true that the implementation of the protocol mechanisms is not always the most time consuming part of the transport activity, it will not be bad to have mechanisms which are well-aligned with the objectives and services of the protocol and which will ease the implementation problems. This does not prevent of course a careful choice of the syntaxes of the PDUs and an in-depth analysis of the memory management.

The Congestion Problem

During the eighties, we have not only seen the widespread usage of the LANs and recently of the high-speed LANs. We have also seen the building of very big internet-based networks following the efforts supported by the US DoD.

Such an internet network may give rise to very difficult problems of congestion which are almost impossible to tackle with the error recovery mechanism which has been introduced in TCP and TP4. The correcting action of the transport protocol in order to recover the lost packets due to the congestion is based on retransmission on time-out. These retransmissions contribute to an increase of the offered load and therefore to increase the congestion problems.

This is why congestion avoidance mechanisms which try to operate the network at the knee of the response time curve [RaJ 88], [ChJ 89] have to be preferred to the congestion control approach. Time-out retransmission must be associated with a reduction of the load on the network [Jai 86].

In the congestion problems, we have also two possible viewpoints. Either the mechanisms are aiming at the protection of the provider of the service, i.e. the network itself, in order to avoid a complete collapse of communication capability. Either, they are aiming at offering, to a service user, the best service possible at any given time. We do not have today a

clear trend between these two views.

Congestion control may be exercised by the network service user, and the slow-start method in TCP [Jac 88] is the best known way to do it, or it may be exercised by the network service itself through some kind of rate control.

In any case, we have to deal with some reduction of the rate exercised either by the network service user or by the network service provider. It is very difficult today to have a clear view of the best solution. There exist a lot of proposals based on different ways to collect the congestion information. The congestion control may indeed be exercised on the basis of the information collected by the user through the observations of the behaviour of the service provider. As measurement bases already proposed, we have the rate of packet loss, the increase of round-trip delay, the occurrence of a timeout on an unacknowledged packet. Another proposal intends to base congestion control on information provided by the network in the form of loss-load curves [WiC 91]. The provider of service may also be involved by setting a special bit in packets going through a congested area. This piece of information is used to report internal congestion to the transport recipient which is able to act on the flow control mechanism in order to ease the congestion problem [RaJ 88], [RaJ 90].

The congestion problem is much more difficult to handle with a connectionless network service. In [Zha 87a,87b], it has been proposed to introduce, on a network association, the idea of a flow which will be used to handle a stream of packets. Setting up a flow, transmitting on a flow and terminating a flow appear similar to what we do for a connection but this flow is used only to handle resources on the way and is not associated with any service property. The protocol ST-II [Top 90] is based on the same idea and try to handle the congestion problem by a resource reservation scheme allowing to offer well defined performance. This is also the congestion problem which push XTP to handle the transport and network layers in an integrated transfer layer.

The deployment of corporate communication networks give rise to corporate internet networks where we face the same problem due to congestion of bridges or routers. The high-speed environment will just make things a little more difficult. For instance, if the high throughput goal requires not to use time-out on data packets, it will be impossible to base congestion control on time-out occurrences.

Increasing the resources to avoid packet loss due to congestion may be done by increasing the buffering capability of the intermediate nodes, but this will drastically increase the round-trip delay and may not be considered as a good solution [Jai 90].

The QoS (Quality of Service) in OSI as well as the ToS (Type of Service) in IP [Pos81], [Alm92] are specified in qualitative terms and are completely useless for the congestion control or the congestion avoidance.

For LANs, the probability that a frame be corrupted by a transmission error is very, very low, but the probability of packet loss due to congestion in the access device to a host, is not to be neglected. With integrated LANs through bridges and routers, the problem may get more complex and the congestion may be responsible for packet losses which will jeopardise the quality of the service.

To tackle this congestion problem, in the new high speed environment, we believe it is essential to have first a better specification of the service user requirements through more adequate QoS parameters associated with a well defined semantics and second a better way to specify the responsibilities of the service users on one side and of the service provider on the other side.

CHANGE IN NETWORK SERVICE CAPABILITY

If it has been possible to integrate the basic functionalities of the LANs in the OSI Reference Model, there exist subnetwork services which are available neither at the network level nor at the transport level.

Multicast Service

One of these services is related to the capability of addressing not only a given service access point, but a group of SAPs through multicast or all SAPs through broadcast. As we will see in the next section, many applications have a direct interest to see this multicast service not limited to the subnetwork level but offered at the network level as well as at the transport level. Some of the applications will require more than the connection-less based

multicast and this will again require an enhancement of the QoS.

Synchronous Service

In FDDI, we have two types of services : the asynchronous service and the synchronous service. The synchronous service offered by FDDI I is a service where the average throughput is guaranteed as soon as it has been reserved.

The stream of bits coming from a constant rate source or from a variable rate source may be packetised and sent through an association offering a jitter on latency which does not compromise the re-synchronization at the acceptor [DRB 87]. This approach was applied in 1988 with the BWN, a 140-Mbps FDDI-like network developed in the framework of the ESPRIT project 73 [DHH 88b]. Videoconferencing with 2-Mbps codecs has been demonstrated on the connectionless BWN and this system was able to operate with a network loaded with more than 100 Mbps of background traffic.

With FDDI I and its synchronous service such an approach is possible even with a higher rate and with a guaranteed average throughput. However such a synchronous service is not available today neither at the network service level nor at the transport service level. Such an extension of the service will require at the transport service, a drastic change in the QoS semantics

Broadband Service

For the B-ISDN, even if the pace of the standardisation of the AAL (ATM Adaptation Layer) is slow, the AAL Type 3-4 is offering a service identical or comparable to SMDS or DQDB. It is not excluded that, through the signalling protocol still under development, the connectionless service will be associated with a guaranteed average rate.

For AAL Type 1, an isochronous service on SDU of more than 48 bytes may not be excluded and, if it happens, it is very likely to be associated with a guaranteed average or peak rate or to something in between.

For AAL Type 5, the data orientation is clear but the extension of the QoS is not yet in its final stage.

A Clear Trend

From the discussion of the new network services, it is obvious that we need to be able to introduce in all the lower layers of the OSI Reference Model, better addressing capabilities such as multicast, as well as to find a way to link the performance such as throughput into the QoS not only in qualitative terms but in quantitative terms.

CHANGES IN APPLICATION REQUIREMENTS

Beside the classical applications such as file transfers and conversational remote access to databases, the eighties have seen new orientations in the applications and these new applications have introduced new requirements for communication.

Client-Server Paradigm

The Client-Server paradigm is one of the important changes which took place during the last decade. From a communication point of view, it implies a low latency and an extended addressing capability.

The search of a server may take place by broadcasting or multicasting a request and the blocking character of the request implies the requirement for a low latency. This low latency is difficult to achieve when using the complete OSI stack. A system such as ANSA [Her 88] has chosen a collapsed OSI architecture based on a convergence sublayer, located in the lower part of the layer 7, acting as a transport service and giving access to the SAP of the subnetwork at the layer 2 interface. Others systems have developed better focused transport service and protocol [ChW 89],[MaB 91].

Today, most of the client-servers are implemented on a local basis, but the clear trend to use this new paradigm at the corporate network level will require a re-evaluation of the architecture presently used. If we hope to see the distributed computing systems integrated in the OSI architecture, it is mandatory for the transport service to offer the needed addressing

capabilities as well as a transaction-oriented service with low response time as requests are most often of the blocking type.

Graphics and Colours

We have seen a constant increase in the power of the microprocessors used for the workstations and PCs but an analysis of the evolution puts into evidence that quite an important part of the power has been invested in the user interface and in a more advanced form of presentation of the results to the users. Multiple windowing, graphics usage as well as use of colours have fundamentally changed the friendliness of the workstations and paved the way for the widespread usage of them.

Of course, the volume of data which has to be exchanged in interactions involving windows, graphics and colours, is much more important than for the classical ASCII pages. The user, however, is not ready to accept to have long delays associated with the dynamic of the data presentation. Therefore the higher volume of data to be exchanged through the communication system on one side and the need to achieve what may be called the “client expected latency” on the other side mean that we will have a request for a high data throughput with a still very important bursty character. It is only if the global network service is offering the high value of the throughput associated with an acceptable latency that this kind of system will be usable on a wide area.

