Université de Liège

Faculté des Sciences Appliquées

Systèmes et Automatique

Institut d'Electricité Montefiore, B28
Université de Liège au Sart Tilman
B-4000 Liège 1 (Belgique)

Esprit Project 5341

OSI95

The OSI95 Transport Service with Multimedia support

***

The OSI95 Connection-Mode Transport Service

Y. Baguette, L. Léonard, G. Leduc and A. Danthine

[BLL 94] Y. Baguette, L. Léonard, G. Leduc, A. Danthine, **The OSI95 Connection-Mode Transport Service**, *The OSI95 Transport Service with Nultimedia Support*, A. Danthine, University of liège, Belgium (ed.), Research Reports ESPRIT - Project 5341 - OSI95 - Volume 1, Springer-Verlag, 1994, pp. 181-198.

Date: July 1994
R94-05

The OSI95 Connection-Mode Transport Service

Yves Baguette[1], Luc Léonard[1], Guy Leduc[2] and André Danthine
[1] Research Assistant of the Belgian National Fund for Scientific Research (F.N.R.S.)
[2] Research Associate of the F.N.R.S.
Université de Liège, Institut d'Electricité Montefiore B28, B-4000 Liège 1, Belgium
Email: baguette@montefiore.ulg.ac.be

We present the main features of the connection-mode service developed as a part of the entire OSI95 Enhanced Transport Service. This connection-mode service results from modifications and enhancements to the standard ISO/IEC Connection-Mode Transport Service. Some of the enhancements are just mentioned in this paper for they are addressed in detail in companion papers.

Keywords: OSI95, Transport Service, Connection-Mode, QoS Semantics, Error Management, Transport Connection Release, QoS Re-negotiation, Out-of-Band Data Transfer.

1 Introduction

This paper deals with the main innovations of the OSI95 Enhanced Connection-Mode Transport Service in comparison with the ISO/IEC Connection-Mode Transport Service [ISO 8072].

These main innovations are:

- the ability to open a uni-directional Transport Connection (TC);
- a modification of the set of the Quality of Service (QoS) parameters used, with the introduction of an additional transit delay jitter performance QoS parameter but also the suppression of other performance QoS parameters (such as the residual error rate, the transfer failure probability, …) that are deemed unmanageable;
- the definition of a new semantics for the performance QoS parameters, consisting in the introduction of the concepts of compulsory QoS and threshold QoS;
- more flexibility of the error control (i.e. the error detection and recovery) on a TC;
- the widening of the TC release facilities to allow the graceful or abrupt release of either direction of data transfer separately from the reverse one on a TC;
- the introduction of a QoS re-negotiation facility;
- the introduction of an out-of-band data transfer facility associated with a TC.

The new semantics for the performance QoS parameters will not be explained in detail here since a specific companion paper [DBL 93] is devoted to this very important issue. Similarly, another specific companion paper [BaL 93] deals with the substantial widening of the TC release facilities consisting in the introduction of a graceful TC release facility and in the enhancement of the standard abrupt TC release facility. All the other aforementionned innovations are examined thoroughly in the remainder of this paper.

2 TC Establishment Facility and Uni-directional TCs

The OSI95 Enhanced Connection-Mode Transport Service is still offered through the classical sequence of four T-CONNECT primitives (request, indication, response and confirm) to establish a TC, allowing a full negotiation of the characteristics of the TC before the beginning of the data transfer.

A significant innovation in comparison with the ISO/IEC Connection-Mode Transport Service, however, is the ability to open uni-directional TCs. This is achieved by means of a new parameter in the T-CONNECT primitives: the "directions to open" parameter. The parameter indicates the direction(s) of data transfer that will be opened on the TC if the TC set-up request is accepted, by selecting one of three possible options:

a) “bidirectional TC”: both directions of data transfer have to be opened;
b) “direct unidirectional TC”: only the direction of data transfer from the calling OSI95 TS user to the called OSI95 TS user (referred to as the calling-to-called direction) has to be opened;
c) “reverse unidirectional TC”: only the direction of data transfer from the called OSI95 TS user to the calling OSI95 TS user (referred to as the called-to-calling direction) has to be opened.

The "directions to open" parameter is not negotiated at all. The decision on which direction(s) is (are) to be opened on a TC is taken by the calling OSI95 TS user only.

Additionally, a new kind of parameter in the T-CONNECT primitives is the TSDU size range parameter. There is a separate TSDU size range parameter for each direction of data transfer that is to be opened. The TSDU size range parameter associated with one direction of data transfer permits the negotiation, between the OSI95 TS users and the OSI95 TS provider, of a TSDU size range (i.e. a minimum and a maximum TSDU size) for this direction. The negotiated TSDU size range will apply to all normal TSDUs submitted by the corresponding sending OSI95 TS user via T-DATA requests. The TSDU size range for either direction, which is negotiated in parallel with the QoS for this direction, is a very useful information for the OSI95 TS provider and the called OSI95 TS user to determine more precisely their ability to provide the requested QoS.

3 Modification of the Set of QoS Parameters

The QoS parameters used in the ISO/IEC Connection-Mode Transport Service include parameters which express TS performance (cf. table 3.1) and two parameters which express other TS characteristics, namely the TC protection and the TC priority.

Phase	Performance criterion: Speed	Accuracy / Reliability
TC establishment	TC establishment delay	TC establishment failure probability (misconnection or TC refusal on the part of the TS provider)
Data transfer	Throughput Transit delay	Residual error rate (corruption, duplication or loss) Resilience of the TC Transfer failure probability
TC release	TC release delay	TC release failure probability

Table 3.1 The performance QoS parameters of the standard ISO/IEC Transport Service

For the OSI95 TS, we have kept the TC protection and TC priority parameters, without any change in their semantics up to now.

Among the performance QoS parameters, we have kept the throughput and transit delay parameters only, but with the new semantics of compulsory and threshold QoS. All other performance QoS parameters of the current ISO/IEC standard have been removed because they have been deemed unmanageable in practice.