Multimedia

By multimedia we mean, as usual, a mixing of text, voice, graphics, image, audio and video. It is however important to make the distinction between two types of multimedia communications.

The videophone is a multimedia terminal as it is able to handle voice and video. But here the terminal is a source of a stream of bits which is very likely to be carried out on a physical or on a virtual circuit from one equipment to another equipment of the same type. The complete OSI architecture is very unlikely to be used as the information exchanged is not represented in machine-processable form, able to be processed, stored, retrieved and disseminated at any time.

The workstation where we will have the text in one window, a fixed image in another window and a video in a third one completed by some audio is an example of equipment requiring the second type of multimedia communication. Here all information which may be processed, stored, retrieved and disseminated are in machine-processable form. Here the video part of such an equipment will be seen as a sequence of data structures which have to be submitted to the presentation device at a rate compatible with the type of presentation we are looking for. As indicated earlier, such a video may be transmitted as a sequence of video frames, each frame being divided into packets at the transport level to be reconstructed at the remote end. The correct reconstruction requires from the transport service a sufficient throughput and a jitter on packet latency which is below some fixed value. Unfortunately, nothing today allows a user of TP4 or of TCP to request a quantitative value for the throughput on its connection, nor to express a limit to the jitter associated with the TSDU transport delay.

With video, the error recovery by retransmission does not work properly as it conflicts with the low jitter requirement. Such a data stream is not flow-controllable and if the transport service is not able to provide the required throughput, it will be reasonable to cancel the connection.

Most of today’s multimedia applications are implemented on a stand-alone workstation. The trend to expand on network in the local area is already there and if some existing systems are working properly, it is very often because the transport service is able, through best effort, to offer a throughput sufficient for the application. This is true for an application implemented on a lightly loaded LAN, but may become less evident when the LAN will be more heavily loaded, when the LAN will be built with bridges and routers and, last but not least, when we will try to have this application running on the wide area internet network. Here, we will face, as we go along these different environments, an increase of the latency and a possible reduction of the throughput.

All applications mentioned in this section require the possibility to associate, to a transport connection, quantitative compulsory QoSs with regard to the throughput, transit delay and delay jitter, and the service provider may have to intervene on the connection if the

performance is not fulfilled.

NEED OF A NEW TRANSPORT SERVICE

In the previous sections, we tried to stress that the improvement of the transport service and protocol does not lie only on the performance, but that it will be necessary to extend the characteristics of the offered services. This opinion was already expressed in [ChW,89] to advocate a new generation of communication systems : *“By ‘new generation’ we mean that the design of protocols, networks and network interfaces should be rethought from the first principles in the context of the new environment. Refining existing designs and tuning conventional approaches in the area are inadequate, given the magnitude of changes that have taken place. At the same time, we need to be extremely careful to identify the key problems of the current standard protocols and come up with a convincing solution that is significantly better than these protocols. That is, we clearly recognise that an incremental improvement over existing protocols, such as the stream-oriented TCP or ISO TP4, would not warrant a new protocol or modifications to the existing protocols, given the enormous investment already in place.”*

It is in complete agreement with this position that started, in October 1990, the ESPRIT II Project, OSI 95¹. The design of a new transport service and a new transport protocol, was the basic goal of the ESPRIT Project OSI 95 which involved Alcatel Bell, BULL, Alcatel Austria, INRIA, Institut National des Télécommunications, Intracom, Olivetti Research, as well as the universities of Madrid, Lancaster and Liège.

THE OSI 95 APPROACH

In the OSI 95 project, the University of Liège, which assumed the technical direction of the project, was responsible for the design and the formal specification of a new transport service and a new transport protocol. Two other tasks were done in parallel in order to better assess the new environment. The first one defined the specific requirements coming from the distributed multimedia applications and ODP systems which may have an impact on the new transport service. The second one defined the data link layer services provided by ATM networks.

The first step of the design of a new transport service has been the enhancement of the QoS (Quality of Service). The term Quality of Service (QoS) is the collective name given to certain characteristics associated with the invocations of (N)-service facilities as observed between (N)-SAPs. The QoS is described in terms of a set of parameters.

Taking into account the previous discussion about network performance, network services and application requirements, we believed in this context that it was essential to define, for the lower layers (4 to 2), a new model of QoS involving a new semantics of the QoS parameters and the definition of new parameters.

The strong requirements of some applications have driven us to define, through the semantics of the QoS, the characteristics of the contract between the service users and the service provider with a well defined behaviour if the requirements cannot be met and it is on the basis of this new semantics of the QoS that the new transport services have been designed [Dan 92a], [Dan 92b], [DBL 92c].

But before presenting our proposal, we will review the situation of the QoS since their introduction in the OSI RM in the beginning of the eighties. We will address mainly the connection-mode service, the problem of the QoS in the connectionless-mode being a subset of the first one.

TYPES OF QoS NEGOTIATIONS

We will restrict our discussion to the peer-to-peer case where the three actors of the negotiation are the calling (service) user, the called (service) user and the service provider. It is however possible to extend the discussion to the 1-to-N, N-to-1 and N-to-N multicast cases.

All negotiations are based on the classical 4-primitives exchange; request, indication,

¹ OSI 95 is the acronym of ESPRIT II Project 5341 “High Performance OSI Protocols with Multimedia Support on HSLANs and B-ISDN”.

response and confirmation.

For some performance parameter such as the throughput, the higher the better. For some other performance parameter such as the transit delay jitter, the smaller the better. For the following discussion we will use the terms “weakening” and “strengthening” a performance parameter to indicate the trend of the modification. Weakening a throughput parameter means reducing its value but weakening a transit delay jitter means increasing its value.

Triangular Negotiation for Information Exchange

In this type of negotiation, the calling user introduces in the request primitive the value of a QoS parameter. This value may be considered as a suggested value because the service provider is free to weaken it as much as it wants before presenting the new value to the called user through an indication primitive. The called user may also weaken the value of the parameter before introducing it in the response primitive. This final value will be included without change by the service provider in the confirm primitive. At the end of the negotiation, the three actors have the same value of this QoS parameter.

Taking into account the freedom for the service provider to weaken the value suggested by the calling user, the service provider will reject directly the request only if it is unable to offer the service whatever the value of the QoS.

The calling user has always the possibility to request a disconnection if it is unsatisfied by the value resulting from the negotiation .

The goal of such triangular negotiation is essentially to exchange information between the three actors.

The ISO Transport Service is using this type of negotiation for the performance-oriented QoS [ISO DIS 8072]. The classes 1 and 3 of the ISO Network Service [ISO 8348] is also based on the same scheme.

Triangular Negotiation for a Bounded Target

In this type of negotiation, the calling user introduces, in the request primitive, two values of a QoS parameter, the target and the lowest quality acceptable. The service provider is not allowed to change the value of the lowest quality acceptable. Here, the service provider is free, as long as it does not weaken it below the lowest quality acceptable, to weaken the target value before presenting the new value of the target and the unchanged value of the lowest quality acceptable to the called user, through an indication primitive. It will be the privilege of the called user to take the final decision concerning the selected value of the target. This selected value of the QoS will be returned by the called user in the response primitive. This selected value will be included without change by the service provider in the confirm primitive. At the end of the negotiation, the three actors have agreed on the value of this QoS parameter.

Here the service provider does not have anymore the possibility to weaken without limit the target value proposed by the calling user, due to the existence of the lowest quality acceptable value. The service provider may have to reject the request if it does not agree to provide a QoS in the requested range.

The called user may also reject the connection attempt if it is not satisfied by the range of values proposed in the indication primitive.

With respect to the target value introduced by the calling user, the only possible modification introduced by the negotiation is the weakening of the target but limited by the lowest quality acceptable value.

The class 2 of the ISO Network Service [ISO 8348] is also based on this scheme.

Triangular Negotiation for a Contractual Value

In this type of negotiation the goal is to obtain a contractual value of a QoS parameter which will bind both the service provider and the service users. Here the calling user introduces, in the request primitive, two values of a QoS parameter, the minimal requested value and the bound for strengthening it. The service provider is not allowed to change the value of the minimal requested value. Here, the service provider is free, as long as it does not weaken it below the minimal requested value, to weaken the bound for strengthening value before presenting the new value of the bound for strengthening and the unchanged value of the minimal requested value to the called user, through an indication primitive. It will be the privilege of the called user to take the final decision concerning the selected requested value.