We have introduced a new interesting performance QoS parameter, viz. the transit delay jitter, with here also the new semantics of compulsory and threshold QoS. Additionally, several QoS parameters have been introduced to deal with error control. The error control QoS parameters are discussed separately in the next chapter.

For each open direction of data transfer of a TC, the throughput is defined as *the ratio of the size (in octets) of a TSDU submitted via a T-DATA request primitive to the time elapsed until the occurrence of the next T-DATA request primitive relating to the same direction on the TC*. The definition of the throughput has been modified in comparison with the one in the current ISO/IEC standard : it relates now to a form of "instantaneous" throughput which is measured at each invocation of the TSDU transfer facility.

After a QoS negotiation (see [DBL 93] for more details on the negotiation scheme), the meaningful values regarding the throughput for an open direction of data transfer are the minimum compulsory value, the threshold value and the maximum value. The following relation must always be verified: 0 ≤ negotiated minimum compulsory throughput value ≤ negotiated threshold throughput value ≤ negotiated maximum throughput value. According to the semantics of compulsory and threshold QoS, the behaviour of the OSI95 TS provider will be time-dependent, as shown by table 3.2 where T_0 is the time at which the last T-DATA request primitive occurred, Δt_{max} = L / minimum compulsory throughput, Δt_{thres} = L / threshold throughput, Δt_{min} = L / maximum throughput and L is the size of the last submitted TSDU in the non-isochronous case (Figure 3.1):

Time interval considered:	Behaviour of the OSI95 TS provider:
In $]\ T_0, T_0 + \Delta t_{min}\ [$,	the OSI95 TS provider is not allowed to accept a T-DATA request interaction.
In $[\ T_0 + \Delta t_{min}, T_0 + \Delta t_{thres}\ [$,	the OSI95 TS provider is allowed to accept or not a T-DATA request interaction.
In $[\ T_0 + \Delta t_{thres}, T_0 + \Delta t_{max}\ [$,	the OSI95 TS provider should be able to accept a T-DATA request interaction at any time, otherwise it has to report its inability to both OSI95 TS users as soon as possible.
In $[\ T_0 + \Delta t_{max}, \infty\ [$,	the OSI95 TS provider should be able to accept a T-DATA request interaction at any time, otherwise it has to shut down the TC as soon as possible.

Table 3.2 Constraints on the OSI95 TS provider as regards the throughput

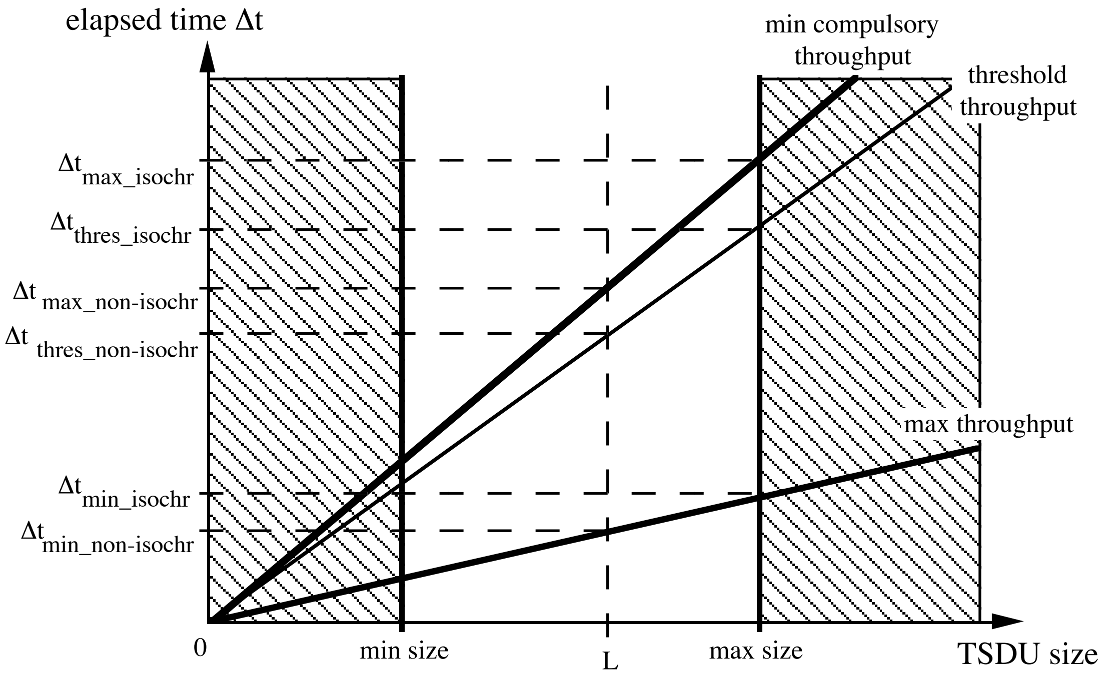

Fig. 3.1 Δt_{min}, Δt_{thres} and Δt_{max} for a non-isochronous or an isochronous traffic

If a sending OSI95 TS user generates an isochronous traffic, that is if it produces TSDUs (possibly of variable length) at a fixed rate, this sending OSI95 TS user may want the OSI95 TS provider to be able to accept the submission of TSDUs at the same rate characterized by the constant time interval Δt_{isochr} between the submission of successive TSDUs whatever their length. For this purpose, the OSI95 TS user may negotiate with the OSI95 TS provider a minimum compulsory (resp. threshold) throughput value given by the ratio of the maximum size authorized for the TSDUs to the constant time interval Δt_{isochr}. With such a minimum compulsory (resp. threshold) throughput value, if T_0 is the time at which the last T-DATA request interaction occurred, the sending OSI95 TS user is assured that the OSI95 TS provider will be ready to accept another T-DATA request interaction at the latest at time $T_0 + \Delta t_{isochr}$ in the case where the TSDU submitted in the interaction at time T_0 is of maximum size and even earlier in the case where the TSDU is of smaller size, or will disconnect (resp. report to the OSI95 TS users) otherwise. However, by forcing the OSI95 TS provider to be able to accept a new T-DATA request primitive before $T_0 + \Delta t_{isochr}$ when the last TSDU is not of maximum size, an unnecessary constraint is imposed on it. In fact, for an actually isochronous traffic, forcing the OSI95 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. This is why for each open direction of data transfer, the sending OSI95 TS user informs the OSI95 TS provider of the type of traffic it intends to submit, by selecting either of two traffic type options: “non-isochronous” or “isochronous”.