This selected value of the QoS will be returned by the called user in the response primitive. This selected value will be included without change by the service provider in the confirm primitive. At the end of the negotiation, the three actors have agreed on the value of this QoS parameter.

The service provider may have to reject the request if it does not agree to provide a QoS in the requested range.

The called user may also reject the connection attempt if it is not satisfied by the range of values proposed in the indication primitive.

With respect to the minimal requested value introduced by the calling user, the only possible modification introduced by the negotiation is the strengthening of the minimal requested value but limited by the bound for strengthening value. The service provider may weaken the bound for strengthening and the called user may strengthen the minimal requested value but to the limit accepted by the service provider.

This scheme of negotiation has been used in OSI 95 for two types of requested values, the compulsory QoS value and the threshold QoS value which will be defined later on.

Bilateral Negotiation versus Triangular Negotiation

For some QoS parameters, the negotiation takes place between the two service users, the service provider being not allowed to modify the value proposed by the service user. However, this bilateral character is more formal than real, the service provider having always the possibility to reject the request.

When not only the service provider but the calling service user is not allowed to modify the value proposed by the calling user, the negotiation is reduced to a “take it or leave it” approach. This type of negotiation may have its merits in special situations.

BEST EFFORT QoS VALUE

Coming back to the result of the negotiation for information exchange, it is clear that it is difficult to attach a strong semantics to the resulting value of the QoS parameter. It is only a value agreeable to all three actors. The QoS parameter does not require a permanent monitoring by the service provider because it is not possible to specify a particular behaviour of the service provider if the real value of the QoS parameter is weaker than the agreed value. The service users do not expect such particular behaviour in such a case. They are just expecting the “*best effort*” from the service provider.

In such a loosely defined environment, if a service user introduces, in a request primitive, the value of a QoS parameter, it is not always clear if this suggested value is related to a boundary or to an average value. In the latter case, the measurement sample or the number of SDUs to be considered is often far from obvious. This is not a great problem if no monitoring has to be done.

For negotiation where the service provider is not allowed to weaken the value suggested by the calling user, it is possible for the service provider to reject the request due to the requested QoS value but, in case of acceptance, nothing will be done if the service provider does not reach the QoS value. In this case, the service user is not even informed about the situation by the service provider as no REPORT indication primitive has been defined.

It is therefore not surprising that in many cases, the QoS is expressed in qualitative terms without any specification of a given value. This confirms the lack of relationship between the QoS parameter and a real performance parameter. The only way for the service users to assess the value of a QoS is to monitor it.

If, to operate in a correct way, an application requires a well defined set of performance parameters, the present approach will not be suitable.

The Monitoring and the Protocol Functions.

Monitoring is not the only way to assess the performance of a service provider. In TP4, the error control function based on time-out and retransmission achieves a very low value of the residual error which will be impossible to assess by monitoring on the life time of most connections.

For the data communication, the ISO transport service has been considered as adequate because it was offering a reliable service. The throughput, the transit delay and the transit

delay jitter were not considered as critical factors of performance.

GUARANTEED QoS VALUE

In some situation, it would be nice to be able to have the concept of a guaranteed QoS, especially when the guarantee value result from a negotiation for a bounded target or from a negotiation for a contractual value. The possibility for the service provider of giving such a guarantee is related to the existence of resources associated with some protocol function for allocating the resources and managing it during the connection.

A residual error rate on a connection may be guaranteed by the service provider if enough buffers and numbering space can be allocated to the connection and if it operates an error control function with known property.

A minimum throughput on a connection may be guaranteed by the service provider if each protocol entities can be allocated enough processing resources and if the underlying service provides the corresponding throughput guarantee.

If the allocation of resources is only partial, leaving room for uncertainty, or if the resource allocation is done on a statistical basis or on an overbooking scheme, it is possible to associate some probability measure with the QoS value. The same is true if the resource allocated to a connection may be recover by another connection of higher priority as in ST-II [Top 90] by the introduction of the precedence mechanism.

We may in these cases either speak of guaranteed QoS with a low probability or we may qualify differently the QoS value when it has to be associated with a probability far from 1. It is was has been done in [Par 92] where a distinction is made between a guaranteed service and a predicted service.

WHAT IF ?

If one associates the semantics of guaranteed QoS to the value resulting from a negotiation for a bounded target or from the negotiation for a contractual value, it is clear that the service provider may have to reject the request if it does not agree to provide a QoS in the requested range.

However if the request is accepted, it less clear what will have to done if during the connection the value of the QoS parameter is weaker that the guaranteed value.

If the service provider do nothing in such case, we are back to the best effort.

The service provider may abort the connection as it is not able to maintain the guaranteed value.

The service provider may also indicate to the service user(s) that it cannot maintain the selected value and leave to the service user(s) the responsibility of aborting.

A NEW SEMANTICS FOR THE QoS

In OSI 95, a new semantics for the performance QoS parameters has been introduced. In this semantics, a parameter is seen as a structure of three values, respectively called "compulsory", "threshold" and "maximal quality", all these values being optional. Each one has its own well defined meaning, and expresses a contract between the service users and the service provider.

This means, and this is the most fundamental difference with the best effort, that the service provider is now submitted to some well defined duties, known by each side. In other words, the rules of the game are clear.

The existence of a contract between the service users and the service provider implies that, in some cases, the service users are also submitted to well defined duties also derived from the application environment.

THE COMPULSORY QoS VALUE

The idea behind the introduction of a "compulsory" QoS value is the following one: *when a compulsory value has been selected for a QoS parameter of a service facility, the service provider will monitor this parameter and abort the service facility when it notices that it cannot achieve the requested service.*

It must be clear that no obligation of results is linked to the idea of compulsory value. The service provider **tries** to respond to the requested service facility and, by monitoring its execution, it will

- either execute it completely without violating the selected compulsory value of the performance parameter;
- or abort it if the selected compulsory value of the performance parameter is not fulfilled.

The compulsory concept reflects the fact that, in some environments (e.g. a lightly loaded LAN), the compulsory QoS value may be achieved without resource reservation. Of course, the same LAN, which does not provide any reservation mechanism or any priority mechanism, may, when heavily loaded, prevent the service provider from reaching the compulsory QoS value and oblige it to abort the execution of the requested service facility.

THE THRESHOLD QoS VALUE

Some service users may find that the solution of aborting the requested service facility, when one of the compulsory QoS values is not reached, is a little too radical. They may prefer to get information about the degradation of the QoS value.

To achieve that we propose to introduce a “threshold” QoS value with the following semantics: *when a threshold value has been selected for a QoS parameter of a service facility, the service provider will monitor this parameter and indicate to the service user(s) when it notices that it cannot achieve the selected value.*

This threshold QoS value may be used without an associated compulsory value. In this case, the behaviour of the service provider is very similar to the one it has to adopt with a compulsory value. The main difference is that, instead of aborting the service facility, it warns, when it notices it is unable to provide the specified value, either or both users depending of the service definition. If the service provider is able to provide a QoS value better than the threshold value, everything is fine.

Threshold QoS versus Best Effort QoS

If the threshold QoS is used without any compulsory QoS, the main difference between the threshold and the best effort is that in the former case, the service provider has the obligation to monitor the parameter and to indicate if the threshold value is not reached.

Threshold and Compulsory QoS values

It is possible to associate, to the same QoS parameter, a threshold and a compulsory QoS values with, of course, the threshold value “stronger” than the compulsory one.

If the performance parameter degrades slowly and continuously, an indication will be delivered to the service users before the abortion of the service facility. Until such a threshold indication occurs, the service user knows that the service facility is not endangered by the current parameter value.

THE MAXIMAL QUALITY QoS VALUE

In most cases, if the service provider is able to offer a “stronger” value of the QoS parameter than the threshold, the service user will not complain about it. But it could happen that the service user wants to put a limit to a “richer” service facility.

A called entity, for instance, may want to put a limit to the data arrival rate or a calling entity may want, for reasons of cost, to prevent the use of too many resources by the service provider.

Such a parameter may be useful to smooth the behaviour of the service provider. Introducing a maximal quality QoS value on a transit delay, i.e. fixing a lower bound to the transit delay values will reduce the transit delay jitter and facilitate the resynchronization at the receiving side.

To achieve that we propose to introduce a “maximal quality” QoS value with the following semantics: *when a maximal quality value has been selected for a QoS parameter of a service facility, the service provider will monitor this parameter and avoid occurrence of interactions with the service users that would give rise to a violation of the selected value.*

It is possible to associate, to the same QoS parameter, a maximal quality, a threshold and a compulsory QoS values with, of course, the maximal quality “stronger” than the threshold value, itself “stronger” than the compulsory value.

QoS PARAMETERS AND INFORMATION PARAMETERS

The introduction of compulsory QoS values implies that the service provider will have a more difficult task to fulfil. It is therefore not surprising that the service user may have to provide the service provider with more information about the characteristics of the elements associated with the request in order to facilitate the decision of rejection or acceptance of the request. Hence, the introduction of the concept of compulsory QoS requires the introduction, in the primitives associated with a request, of additional parameters. These additional parameters may be designated as information parameters to distinguish them from the QoS parameters proper.

Information values of QoS parameters may also have to be introduced to control the negotiation process in the case of compulsory or threshold QoS to preserve the semantics associated with the negotiated value.

THE NEGOTIATED QoS PARAMETERS

Depending upon the type of facilities, these QoS parameters will be submitted to a negotiation or not. In case of a negotiation, some rules related to the rights of strengthening or weakening (lowering) the QoS values are defined for each type of value and for each participant in the negotiation. This has to be done to keep the final result coherent with the meaning of the parameter value.

Negotiation of a Compulsory QoS value

When a negotiation occurs on a compulsory value, the selected compulsory value of a performance parameter may be different from the value proposed first by the calling service user. However, in order to remain coherent with the semantics attached to a compulsory value, some rules on the negotiation procedure must be imposed.

First, it must be clear that if a service user introduces a compulsory QoS value for a performance parameter to be negotiated, the only possible modification is the strengthening of this compulsory value. In particular, it is absolutely excluded for the service provider to modify this value in order to relax the requirement

However the calling service user may not be interested in an unlimited strengthening of the proposed compulsory QoS value. It introduces therefore in the request primitive, a second value which fixes a bound indicating to what extent the proposed compulsory QoS value may be strengthened.

When the service provider analyses the request of the calling service user, it has to decide whether it rejects it or not (it can already do so as it knows that the request could only be strengthened).

This reject decision may be based:

- on the rejection by the underlying service provider of the underlying request with related compulsory values;
- on the overall knowledge of the capability of the underlying service;
- on the past history of other associations of the same type involving the same set of service users. An interesting way to be explored is the feasibility for the service provider to rely on the information collected by the management entities of the service.

In the latter case, it has to examine the bound of strengthening. This bound may be made weaker (brought closer to the compulsory value) by the service provider, before issuing the indication primitive to the called service user, in such a way to give, to the called service user, the range of compulsory values acceptable by both the calling service user and the service provider.

The service provider does not have to strengthen the compulsory QoS value which must be seen as the expression of the requirements of the service users.

After receiving the indication primitive, the called service user may accept or reject the request. If it accepts it, it may modify (strengthen) the compulsory QoS value up to the value

of the bound and return it in its response. In this case the negotiation is completed and the service provider may confirm the acceptance of the request and provide the final selected compulsory QoS value to the calling service user.

If the negotiation is successful, the bound is of no interest anymore (the bound is an example of information values mentioned earlier) and the selected compulsory QoS value reflects now the final and global request to the service provider from both service users.

Negotiation of a Threshold QoS value

The negotiation procedure of a threshold value is similar to the negotiation procedure of a compulsory value. Here also the only possible modification is the strengthening of the threshold value. Here also the calling service user introduces, in the request primitive, an information parameter which fixes a bound indicating to what extent the proposed threshold QoS value may be strengthened.

If a compulsory and a threshold values are associated to the same QoS parameter, there exist a set of order relationship between the compulsory, the threshold and their bounds values which must be verified in the request primitive and maintained during the negotiation.

Negotiation of a Maximal Quality QoS value

If a service user introduces a maximal quality QoS value for a performance parameter, the only possible modification is the weakening of this maximal quality QoS value. This value can be weakened during the negotiation by the service provider that indicates by this way the limit of the service it may provide and by the called service user.

If the maximal quality is the only value associated with a given QoS parameter, no bound will be introduced in the request primitive and the negotiation will result in the selection of the weakest of the maximal quality values of the service users and the service provider. This is an example of negotiation for information exchange.

If the maximal quality value and a compulsory or/and a threshold values are associated to the same QoS parameter, there exist an order relationship between the maximal quality value and the bound value on the threshold (or the bound value on the compulsory value if no threshold value is specified) which must be verified in the request primitive and preserved during the negotiation.

THE NON-NEGOTIATED QoS PARAMETERS

The Non-negotiated Compulsory QoS

When a service user asks for a service facility with a performance parameter at a given compulsory value, the service provider will reject the request, without even trying to provide the requested service, when the service provider knows that it will be unable to succeed.

This decision may be based:

- on the rejection by the underlying service provider of the underlying request with related compulsory values;
- on the overall knowledge of the capability of the underlying service;
- on the past history of other associations of the same type involving the same set of service users. An interesting way to be explored is the feasibility for the service provider to rely on the information collected by the management entities of the service.

The Non-negotiated Threshold and Maximal Quality QoSs

When a service user asks for a service facility with a performance parameter at a given threshold value, the service provider will accept the request. The same is true for a service facility with a performance parameter at a given maximal quality value.

THE DEFINITION OF THE PERFORMANCE QoS PARAMETERS

Remember that the QoS refers collectively to certain characteristics associated with the use of the available service facilities, as observed between the SAPs. The only aspects,

regarding the use of the service facilities, that are observable at the SAPs are the interactions which take place by means of service primitives.

The QoS is specified in terms of QoS parameters. The performance QoS parameters, in particular, are those QoS parameters whose definition is based on a performance criterion to be assessed.

It is very important to have a clear view about the way to define a performance QoS parameter. Of course the assessment of a performance criterion requires the introduction of timing considerations. Since the only events observable by a service user, when it is using a service facility, are the primitives that occur at its SAP, the only notion of time which can be relied on to introduce timing considerations is the notion of time of occurrence of a service primitive at a SAP. Such a time is an absolute time which is not usable as such.

Time has to appear in the definition of the performance QoS parameters in the form of time intervals between occurrences of service primitives. Thus, any performance QoS parameter has to be defined in relation with the occurrences of two or more related service primitives. These primitives may be related in very different ways. They may be related just by the fact that they pertain to a same connection, or by the fact that they are occurring successively at a same SAP, or by the fact that one or several parameters of one of the primitives occurring at a given SAP have been replicated in the other primitive(s) occurring at peer SAP(s), etc. The figure 2 give example of time intervals usable in relation with the performance parameter definition. More complex relationships resulting from combinations of several basic relationships are also possible.

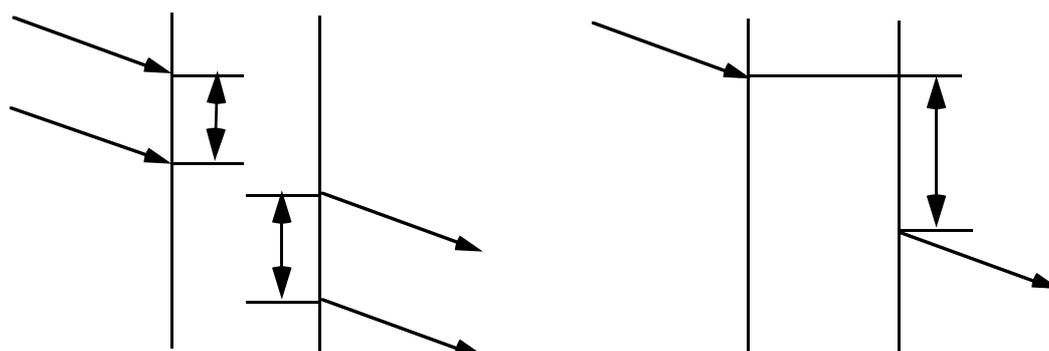


Figure 2. Time intervals usable in relation with the performance parameter definition

Transit Delay Definition

For example, the **transit delay** associated with an invocation of a peer-to-peer SDU transfer facility may be defined as *the time interval between the occurrence, at the SAP of the sending user, of the DATA request primitive that conveys the SDU and the occurrence, at the peer SAP, of the corresponding DATA indication primitive*. This definition may be used in the connectionless-mode case as well as in the connection-mode case. Here, the DATA request and indication primitives are related by the fact that the SDU parameter of the DATA indication has been replicated from the SDU parameter of the DATA request.