The traffic type indicator is not negotiated at all: the traffic type option is imposed by an OSI95 TS user for its output direction of data transfer on a TC. If the traffic type option is fixed to “non-isochronous”, then the Δt_{max}, Δt_{thres} and Δt_{min} are calculated as explained previously, i.e. by taking the size L of the previous submitted TSDU into account (Figure 3.1). By contrast, if the traffic type option is fixed to “isochronous”, then the Δt_{max}, Δt_{thres} and Δt_{min} are simply taken equal to Δt_{max_isochr}, Δt_{thres_isochr} and Δt_{min_isochr} (Figure 3.1) as if all the TSDUs were of maximum size.

For each open direction of data transfer of a TC, the transit delay is defined as *the time elapsed between the occurrence of a T-DATA request primitive at a TSAP and the occurrence of the corresponding T-DATA indication primitive at the peer TSAP*, just like in the current ISO/IEC standard. This definition means that a measure of the transit delay is associated with each invocation of the TSDU transfer facility.

After a QoS negotiation, the meaningful values regarding the transit delay for an open direction of data transfer are the maximum compulsory value and the threshold value. The following relation must always be verified: 0 ≤ negotiated threshold transit delay value ≤ negotiated maximum compulsory transit delay value. The maximum compulsory value for the transit delay is such that it will never be exceeded at any invocation of the TSDU transfer facility while it applies: if the OSI95 TS provider cannot deliver a T-DATA indication anymore without violating this value, it has to shut down the TC as soon as possible (via T-DISCONNECT primitives). The threshold value for the transit delay is such that every time the OSI95 TS provider exceeds this QoS value for a particular invocation of the TSDU transfer facility while this value applies, it has to report this fact to the OSI95 TS users as soon as possible (via T-REPORT indication primitives) but the TC is not shut down.

For each open direction of data transfer of a TC, the transit delay jitter is defined as *the difference between the longest and the shortest transit delays observed on this direction since the successful completion of the last QoS negotiation on the sending OSI95 TS user's side*. According to this definition, a measure of the transit delay jitter is thus associated with each invocation of the TSDU transfer facility, but this measure of the transit delay jitter depends upon the transit delays observed for all invocations of the TSDU transfer facility (including the current one) since the successful completion of the last QoS negotiation on the sending OSI95 TS user's side.

After a QoS negotiation, the meaningful values regarding the transit delay jitter for an open direction of data transfer are the maximum compulsory value and the threshold value. The following relation must always be verified: 0 ≤ negotiated threshold transit delay jitter value ≤ negotiated maximum compulsory transit delay jitter value. The role of these two values for the transit delay jitter is similar to the one of the corresponding values for the transit delay.

4 Flexibility of the Error Control on a TC

The ISO/IEC Connection-mode Transport Service is always error-free during the data transfer phase, in the sense that there are normally no losses, no corruptions of contents and no duplications of transferred TSDUs except for what is tolerated by the residual error rate QoS parameter.

Even if the connection mode has to provide a full error recovery capability to many applications, and to bulk data transfer applications in particular, such a capability may be quite unsuitable in case of multimedia applications for example. In several multimedia applications indeed, the receiving end is able to manage some losses but has stringent transit delay and transit delay jitter requirements, so that a full error recovery capability at the transport service level based on retransmissions at the protocol level is not viable. Our approach in OSI95 has been to permit a tune-up (which may lead ultimately to a suppression) of the error control [BLL 92]. This has been done by means of a new QoS parameter: the error control policy.

4.1 Description of the Possible Error Control Policies

An error control policy is negotiated for either open direction of data transfer on a TC regardless of and separately from the negotiation of the error control policy for the reverse direction (if the TC is bi-directional). The error control policy for a direction of data transfer consists in the combination of a TSDU corruptions policy and a TSDU losses policy [Bag 93a].

The TSDU losses policy is specified qualitatively by selecting one of two options:

a) losses of (parts of) TSDUs not accepted;
b) losses of (parts of) TSDUs accepted but indicated.

The TSDU corruptions policy is also specified qualitatively by selecting one of three options:

a) corruptions of contents in (parts of) TSDUs not accepted;
b) corruptions of contents in (parts of) TSDUs accepted but indicated;
c) corruptions of contents in (parts of) TSDUs accepted and not indicated.

The six possible combinations result in six different error control policies. Table 4.1 shows the error rates that are meaningful in each of the six error control policies.

Error control policy		LOSSES	
		Losses not accepted	**Losses accepted but indicated**
CORRUPTIONS	**Corruptions not accepted**	1 — —	4 Lost TSDU error rate
	Corruptions accepted but indicated	2 Corrupted TSDU error rate	5 Corrupted TSDU error rate Lost TSDU error rate
	Corruptions accepted and not indicated	3 — —	6 Lost TSDU error rate

Table 4.1 The six possible error control policies and the corresponding meaningful error rates

When losses (resp. corruptions of contents) of (parts of) TSDUs are not accepted, all losses (resp. corruptions of contents) have to be recovered. When losses (resp. corruptions of contents) of (parts of) TSDUs are accepted provided they are indicated, these losses (resp. corruptions of contents) do not need to be recovered, within the limits expressed by the lost (resp. corrupted) TSDU error rate requirements. They do not need to be recovered but, of course, they may be. However, they necessarily need to be detected so that they can be indicated. By contrast, when corruptions of contents in (parts of) TSDUs are accepted without having to be indicated, these corruptions of contents do not need to be recovered or even detected but, of course, they may be.