Transit Delay Jitter Definition

In connection mode, it is conceivable that a performance parameter calculated at each invocation of a service facility intervenes in the definition of a more complex performance parameter whose calculation is based on a certain number of invocations of this service facility. A **transit delay jitter** may be defined for one direction of data transfer of a connection as *the difference between the longest and the shortest transit delays observed on this direction since the connection establishment*. Here, each pair formed by a DATA request and the corresponding DATA indication primitives is related to the other pairs of primitives by the fact that they pertain to a same direction of data transfer on a connection.

Throughput Definition

An example of a performance parameter whose definition is based on the time of occurrence of service primitives related by the fact that they occur at the same SAP is the throughput of a direction of data transfer on a connection. Unlike the definitions of the previous two performance parameters, the definition of the throughput does not only use time intervals between occurrences of service primitives, but uses an additional quantity, namely the length of the transferred SDUs.

In the specifications of the ISO services [ISO 8072], the throughput for one direction of transfer of a connection is defined in terms of a sequence of n (with $n \geq 2$) successfully transferred SDUs.

When particularised to the case $n = 2$, this definition becomes:

“the throughput for one direction of transfer is defined to be the smaller of:

- a) the number of SDU octets contained in the last transferred SDU divided by the time interval between the previous and the last T-DATA requests; and*
- b) the number of SDU octets contained in the last transferred SDU divided by the time interval between the previous and the last T-DATA indications”*,

where the throughput measured in a) and b) may be referred to respectively as the sending user's throughput and the receiving user's throughput for the direction of transfer considered. When $n = 2$, the definition corresponds to an instantaneous throughput.

The figure 3 represent the ISO definition of the sending user's throughput when $n = 2$.

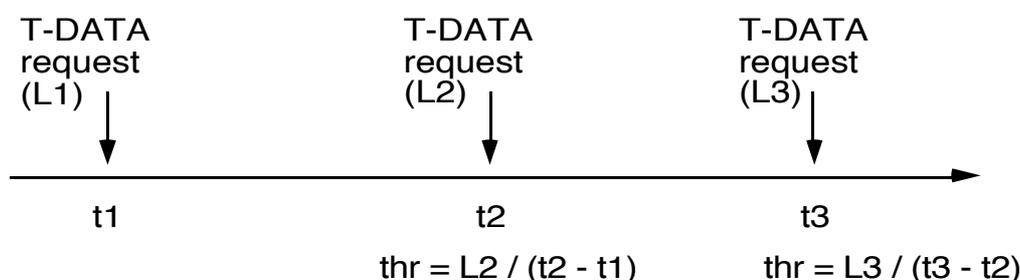


Figure 3. ISO definition of the sending user's throughput

An alternative way to define the throughput associated with an invocation of the SDU transfer facility on one direction of transfer of a connection is the following: the **global throughput** for one direction of transfer is still defined to be the smaller of the **sending user's throughput** and the **receiving user's throughput** but these two throughputs are now defined differently from the ISO's view:

- the **sending user's throughput** is now defined as *the number of SDU octets contained in the last transferred SDU divided by the time interval between the last and the next T-DATA requests.*
- the **receiving user's throughput** is now defined as *the number of SDU octets contained in the last transferred SDU divided by the time interval between the corresponding last and next T-DATA indications.*

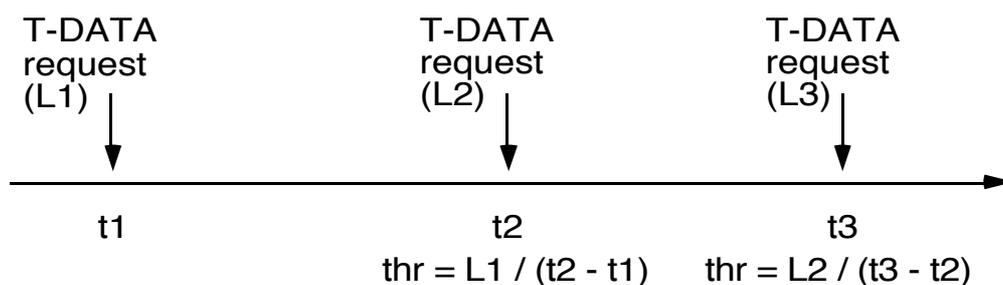


Figure 4. OSI 95 definition of the sending user's throughput

This definition corresponds also to an instantaneous throughput. It is this definition of

the throughput that has been adopted in the framework of the OSI 95 Enhanced Transport Service Definition [DBL 92b, DBL 92c, BLL 92b]. The figure 4 represent the OSI 95 definition of the sending user's throughput

In comparison with the ISO's one, the latter definition of the global throughput has a great advantage. Indeed, this definition is much more convenient for the service provider that has to check its ability to correctly fulfil its contract, in particular as regards the selected compulsory and threshold throughput values, during the data transmission phase.

With the latter definition, the behaviour of the service provider between two occurrences of DATA request primitives on the sending side is influenced only by the time that is elapsing. This is due to the fact that the length of the SDU that is used for the calculation of the current throughput value (i.e. the value associated with the last invocation of the SDU transfer facility) is known. Only the time interval between the last DATA request and the next expected DATA request is unknown until the occurrence of this next primitive. This means that the constraints which the service provider has to obey on the sending side will be expressed only in terms of the time already elapsed since the last DATA request, regardless of the possible lengths of SDU in the next DATA request. Of course a similar conclusion is true on the receiving side.

With the ISO's definition, the monitoring of a compulsory or a threshold throughput value is more intricate. In fact, the behaviour of the service provider between two occurrences of DATA request primitives on the sending side depends not only upon the time that is elapsing but also upon the possible lengths of SDU in the next expected DATA request. This results from the fact that the length of the SDU that is to be used for the calculation of the current throughput value (i.e. the value associated with the next invocation of the SDU transfer facility) is obviously unknown, as well as the time interval between the last DATA request and the next expected DATA request, until the occurrence of this next primitive. This means that the constraints which the service provider has to obey on the sending side will be expressed in terms of the time already elapsed since the last DATA request but also in terms of the possible lengths of SDU in the next DATA request. Of course a similar conclusion is true on the receiving side.

EXAMPLES

The purpose of the following examples is to illustrate the use of the various values of a performance QoS parameter.

Data Transfer in a Connectionless-mode Service

If we introduce, in an UNITDATA request, a compulsory value for the transit delay, the time between an UNITDATA request and the corresponding UNITDATA indication will never exceed this value: either it will be shorter, or no UNITDATA indication will ever occur. In the latter case, no indication of failure will be given to the sending user, which is coherent with the semantics of a connectionless transmission.

Introducing a threshold value for the transit delay in an UNITDATA request may appear meaningless as the semantics of this transmission service implies that the sending user does not receive any further information about its request. However, it is possible to envisage the signalling of the threshold violation in the UNITDATA indication on the receiving side.

Introducing a maximal quality value for the transit delay in an UNITDATA request seems difficult to justify as there is very little rationale for a sending user in a connectionless-mode to prevent the transit delay to be lower than a given value.

Data Transfer in an Acknowledged Connectionless-mode Service

In an acknowledged connectionless-mode, such as the one introduced in [Dan 92a, DBL 92a], it is possible to introduce a threshold value for the transit delay. A violation of the threshold value (actual transit delay value above the threshold value) could be signalled to the sending user by a special parameter of the confirm primitive as well as to the receiving user by a special parameter of the indication primitive.

In an acknowledged connectionless-mode it is also possible to introduce a compulsory value for the transit delay. A violation of the compulsory value will prevent the ACKDATA

indication to be issued and the sending user will receive the notification of non-delivery via the REJECT- ACKDATA.indication primitive.

Introducing a maximal quality value for the transit delay appears difficult to justify as there is very little rationale for a sending user in an acknowledged connectionless-mode to prevent the transit delay to be lower than a given value.

Data Transfer in a Connection-mode Service

If a compulsory value for the transit delay has been negotiated during the connection establishment, the time between a DATA request and the corresponding DATA indication will never exceed this value: either it will be shorter, or no DATA indication will ever occur. In the latter case, the service provider will abort the connection by issuing a DISCONNECT indication with the reason of the abortion.

It is possible to negotiate, during the connection establishment, a threshold value for the transit delay. If the time between a DATA request and the corresponding DATA indication is above the threshold value, the service provider will signal it to the service user(s) by a REPORT indication primitive.

Introducing a maximal quality value is not anymore meaningless in a connection-mode transmission. Preventing the transit delay to be lower than a given value implies that the DATA indication may be delayed at the receiving user SAP and this may be an efficient means to reduce the transit delay jitter. With a maximal quality value corresponding to a minimum value of the transit delay and with a compulsory value corresponding to a maximum value of the transit delay, it is possible to introduce an upper bound on the transit delay jitter.

A More Complex Example of Negotiated QoS Parameter

As a more complex example of a negotiated QoS parameter, let us consider the throughput on a direction of transfer of a connection with the associated compulsory, threshold and maximal quality values.

The compulsory is the smallest value and represents the throughput that the service provider must be able to offer to the sending and to the receiving users. If at anytime during the whole life of the connection, the service provider happens to be unable to maintain this throughput a DISCONNECT indication will occur.

If the traffic on the connection is dedicated to a video channel, the compulsory value may represent the throughput which allows to maintain the synchronisation between the sending and the receiving users.

However, in the most general case, two particular situations have to be considered. Both are dealing with the responsibilities of the users on the one hand and of the service provider on the other hand when the selected values are not matched.

At the sending SAP, the compulsory value may be violated because no DATA request occurred in due time owing to the behaviour of the sending user only (i.e. the service provider was ready to accept a DATA request in due time). This does not imply a failure of the service provider but only results from the behaviour of the sending user. In this case, it may be reasonable to believe that the service provider does not have to issue a DISCONNECT indication.

If at the receiving SAP, no DATA indication occurs due to the receiving user, the compulsory QoS value will also be violated. Here however it seems reasonable to recognise that the compulsory value does not only constrain the service provider but also constrains the receiving user. Its acceptance of the compulsory value implies that it believes to be able to support the compulsory value. Therefore a DISCONNECT indication seems appropriate. This asymmetrical behaviour may not be acceptable in every application environment.

It is possible, for the threshold value to apply the same rules as for the compulsory value, issuing, of course, a REPORT indication instead of a DISCONNECT indication. However, taking into account the warning characteristic associated with the threshold value, the service provider may also issue a REPORT indication to the service users as soon as the threshold value is violated.

The maximal quality QoS value is the highest value of the three. The service provider, by controlling the occurrence of the interactions with the service users, prevents the sending user's throughput and the receiving user's throughput from crossing the limiting value.

An interesting point to mention is that the combination of the compulsory and maximal quality values helps to make more precise the model of the service provider. The internal

queues may build up at a maximum pace linked to the difference between the maximal quality and the compulsory values of the throughput.

The Connection-mode Transport Service of OSI 95

For the connection-mode transport service of OSI 95, we used the structure of three values, compulsory, threshold and maximal quality.

The introduction of compulsory QoS values in the QoS negotiation implies that the TS provider will have a more difficult task to fulfil. It is therefore not surprising that the TS user may have to provide the TS provider with more information about the characteristics of the sequence of TSDUs it intends to submit. Requesting a throughput of 2 Mb/s with TSDUs of 10 Kbytes is different from requesting a throughput of 2 Mb/s with TSDUs of 40 bytes. Hence, the introduction of the concept of compulsory QoS requires the introduction, in the primitives associated with the opening of a TC, of additional parameters such as the maximum and the minimum size of TSDU. As indicated earlier, these additional parameters will be designated here as information parameters to distinguish them from the QoS parameters proper.

Assume that through a successful negotiation we have the three values for the throughput defined as indicated at the figure 4. It may be of interest to see how the service provider will operate.

The minimum compulsory value for the throughput is the value that the TS provider must be able to maintain during the whole TC lifetime. Otherwise it must immediately shut down the TC. So, if a T-DATA.request interaction occurs at time T_0 , with a TSDU of size L , the TS provider must be ready to offer another T-DATA.request interaction at the latest at time $T_0 + \Delta t_{\max}$, where Δt_{\max} is given by the ratio of L to the minimum compulsory value for the throughput. Besides, the TS provider must be able to offer the first T-DATA.request interaction immediately after the receipt of the T-CONNECT.response on the called side and the issuance of the T-CONNECT.confirm on the calling side.

The maximum value for the throughput is a value that may never be exceeded for the duration of the TC. So, if a T-DATA.request interaction occurs at time T_0 , with a TSDU of size L , the TS provider shall offer another T-DATA.request interaction at the earliest at time $T_0 + \Delta t_{\min}$, where Δt_{\min} is given by the ratio of L to the maximum value for the throughput. The maximum value for the throughput can thus be used for rate control purposes.

Figure 5 shows the relation between the TSDU size L and the Δt_{\min} and Δt_{\max} .

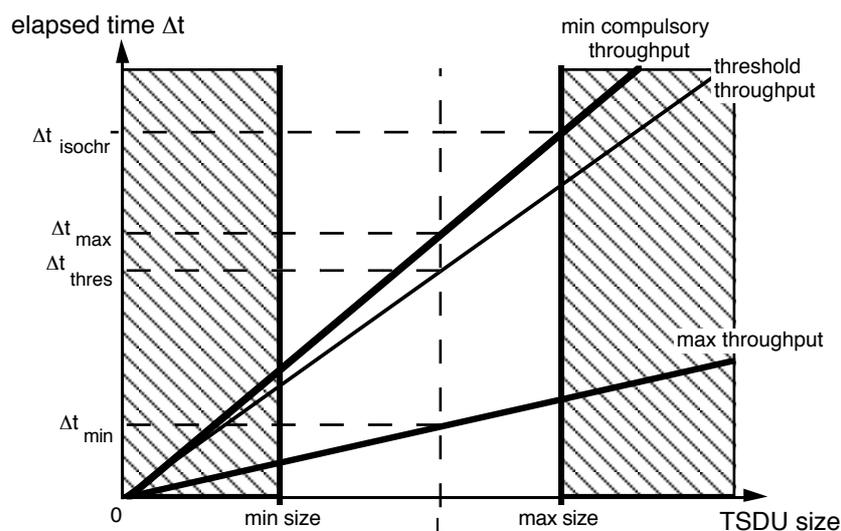


Figure 5. Relation between the TSDU size and the Δt_{\min} and Δt_{\max}

If the TS provider is not able to offer another T-DATA.request interaction at the latest at time $T_0 + \Delta t_{\max}$, the minimum compulsory throughput cannot be achieved and the TS provider must shut down the TC. This does not mean that the throughput associated with a particular invocation of the data transfer facility may never go under the minimum compulsory

value. This may happen if the offer of the TS provider for another T-DATA.request interaction is not matched in time by the sending TS user. In such a case however, the inability of maintaining the throughput at or above the minimum compulsory value is entirely the responsibility of the sending TS user. Our choice of an “instantaneous” throughput measured at each invocation of the data transfer facility has the advantage of putting a clear separation between the responsibilities of the TS user and those of the TS provider. The TS provider does not have to disconnect in case of a late submission of TSDU from the part of the TS user.

If a sending TS user generates an isochronous traffic, that is if it produces TSDUs (possibly of variable length) at a fixed rate, this sending TS user may want the TS provider to be ready to accept the submission of TSDUs at the same rate characterised by the constant time interval Δt_{isochr} between the submission of successive TSDUs whatever their length. For this purpose, the TS user may negotiate with the TS provider a minimum compulsory throughput given by the ratio of the maximum size authorised for the TSDUs to the constant time interval Δt_{isochr} . With such a minimum compulsory throughput, if T_0 is the time at which the last T-DATA.request occurred, the sending TS user is assured that the TS provider will be able to offer another T-DATA.request interaction at the latest at time $T_0 + \Delta t_{\text{isochr}}$ in case the TSDU submitted in the interaction at time T_0 is of maximum size or even earlier in case the TSDU is of smaller size, and will disconnect otherwise.