Error control policy (1) is the classical full error recovery policy (i.e. no errors at all are tolerated). When this policy is adopted for an open direction of data transfer, the OSI95 TS provider has to preserve unchanged the boundaries and the contents of all submitted TSDUs. That is, any TSDU delivered to the receiving OSI95 TS user via a T-DATA indication shall have the same number of octets and the same value for each octet as the TSDU received from the sending OSI95 TS user in the corresponding T-DATA request.

With error control policy (2), all losses of (parts of) TSDUs have to be recovered, whereas corruptions of contents in (parts of) TSDUs do not need to be recovered, within the limits expressed by the corrupted TSDU error rate requirements, but need to be detected. With error control policy (3), all losses of (parts of) TSDUs have to be recovered, whereas corruptions of contents in (parts of) TSDUs do not need to be recovered or even detected. With error control policy (4), losses of (parts of) TSDUs do not need to be recovered, within the limits expressed by the lost TSDU error rate requirements, but need to be detected, whereas all corruptions of contents in (parts of) TSDUs have to be recovered. With error control policy (5), losses and corruptions of contents of (parts of) TSDUs do not need to be recovered, within the limits expressed respectively by the lost and corrupted TSDU error rate requirements, but both need to be detected. Finally, with error control policy (6), losses of (parts of) TSDUs do not need to be recovered, within the limits expressed by the lost TSDU error rate requirements, but need to be detected, whereas corruptions of contents in (parts of) TSDUs do not need to be recovered or even detected.

In the case where a whole submitted TSDU is lost through the OSI95 TS provider and the loss is not recovered by the OSI95 TS provider, a zero-length TSDU shall be delivered to the receiving OSI95 TS user instead. In all other cases (i.e. when a non-zero-length TSDU is delivered), the TSDU delivered to the receiving OSI95 TS user via a T-DATA indication shall have the same number of octets as the TSDU received from the sending OSI95 TS user in the corresponding T-DATA request, but the value of some octets may have been altered by the OSI95 TS provider:

- If a part of a submitted TSDU (not the whole TSDU) is lost through the OSI95 TS provider, and if the loss is not recovered by the OSI95 TS provider, then each octet of this part in the delivered TSDU will have a value that is either any value or a specific default value (called dummy character value), the value being provided by the OSI95 TS provider regardless of the correct value of the octet in the submitted TSDU. In the worst case, all the octets of this part will have an incorrect value in the delivered TSDU.
- If a part of a submitted TSDU (possibly the whole submitted TSDU) is corrupted through the OSI95 TS provider, and if the corruption of contents is not recovered by the OSI95 TS provider, then at least one octet of this part in the delivered TSDU will have a value different from the correct value in the submitted TSDU. In the worst case, all the octets of this part will have an incorrect value in the delivered TSDU. In other words, a number of octets between 1 and the length of this part inclusive have an incorrect value in the delivered TSDU.

As already said above, when error control policy (4) applies, losses of (parts of) TSDUs do not need to be recovered, within the limits expressed by the lost TSDU error rate requirements, but need to be detected, whereas all corruptions of contents in (parts of) TSDUs have to be recovered. One possible way to recover from corruptions of contents of (parts of) TSDUs in such a case is to consider and process the corrupted parts as if they had been lost. Other ways are conceivable however.

The error control policy for either direction of data transfer on a TC is negotiated between the OSI95 TS users only. This is because the choice of the error control policy must depend upon the needs of the application only, and must not be influenced by the OSI95 TS provider which has to be able to provide full error recovery if required. The error control policy proposed by the calling OSI95 TS user in the T-CONNECT request (and replicated unchanged in the T-CONNECT indication) may only be strengthened by the called OSI95 TS user in the T-CONNECT response (which means that, in table 4.1, the square corresponding to the policy finally selected by the called OSI95 TS user may not be on the right or under the square corresponding to the policy proposed by the calling OSI95 TS user). This introduces a partial order among the error control policies.

4.2 Role of the Other Error Control QoS Parameters

In OSI95, besides the error control policy, there are four other error control QoS parameters: the lost TSDU error rate, the corrupted TSDU error rate, the dummy replacement indicator and the dummy character. Each of these parameters is specified for either open direction of data transfer on a TC regardless of and separately from its specification for the reverse direction (if the TC is bi-directional) [Bag 93a].

The successful completion of a QoS (re-)negotiation corresponds to the occurrence of a T-CONNECT confirm or a T-RENEGOTIATE confirm (see section 5) on the calling side, and to the occurrence of a T-CONNECT response or a T-RENEGOTIATE response (see section 5) on the called side. All the TSDUs submitted via T-DATA requests by a same sending OSI95 TS user between two successive successful completions of a QoS (re-)negotiation make up a sequence of TSDUs. The TSDUs of a sequence can be numbered conceptually in the order in which they are submitted to the OSI95 TS provider with consecutive integers starting from zero. Then, a sequence of TSDUs can be partitioned into successive N-subsequences of TSDUs. Given an integer $N \geq 1$, an N-subsequence of TSDUs is a series of N TSDUs that belong to a same sequence, that have consecutive increasing numbers and such that the number of the first TSDU of the series is a multiple of the integer N. A TSDU error rate is expressed as a ratio of N1 to N (N1 and N being integers and verifying $0 \leq N1 \leq N$ with $N \geq 1$) where N1 is the number of TSDUs in error in a series of N TSDUs.

More precisely, for an open direction of data transfer on a TC, the lost (resp. corrupted) TSDU error rate is defined as the number of TSDUs, out of an N-subsequence (resp. N'-subsequence) of TSDUs submitted by the sending OSI95 TS user to the OSI95 TS provider, which are delivered to the receiving OSI95 TS user with an indication that (parts of) their contents has been lost (resp. corrupted). It has to be stressed that the size N of the N-subsequences of TSDUs intervening in the measure of the lost TSDU error rate may be different from the size N' of the N'-subsequences of TSDUs intervening in the measure of the corrupted TSDU error rate. As previously shown in table 4.1, for either open direction of data transfer, the lost (resp. corrupted) TSDU error rate parameter is meaningful only if the error control policy negotiated for this direction specifies that losses (resp. corruptions of contents) of (parts of) TSDUs are accepted but have to be indicated.