However, by forcing the TS provider to be able to accept a new T-DATA.request before $T_0 + \Delta t_{\text{isochr}}$ when the last TSDU is not of maximum size, we impose an unnecessary constraint on it. In fact, for a really isochronous traffic, forcing the TS provider to be able to offer a T-DATA.request interaction every Δt_{isochr} units of time, whatever the size of the TSDU submitted in the previous interaction, is a sufficient constraint. That is why we have decided to add a traffic type indicator to the QoS components pertaining to the throughput parameter. If this traffic type indicator is “non-isochronous”, then the Δt_{max} is calculated as explained above, i.e. by taking the size L of the previous submitted TSDU into account. By contrast, if the traffic type indicator is “isochronous”, then the Δt_{max} is simply taken equal to Δt_{isochr} as if all the TSDUs were of maximum size (Figure 5).

Obviously, if the negotiated minimum compulsory throughput is equal to zero, Δt_{max} is becoming infinite and we come back to the concept of best effort QoS.

To respond to certain requirements, we have also introduced, in our list of QoS components pertaining to the throughput parameter, the idea of a threshold value for the throughput that would be slightly above the minimum compulsory value (Figure 5). If the TS provider was not able to maintain the threshold throughput, i.e. if it was not ready to offer another T-DATA.request interaction at time $T_0 + \Delta t_{\text{thres}}$, then it would have to indicate its inability to the TS users but the TC would remain open.

In summary, the behaviour of the TS provider will be time-dependent according to the following table where T_0 is the time at which the last T-DATA.request occurred:

In $]T_0, T_0 + \Delta t_{\text{min}} [$,	the provider will not accept a T-DATA.request.
In $[T_0 + \Delta t_{\text{min}}, T_0 + \Delta t_{\text{thres}} [$,	the provider may accept or not a T-DATA.request.
In $[T_0 + \Delta t_{\text{thres}}, T_0 + \Delta t_{\text{max}} [$,	the provider must accept a T-DATA.request, or otherwise must indicate it to the TS user.
In $[T_0 + \Delta t_{\text{max}}, \infty [$,	the provider must accept a T-DATA.request or otherwise must start the release of the TC.

In this table, $\Delta t_{\text{max}} = L / \text{minimum throughput}$,

where L is either the size of the last TSDU (in the non-isochronous case) or the maximum TSDU size (in the isochronous case).

CONCLUSION

We present in this paper an enhancement of the QoS model able to match the communication requirements of the new application environment.

In our proposed semantics, a QoS parameter is seen as a structure of three values, respectively called “compulsory”, “threshold” and “maximal quality”, all these values being

optional. Each one has its own well defined meaning, and expresses a contract between the service users and the service provider.

This means, and this is the most fundamental difference with the present situation, that the service provider is now submitted to some well defined duties, known by each side. In other words, the rules of the game are clear. These duties correspond to services that we found useful for a service user in the above-mentioned current environment.

The existence of a contract between the service users and the service provider implies that, in some cases, the service users are also submitted to well defined duties also derived from the application environment.

We have applied the ideas of this paper in the design of the OSI 95 Enhanced Transport Service [Dan 92b, BLL 92a, DBL 92b, BLL 92b] and we are currently working in the framework of the CIO² project on the integration of the new QoS semantics into the protocol developed to provide this enhanced service

The goal of the OSI 95 project was not only to specify new transport service and protocol. It is also to use LOTOS as formal description technique, from the very beginning of the design process. The methodology for the design of the service and the protocol has been based on the constraint-oriented style in order to allow an incremental development of the protocol. This method allows the addition of, the removal of or a modification in a mechanism without jeopardising the other part of the work. This methodology has been tested successfully on TP4 [Led 92] and has been applied for the service and the protocol [BLL 92b].

The work done in OSI 95 and in CIO has been introduced in the ISO in the framework of the ECFF (Enhanced Communication Functions and Facilities) new working group of ISO/SC6 [BLL 92a], [DBL 92a], [DBL 92b], [BLL 93], [Bag 93], [DBL 93]. The document [DBL 93] is presently circulated in the National Bodies and Liaison Organisation of SC6 for review and comments.

REFERENCES

- [Alm 92] P. ALMQUIST, **Type of Service in the Internet Protocol Suite**, Network Working Group RFC 1349, July 1992
- [Bag 93] Y. BAGUETTE, **Enhanced Transport Service Definition (Informal specification in English - Version 1)**, ISO/IEC JTC1/SC6/WG4/N822, 19 Apr. 1993, 85 p. (R2060/ULg/CIO/DS/P/003/b1, Feb. 1993, SART 93/03/15)
- [BLL 92a] Y. BAGUETTE, L. LEONARD, G. LEDUC, A. DANTHINE, **Belgian National Body Contribution - Four Types of Enhanced Transport Services and their LOTOS Specifications**, SC6 plenary meeting, San Diego, July 8-22, 1992, ISO/IEC JTC1/SC6 N7323, May 1992, 58 p. (OSI 95/ULg/A/22/TR/P, May 1992, SART 92/12/09).
- [BLL 92b] Y. BAGUETTE, L. LEONARD, G. LEDUC, A. DANTHINE, O. BONAVENTURE, **OSI 95 Enhanced Transport Facilities and Functions**, OSI 95/Deliverable ULg-A/P, Dec.1992, 277 p. (SART 92/25/05)
- [BLL 93] Y. BAGUETTE, L. LEONARD, G. LEDUC, A. DANTHINE, **OSI 95 Enhanced Transport Services**, ISO/IEC JTC1/SC6/WG4/N821, 19 Apr. 1993, 107 p. (R2060/ULg/CIO/DS/S/002/b1, Jan. 1993, SART 93/02/15)
- [CCC 91] COMER M.H, CONDRY M.W., CATTANACH S., CAMPBELL R., **Getting the Most for Your Megabit**, *Computer Communication Review*, 1991, Vol. 21, nr 3, pp. 5-12
- [Che 88] D. CHERITON, **VMTP : Versatile Message Transaction Protocol - Protocol Specification**, Network Working Group RFC 1045, Feb. 1988
- [Ches 89] CHESSON G., **XTP/PE Design Considerations**, *Protocols for High-Speed Networks*, Zurich, May 9-11, 1989, Rudin H. & Williamson R. Ed., Elsevier (North Holland), pp. 27-32
- [ChJ 89] CHIU D-M, JAIN R., **Analysis of the Increase and Decrease Algorithms for Congestion Avoidance in Computer Networks**, *Computer Networks and ISDN Systems*, 1989, Vol. 17, nr 1, pp. 1-14
- [ChW 89] CHERITON D.R., WILLIAMSON C.L., **VMTP as the Transport Layer for High-Performance Distributed Systems**, *IEEE Communications Magazine*, 1989, Vol. 27, n° 6, pp. 37-44
- [CJR 89] CLARK D.D., JACOBSON V., ROMKEY J., SALWEN H., **An Analysis of TCP**

² CIO is the acronym of RACE Project 2060 "Coordination, Implementation and Operation of Multimedia Services".

- Processing Overhead**, *IEEE Communications Magazine*, 1989, Vol. 27, nr 6, pp. 23-29
- [CLZ 87] CLARK D.D., LAMBERT M.L., ZHANG L., **NETBELT : A Bulk Data Transfer Protocol** Network Working Group RFC 998, March 1987
- [Dan 89] DANTHINE A., **Communication Support for Distributed Systems - OSI versus Special Protocols**, *Proceedings of the 11th World Computer Congress on Information Processing*, San Francisco, 28 August-1 September 1989, pp. 181-190
- [Dan 90] DANTHINE A., **10 Years with OSI**, *INDC-90*, Lillehammer, Norway, 26-29 March 1990, (Ed. : D. KHAKHAR, F. ELIASSEN), Elsevier (North Holland), 1990, pp. 1-16
- [Dan 92a] A. DANTHINE, **A New Transport Protocol for the Broadband Environment**, *IFIP Workshop on Broadband Communications*, Estoril, January 20-22, 1992, A. Casaca, ed., Elsevier (North-Holland), 1992, pp. 337-360.
- [Dan 92b] A. DANTHINE, **Esprit Project OSI 95 - New Transport Services for High-Speed Networking**, *3rd Joint European Networking Conference*, Innsbruck, May 11-14 1992, also in: *Computer Networks and ISDN Systems 25 (4-5)*, Elsevier Science Publishers (North-Holland), Amsterdam, November 1992, pp. 384-399.
- [DBL 92a] A. DANTHINE, Y. BAGUETTE, G. LEDUC, **Belgian National Body Contribution - Issues Surrounding the Specification of High-Speed Transport Service and Protocol**, *SC6 interim meeting, Paris, February 10-13, 1992*, ISO/IEC JTC1/SC6 N7312, January 1992 58 p. (OSI 95/ULg/A/15/TR/P/V2, 46 p., January 1992, SART SART 92/04/05).
- [DBL 92b] A. DANTHINE, Y. BAGUETTE, G. LEDUC, L. LEONARD, **Belgian National Body Contribution - The Enhanced Connection-Mode Transport Service of OSI 95**, *SC6 plenary meeting, San Diego, July 8-22, 1992*, ISO/IEC JTC1/SC6 N7759, June 1992 16 p. (OSI 95, OSI 95/ULg/A/24/TR/P, 16 p., June 1992, SART 92/14/05).
- [DBL 92c] A. DANTHINE, Y. BAGUETTE, G. LEDUC, L. LEONARD, **The OSI 95 Connection-mode Transport Service - The Enhanced QoS**, *IFIP Conference on High Performance Networking*, Liège, December 16-18, 1992, to appear in: A. Danthine, O. Spaniol, eds., *High Performance Networking*, Elsevier Science Publishers (North-Holland), Amsterdam, 1993.
- [DBL 93] A. DANTHINE, Y. BAGUETTE, L. LEONARD, G. LEDUC, **An Enhancement of the QoS Concept**, ISO/IEC JTC1/SC6/N8010, Jan. 1993, 16 p. (OSI 95/ULg/A/28/TR/P, Jan. 1993)
- [DHH 88a] DANTHINE A., HENQUET P., HAUZEUR B., CONSTANTINIDIS C., FAGNOULE D., CORNETTE V., **Access Rate Measurement in the BWN Environment**, *Proceedings of High Speed Local Area Networks II*, Liège, 14-15 April 1988, Elsevier Science Publishers BV(North Holland), 1990, pp. 89-115
- [DHH 88b] DANTHINE A., HAUZEUR B., HENQUET P., CONSTANTINIDIS C., FAGNOULE D., CORNETTE V., **Corporate Communication System by LAN Interconnection**, *EURINFO 88 - Information Technology for Organisational Systems*, Athènes, 16-20 May 1988 Bullinger H.J. & al., Ed., Elsevier (North Holland), 1988, pp. 315-326
- [DRB 87] DE PRYCKER M., RYCKEBUSCH M., BARRY P., **Terminal Synchronization in Asynchronous Networks**, *IEEE Conference on Communications*, Seattle, 1987
- [GKW 89] GIARRIZZO D., KAISERWERTH M., WICKI T., WILLIAMSON R.C., **High-Speed Parallel Protocol Implementation**, *Protocols for High-Speed Networks*, Zurich, May 9-11, 1989, Rudin H. & Williamson R. Ed., Elsevier (North Holland), pp. 165-180
- [Her 88] HERBERT A., **Communications Aspects in ANSA**, *Computer Standards & Interfaces*, 8 (1988) pp 49-56
- [Hes 89] HEATLEY S., STOKESBERRY D., **Analysis of Transport Measurements Over a Local Area Network**, *IEEE Communications Magazine*, 1989, Vol. 27, nr 6, pp. 16-22, also in *Protocols for High-Speed Networks*, Zurich, May 9-11, 1989
- [ISO 8348] ISO/IEC JTC1, **Information Technology - Telecommunications and Information Exchange between Systems - Network Service Definition for Open Systems Interconnection**, ISO/IEC JTC1/SC6 N7558, 21 Sep. 1992
- [ISO DIS 8072] ISO/IEC JTC1, **Information technology - Transport service definition for Open Systems Interconnection**, ISO/IEC DIS 8072 (DIS ballot terminates on April 1993).
- [Jac 88] JACOBSON V., **Congestion Avoidance and Control**, *SIGCOMM 88*, Stanford, Aug. 16-19, 1988, pp. 314-329
- [Jai 86] JAIN R., **A Timeout-Based Congestion Control Scheme for Window Flow-Controlled Networks**, *IEEE Journal on Selected Areas in Communications*, Vol. SAC-4, nr 7, Oct. 1986, pp. 1162-1167
- [Jai 90] JAIN R., **Congestion Control in Computer Networks : Issues and Trends**, *IEEE Network Magazine*, 1990, Vol. 4, nr 3, pp. 24-30
- [JSB 90] JAIN N., SCHWARTZ M., BASHKOW T.R., **Transport Protocol Processing at GBPS**

- Rates**, *SIGCOMM 90*, Philadelphia, Sep. 24-27, 1990, *Computer Communications Review*, 1990, Vol. 20, nr 4, pp. 188-199
- [Led 92] LEDUC G., **A Methodology for the Design of Large LOTOS Specifications and its Application to ISO 8073**, Report from the ESPRIT II Project OSI 95, OSI 95/ULG/A/16/TR/R/V2, April 1992.
- [Mab 91] MAFLA E., BHARGAVA B., **Communication Facilities for Distributed Transaction-Processing Systems**, *Computer*, 1991, Vol. 24, nr 8, p. 61
- [Mei 91] MEISTER B.W., **A Performance Study of the ISO Transport Protocol**, *IEEE Transactions on Computers*, 1991, Vol. 40, nr 3, pp. 253-262
- [NRS 90] NETRAVALI A.N., ROOME W.D., SABNANI K., **Design and Implementation of a High-Speed Transport Protocol**, *IEEE Transactions on Communications*, 1990, Vol. 38, nr 11, pp. 2010-2024
- [Par 92] C. PARTRIDGE, **A proposed Flow Specification**, Network Working Group RFC 1363, Sep. 1992
- [PE 91] PROTOCOL ENGINES, **XTP Protocol Definition — Revision 3.6 — Editor's Second Draft**, November 4, 1991.
- [Pos 81] J. POSTEL, Ed., **Internet Protocol - Darpa Internet Protocol Specification**, Network Working Group RFC 791, Sep. 1981
- [RaJ 88] RAMAKRISHNAN K.K., JAIN R., **A Binary Feedback Scheme for Congestion Avoidance in Computer Networks with a Connectionless Network Layer**, *SIGCOMM 88*, Stanford, Aug. 16-19, 1988, pp. 303-313
- [RaJ 90] RAMAKRISHNAN K.K., JAIN R., **A Binary Feedback Scheme for Congestion Avoidance in Computer Networks**, *ACM Transactions on Computer Systems*, 1990, Vol. 8, nr 2, pp. 158-181
- [Svo 89a] SVOBODOVA L., **Measured Performance of Transport Service in LANs**, *Computer Networks and ISDN Systems*, 1989, Vol. 18, nr 1, pp. 31-45
- [Svo 89b] SVOBODOVA L., **Implementing OSI Systems**, *IEEE Journal on Selected Areas in Communications*, 1989, Vol. 7, nr 7, pp. 1115-1130
- [Top 90] C. TOPOLCIC, Ed., **Experimental Internet Stream Protocol, Version 2 (ST-II)**, Network Working Group RFC 1190, Oct. 1990
- [Wha 89] WHALEY A.D., **The XPRESS Transfer Protocol**, *14th Conference on Local Computer Networks*, Minneapolis, Oct. 10-12, 1989, pp. 408-414
- [WiC 91] WILLIAMSON C.L., CHERITON D.R., **Loss-Load Curves : Support for Rate-Based Congestion Control in High-Speed Datagram Networks**, *SIGCOMM 91*, Sept. 1991, pp. 17-23
- [Zha 87a] ZHANG L., **Some Thoughts on the Packet Network Architecture**, *ACM Computer Communication Review*, 1987, Vol. 17, nr 1-2, pp. 3-17
- [Zha 87b] ZHANG L., **Designing a New Architecture for Packet Switching Communication Networks**, *IEEE Communications Magazine*, Sept. 1987, Vol. 25, nr 9, pp. 5-12
- [Zitt 91] ZITTERBART M., **High-Speed Transport Components**, *IEEE Network Magazine*, 1991, Vol. 5, nr 1, pp. 5