The semantics adopted for the lost and corrupted TSDU error rates is the one adopted for the performance QoS parameters (i.e. the throughput, the transit delay and the transit delay jitter), giving the possibility to negotiate compulsory and threshold QoS values. The sole difference in comparison with the performance QoS parameters is in the negotiation scenario: the OSI95 TS provider may not intervene at all in the negotiation, so that the proposal of the calling OSI95 TS user in the T-CONNECT request as regards the lost and corrupted TSDU error rates is replicated unchanged in the T-CONNECT indication.

Moreover, for either open direction of data transfer on a TC, the dummy replacement indicator parameter is meaningful only if the error control policy negotiated for this direction specifies that losses of (parts of) TSDUs are accepted provided they are indicated. Either of two dummy replacement options may be selected by the receiving OSI95 TS user: "replacement with dummy character" or "no replacement with dummy character". The dummy replacement indicator is not negotiated at all: the dummy replacement option is imposed by an OSI95 TS user (in the T-CONNECT request for the calling OSI95 TS user and in the T-CONNECT response for the called OSI95 TS user) for its input direction of data transfer on a TC.

If the dummy replacement option is fixed to "replacement with dummy character" by the receiving OSI95 TS user, then, in every TSDU received with an indication of loss, all lost octets in the zone of the TSDU affected by errors have a default value (called the dummy character value) provided by the OSI95 TS provider regardless of the correct value of the octets in the submitted TSDU. By contrast, if the dummy replacement option is fixed to "no replacement with dummy character", then, in every TSDU received with an indication of loss, each lost octet in the zone of the TSDU affected by errors may have any value provided by the OSI95 TS provider regardless of the correct value of the octet in the submitted TSDU.

Lastly, for either open direction of data transfer, the dummy character parameter is meaningful only if the dummy replacement indicator is meaningful and if the dummy replacement option is set to "replacement with dummy character". Any dummy character value between 0 and 255 (8-bit unsigned value) may be selected by the receiving OSI95 TS user. The dummy character is not negotiated at all: its value is imposed by an OSI95 TS user for its input direction of data transfer on a TC.

4.3 Use of Transfer Status Parameters in the T-DATA indication

Two transfer status parameters, viz. the error type indicator and the affected zone, have been inserted in the T-DATA indication to cope with error control [Bag 93a].

The possible error type indicator options in a T-DATA indication depend upon the error control policy negotiated for the corresponding direction of data transfer:

a) with error control policies (1) & (3), the sole possible error type indicator option is "no errors signalled in the TSDU";
b) with error control policies (4) & (6), two error type indicator options are possible: "no errors signalled in the TSDU" or "losses signalled in the TSDU";
c) with error control policy (2), two error type indicator options are possible: "no errors signalled in the TSDU" or "corruptions of contents signalled in the TSDU";
d) with error control policy (5), four error type indicator options are possible: "no errors signalled in the TSDU", "losses signalled in the TSDU", "corruptions of contents signalled in the TSDU" or "both losses and corruptions of contents signalled in the TSDU".

Any TSDU delivered to the receiving OSI95 TS user with the error type indicator indicating "losses signalled in the TSDU" or "both losses and corruptions signalled in the TSDU" increments by 1 the number of (partly) lost TSDUs pertaining to the current N-subsequence of TSDUs. In the same way, any TSDU delivered to the receiving OSI95 TS user with the error type indicator indicating "corruptions signalled in the TSDU" or "both losses and corruptions signalled in the TSDU" increments by 1 the number of (partly) corrupted TSDUs pertaining to the current N'-subsequence of TSDUs.

The affected zone parameter in a T-DATA indication is not meaningful when the associated error type indicator indicates “no error signalled in the TSDU”, or when the whole submitted TSDU is considered to be lost and a zero-length TSDU is delivered instead (which is necessarily accompanied by an error type indicator that indicates “losses signalled in the TSDU”). In all other cases, this parameter shall delimit the zone affected by errors, by giving a pair of integers which are the number of the first octet of the zone and the number of the last octet of the zone (the octets of the submitted TSDU are supposed to be numbered consecutively starting from 1).

Even when the corrupted and lost parts of a TSDU are not adjacent, the affected zone in a delivered TSDU is a continuous block of octets. The affected zone is the smallest continuous block of octets containing all corrupted and lost parts of the TSDU. In the example of figure 4.1, the affected zone is determined by the pair of integers (b, e), despite the fact that there are parts of the TSDU whose integrity is preserved between the octets numbered b and e. If it is deemed useful, the transfer status information could be refined in a future definition, in order to give details about the corrupted, lost and safe parts in the affected zone.

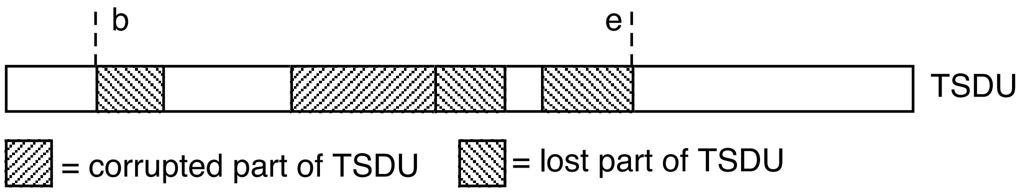

Fig. 4.1 Affected zone parameter in the transfer status

5 QoS Re-negotiation Facility

In the ISO/IEC Connection-Mode Transport Service, the QoS is negotiated once and for all at the TC establishment. This means that the negotiated QoS always applies for the whole TC lifetime.

To respond to certain requirements formulated in [LCH 93], we have decided to introduce a QoS re-negotiation facility in our OSI95 Enhanced Connection-Mode Transport Service definition.

5.1 Who Will Be Allowed to Start a Re-negotiation ?

The introduction of such a facility poses a first problem: who will be allowed to initiate a re-negotiation ? If both users of a given TC are permitted to start a re-negotiation, we will face a major difficulty when they invoke the facility at the same time. Let us remind that, as regards the TC establishment, colliding TC set-up requests lead to the establishment of distinct TCs. Bypassing the difficulty in a similar way is of course impossible in case of re-negotiation.

An obvious solution to the problem is to suppress it by precluding one of the two OSI95 TS users to initiate a re-negotiation on the TC. Therefrom comes another question: which user should have the right to start a re-negotiation whereas the other one would not have this right ? This could be decided during the TC set-up phase, through a form of negotiation by way of the T-CONNECT primitives. The sole OSI95 TS user authorized to initiate a re-negotiation could also be always the same one, either the calling or the called one. At first glance, it would be preferable to elect the calling OSI95 TS user, i.e. the initiator of the TC establishment, as the only possible initiator of a QoS re-negotiation on the TC. Whichever the OSI95 TS user authorized to start the re-negotiation, the other one could anyway use the out-of-band data transfer facility (see section 6) to indicate to its peer its wish to have a re-negotiation.

Another solution to the aforementioned problem is to allow both OSI95 TS users to initiate the re-negotiation, and thus to accept the possibility of colliding re-negotiation requests, while giving the precedence to one of them. In case of collision, the re-negotiation request from the OSI95 TS user that has not the precedence is simply discarded and the other request is processed normally. The different alternatives, proposed in the context of the previous solution to choose the only OSI95 TS user that should be authorized to start the re-negotiation, remain conceivable to choose the OSI95 TS user that should have the precedence.

In the current version of our connection-mode service, we have decided that only the initiator of the TC set-up, i.e. the calling OSI95 TS user, will have the right to start a QoS re-negotiation.

5.2 Sequence of Primitives for a Successful Re-negotiation

The sequence of OSI95 TS primitives for QoS re-negotiation has been inspired by the classical four-primitive sequence used for TC set-up. Nevertheless, because of the wish not to suspend the TSDU transfers over the TC while a QoS re-negotiation takes place, a fifth primitive named T-NEW-QOS indication has been added (Figure 5.1).

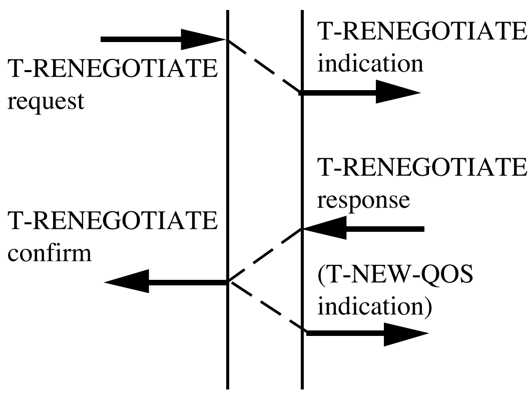

Fig. 5.1 Time sequence diagram for a successful re-negotiation

On the calling side, both the output and input traffics switch from the old QoS to the new one upon receipt of the T-RENEGOTIATE confirm. On the called side, the output traffic switch from the old QoS to the new one upon invocation of the T-RENEGOTIATE response but the additional T-NEW-QOS indication is necessary to indicate when the QoS of the input traffic changes.

The QoS re-negotiation primitives are used to re-negotiate the QoS as well as the TSDU size range in either or both open direction(s) of data transfer on an established TC. If the established TC is bidirectional and both directions of data transfer are still considered as open by the calling OSI95 TS user, this calling OSI95 TS user may choose to re-negotiate the QoS and the TSDU size range in either or both open direction(s). In all other cases, the re-negotiation will apply to the only direction still considered as open by the calling OSI95 TS user. The calling OSI95 TS user may never initiate a QoS re-negotiation when it considers both directions of data transfer as closed or closing[1].

The parameters of the T-RENEGOTIATE primitives are thus identical to those of the T-CONNECT primitives, except that the unnecessary address parameters are removed and that the "directions to open" parameter becomes a "directions to re-negotiate" parameter. The relations which have to be verified between the parameters of the different T-RENEGOTIATE primitives are summarized in table 5.1.

	PP-ECOTS primitives:			
Parameters:	T-RENEG… request	T-RENEG… indication	T-RENEG… response	T-RENEG… confirm
Directions to re-negotiate	M	M	M	M
Quality of service {…}	M	M*(N)*	M*(N)*	M*(=)*
TSDU size range	M	M*(N)*	M*(N)*	M*(=)*
OSI95 TS user-data	U	C*(=)*	U	C*(=)*

<u>Keys:</u> {…} : generic denomination gathering several parameters;
M : presence of this parameter or of this group of parameters in the primitive is mandatory;
U : inclusion of this parameter or of this group of parameters in the request or response primitive is a choice made by the OSI95 TS user;
C : presence of this parameter or of this group of parameters in the indication or confirm primitive is conditional, depending upon its inclusion in the preceding request or response primitive;
(=) : the value or option of the parameter in the primitive is identical to that of the corresponding parameter in the preceding primitive;
(N) : each value or option of the parameter in the primitive is obtained from the values or options of the corresponding parameter in the preceding primitive in accordance with the negotiation rules;

Table 5.1 Relations between the parameters of the different T-RENEGOTIATE primitives

[1] When an OSI95 TS user has invoked the graceful TC release facility (see [2] for more details) to release its output direction of data transfer gracefully, then, while it is waiting for the confirmation of the graceful release, it does not regard this direction as open any more, neither does it regard this direction as closed yet. Its output direction is considered to be in an intermediate state between the open and closed states, referred to as the closing state.

The T-NEW-QOS indication has no parameter. As its sole function is to indicate to the called OSI95 TS user when the QoS and TSDU size range characteristics relating to the calling-to-called direction of data transfer switch from the old ones to the new ones, this T-NEW-QOS indication shall be delivered only if the calling-to-called direction appears in the "directions to re-negotiate" parameter of the T-RENEGOTIATE confirm and if the called OSI95 TS user does not consider this direction as closing or closed before the delivery of the primitive.

5.3 Possible Sequences of Primitives for an Unsuccessful Re-negotiation

The QoS re-negotiation of any direction of data transfer has three possible outcomes: it succeeds, it is rejected by either the OSI95 TS provider or the called OSI95 TS user, or it is stopped if it becomes meaningless.

The QoS re-negotiation procedure may fail either due to the inability of the OSI95 TS provider to accept the TC with the new requested characteristics or due to the unwillingness or the inability of the called OSI95 TS user to accept the TC with the new requested characteristics. These two cases are described in figures 5.2 & 5.3. When the QoS re-negotiation involves both directions, the rejection of the proposal made by the calling OSI95 TS user for one direction will inescapably cause the rejection of the whole re-negotiation. When a QoS re-negotiation is rejected, the TC is of course maintained with its old QoS and TSDU size range characteristics.

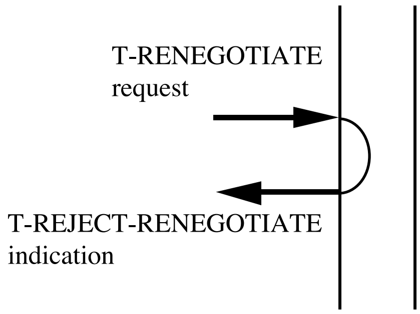

Fig. 5.2 Time sequence diagram for the rejection of a re-negotiation attempt by the OSI95 TS provider

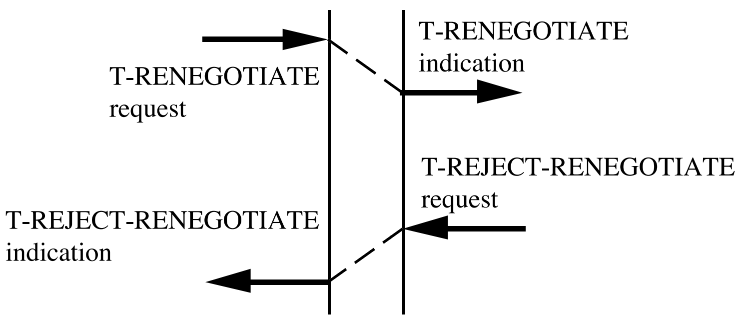

Fig. 5.3 Time sequence diagram for the rejection of a re-negotiation attempt by the called OSI95 TS user

The parameters of the T-REJECT-RENEGOTIATE primitives, and the relations which have to be verified between them, are summarized in table 5.2.

The reject reason is one of the following:

a) remote OSI95 TS user initiated;
b) OSI95 TS provider initiated due to any internal reason which may be of transient or permanent nature.

	PP-ECOTS primitives:	
Parameters:	T-REJECT-RENEG… request	T-REJECT-RENEG… indication
Reject reason		M
OSI95 TS user-data	U	C(=)

Keys: M : presence of this parameter or of this group of parameters in the primitive is mandatory;
U : inclusion of this parameter or of this group of parameters in the request or response primitive is a choice made by the OSI95 TS user;
C : presence of this parameter or of this group of parameters in the indication or confirm primitive is conditional, depending upon its inclusion in the preceding request or response primitive;
(=) : the value or option of the parameter in the primitive is identical to that of the corresponding parameter in the preceding primitive;
Blank : the parameter is absent.

Table 5.2 Relations between the parameters of the T-REJECT-RENEGOTIATE primitives

Any of the time sequence diagrams in figures 5.1, 5.2 & 5.3 may terminate before completion if the QoS re-negotiation is stopped. A QoS re-negotiation is stopped in the case where all direction(s) of data transfer indicated initially in the directions to re-negotiate parameter of the T-RENEGOTIATE request is (are) no more considered as open by the OSI95 TS user which should have issued or received the next expected primitive of the sequence.

The calling OSI95 TS user may not initiate a first QoS re-negotiation before the T-CONNECT confirm primitive has been issued. Moreover, the calling OSI95 TS user may not initiate a subsequent QoS re-negotiation until a T-RENEGOTIATE confirm or a T-REJECT-RENEGOTIATE indication primitive has been received in response to the previous T-RENEGOTIATE request primitive, indicating that the previous re-negotiation is successful or rejected, or until it considers the direction(s) of data transfer involved in the previous re-negotiation as closing or closed.

6 Out-of-Band Data Transfer Facility

In addition to the transfer of normal TSDUs, the OSI95 Enhanced Connection-Mode Transport Service provides for an exchange of out-of-band TSDUs between the OSI95 TS users. Like the exchange of normal TSDUs, the exchange of out-of-band

TSDUs may not take place before the successful completion of the TC set-up phase. Unlike the exchange of normal TSDUs, the exchange of out-of-band TSDUs is allowed even when the corresponding direction of data transfer is not considered as open. Transfers of out-of-band TSDUs can always occur in both directions simultaneously as long as the TC remains established, i.e. as long as both directions of data transfer of the TC are not considered as closed.

Since the interest of transmitting out-of-band TSDUs over TCs with no acknowledgement procedure at all seems quite limited to us, we have elected an out-of-band TSDU transfer facility which is very similar to the Enhanced Acknowledged Connectionless-mode Transport Service defined in detail in [Bag 93b]. In fact, the out-of-band TSDU transfer facility differs from this acknowledged connectionless-mode service only by the fact that the transmitted out-of-band TSDUs are automatically associated with an established TC by means of the TCEP (Transport Connection Endpoint) identification mechanism. This is why the T-OOB-DATA primitives, unlike the T-ACKDATA primitives (cf. [Bag 93b]) do not need any address parameter: all addressing information is provided in the T-CONNECT primitives, and the other connection-mode primitives makes use of the TCEP identification mechanism to identify the TC to which they pertain.

The QoS associated with an out-of-band TSDU is specified by the sending OSI95 TS user regardless of the QoS negotiated on the TC. The QoS associated with an out-of-band TSDU may thus be lower or higher than that for the TSDUs transferred via the T-DATA primitives. In particular, the monitoring of the compulsory and threshold QoS negotiated on the TC does not apply to the out-of-band TSDU transfers. An out-of-band TSDU is conveyed by the OSI95 TS provider exactly in the same way as would be an acknowledged connectionless-mode TSDU submitted with the same QoS in the Enhanced Acknowledged Connectionless-mode Transport Service.

With respect to error control, the OSI95 TS provider always preserves the integrity and the boundaries of the submitted out-of-band TSDUs. On the other hand, it may lose some submitted out-of-band TSDUs, may deliver out-of-band TSDUs in an order different from that in which they have been submitted and may duplicate out-of-band TSDUs. Out-of-band TSDUs may also pass or be passed by normal TSDUs.

Clearly, an identification mechanism has to be provided at the OSI95 TS interface if an OSI95 TS user and the OSI95 TS provider need to distinguish between several invocations of the out-of-band data transfer facility relating to a same TC. For this purpose, the TAEP (Transport Association Endpoint) identification mechanism which is also used by the OSI95 Acknowledged and Request/Response Connectionless-Mode Transport Services (cf. [Bag 93b]) is provided locally. Here however, unlike what is done in the connectionless-mode services where the TAEP identification mechanism is used at the TSAP level, this identification mechanism is used at the TCEP level!

Even though it has not exactly the same properties, the out-of-band data transfer facility may be an interesting alternative to the expedited data transfer facility. A recommended use of the out-of-band data transfer facility is in support of the QoS re-negotiation. With the current QoS re-negotiation facility, only the OSI95 TS user that initiated the TC may start a re-negotiation. When the other OSI95 TS user wants to re-negotiate the QoS on the TC, it could inform the OSI95 TS user that has the right to re-negotiate and provide this OSI95 TS user with the relevant information by way of the out-of-band channel corresponding to the TC.

7 Conclusions

A lot of modern applications (multimedia applications for instance) may not be completely satisfied with the capabilities currently available at the transport service interface in the networks, especially for what regards the multiple QoS aspects. In this light, the modifications and enhancements in the field of QoS has been a very important facet of our work on a new connection-mode transport service within the OSI95 project. The removal of QoS parameters deemed unmanageable in practice, the change in the definition of the throughput, the introduction of a transit delay jitter, the definition of a new semantics for the performance QoS parameters, the introduction of specific QoS parameters to deal with error control, the introduction of a TSDU size range information parameter and finally the definition of a QoS re-negotiation facility all aim at better meeting the needs of modern applications.

The ability to open uni-directional TCs, as well as bi-directional TCs, in the OSI95 Connection-Mode Transport Service reflects our choice to adopt the concept of "simplex virtual crcuit" (corresponding to one direction of data transfer) as the basic building block for a TC. By the way, this choice has had consequences on the design of the TC release facilities (see [BaL 93] for the details).

Lastly, despite the fact that it has not exactly the same properties, the ouf-of-band data transfer facility has taken the place of the classical expedited data transfer facility in the current version of the OSI95 Connection-Mode Transport Service. The requirements of the applications in the field of priority data transfer need further investigations however, and therefore it remains an open question to know whether the out-of-band data transfer facility (if kept in a next version) should remain a substitute for or should come in addition to the expedited data transfer facility.

References

[Bag 93a] Y. Baguette, "Enhanced Transport Service Definition (Informal specification in English - Version 1)", Internal ref.: S.A.R.T. 93/03/15 (also circulated as ISO/IEC JTC1/SC6/WG4 N 822), Université de Liège, Institut Montefiore B28, B-4000 Liège, Belgium, Feb. 5, 1993.

[Bag 93b] Y. Baguette, "The OSI95 Connectionless-Mode Transport Services", *The OSI95 Transport Service with Multimedia Support on HSLAN's and B-ISDN*, A. Danthine, ed., Springer Verlag, 212-224

[BaL93] Y. Baguette, G. Leduc, "The Connection Release Facilities in the OSI95 Transport Service", *The OSI95 Transport Service with Multimedia Support on HSLAN's and B-ISDN*, A. Danthine, ed., Springer Verlag, 199-211

[BLL 92] Y. Baguette, L. Léonard, G. Leduc, A. Danthine, O. Bonaventure, "OSI95 Enhanced Transport Facilities and Functions", Technical report for ESPRIT II project 5341 (OSI95), Ref.: OSI95/Deliverable ULg-A/P/V1, Université de Liège, Institut Montefiore B28, B-4000 Liège, Belgium, Dec. 1992.

[DBL 92] A. Danthine, Y. Baguette, G. Leduc, "Issues Surrounding the Specification of High-Speed Transport Service and Protocol", Technical report for ESPRIT II project 5341 (OSI95), Ref.: OSI95/ULg/A/15/TR/P/V2 (also submitted as ISO/IEC JTC1/SC6 N 7312 through the Belgian National Body at the SC6 interim meeting, Paris, France, Feb. 10-13, 1992), Université de Liège, Institut Montefiore B28, B-4000 Liège, Belgium, Jan. 1992.

[DBL 93] A. Danthine, O. Bonaventure, G. Leduc, "The QoS Enhancements in OSI95", *The OSI95 Transport Service with Multimedia Support on HSLAN's and B-ISDN*, A. Danthine, ed., Springer Verlag, 125-150

[ISO 8072] ISO/IEC JTC1: "Information Processing Systems - Open Systems Interconnection - Transport Service Definition", ISO 8072, 1986.

[ISO 8807] ISO/IEC JTC1: "Information Processing Systems - Open Systems Interconnection - LOTOS - A Formal Description Technique Based on the Temporal Ordering of Observational Behaviour", ISO 8807, 1989.

[LCH 93] H. Leopold, A. Campbell, D. Hutchinson, N. Singer, "Distributed Multimedia Communication System Requirements", *The OSI95 Transport Service with Multimedia Support on HSLAN's and B-ISDN*, A. Danthine, ed., Springer Verlag, 64-81

[LeD 93] G. Leduc, A. Danthine, "On the Provision of a Fast Connect Facility in a Connection-Mode Transport Service",*The OSI95 Transport Service with Multimedia Support on HSLAN's and B-ISDN*, A. Danthine, ed., Springer Verlag, 225-238