24

Service Specific Connection Oriented Protocol (SSCOP) with no credit limitation

Gérard Hébuterne
Institut National des Télécommunications,
gerard.hebuterne@int-evry.fr

Wei Monin
France Telecom/CNET
wei.monin@lannion.cnet.fr

Abstract

An error correction procedure is currently being defined by ITU-T under the name of "Service Specific Connection Oriented Protocol". It is to be used in the B-ISDN in conjunction with AAL5 in order to provide an error-free data link layer. We study the performances of the procedure in terms of end-to-end delay. An analytical model is presented, which allows to address various traffic engineering aspects for realistic loss probability figures, and especially the influence of the Poll period.

Keywords: Error correction, Selective retransmission, SSCOP, B-ISDN.

1 Introduction

According to ITU-T Recommendation I.121 [4], the future Broadband ISDN is to be based upon ATM, which has thus to be seen as an universal and powerful transport network. To achieve this goal, and to attain its maximum efficiency, the ATM layer only provides basic transfer functions. Especially, no error correction is performed at this level. This is because such a function is not of "universal" nature: real time services must not rely on error correction, for instance.

L. Mason et al. (eds.), *Broadband Communications*

However, there does exist services which strongly rely on error-free connections. Signaling is one of such services, and ITU-T has undertaken the definition of the "Service Specific Connection Oriented Procedure" (SSCOP). SSCOP is to be used as Service Specific Convergence Sublayer of the ATM Application Layer type 5 (AAL5). Although SSCOP has been devised first for signaling needs, its use is by no means restricted and any data service based upon the AAL5 may rely on it.

In order to operate an error-free data connection, various engineering rules have to be devised. Quantiles of end-to-end delay are probably the most important Traffic-related QOS figures – allowing to define response time of signaling functions; also, their value are necessary for a correct dimensioning of timeouts, etc. A preliminary simulation study [3] allowed to characterize the protocol behaviour for high (and rather unrealistic) loss figures.

In this paper, we give an analytical derivation for the probability distribution of the end-to-end transfer delay for the Selective Retransmission SSCOP. The expression allows to estimate the quantiles of the delay, as a function of loss probability, of the round-trip delay and of the POLL period, for loss figures corresponding to what can be expected from nowadays networks.

2 Protocol Description

The description given here conforms with the ITU-T references of SG XI, May 1993 [5], both as SDL or textual parts. Only the data transfer part is of concern.

SSCOP ensures reliable data transfer between users of the ATM Application Layer (AAL). When the connection is established, the Sender may send data as long as it has credits available, which correspond to available buffer capacity at the Receiver's end. In case of lost data, the Receiver requests retransmission to the Sender, according to a selective retransmission scheme (see e.g. [1]).

The protocol makes use of three types of Protocol Data Units (PDUs): Sequences Data PDUs (SD PDU), Poll PDUs and Stat PDUs. SD PDUs, refered to as data packets or frames, are variable-length packets which carry user's data, and are identified by their sequence number. Poll PDUs are periodically generated by the Sender to request information about the state of the Receiver, and especially about possibly lost SD PDUs. Stat PDUs are either "sollicited" (sent by the Receiver on receipt of a Poll PDU) or "unsollicited" (sent by the Receiver on detection of a loss, by detecting a gap in sequence numbers).

2.1 Data Transfer using SSCOP

1. Upon arrival of a data block, a SD PDU is generated and buffered. As soon as the Sender is available, the SD PDU is sent. At the same time, it is copied in a retransmission buffer, along with the current value of the Poll period.

2. The Sender periodically generates and sends a Poll PDU. Each Poll PDU is numbered and carries the number of the last sent PDU.

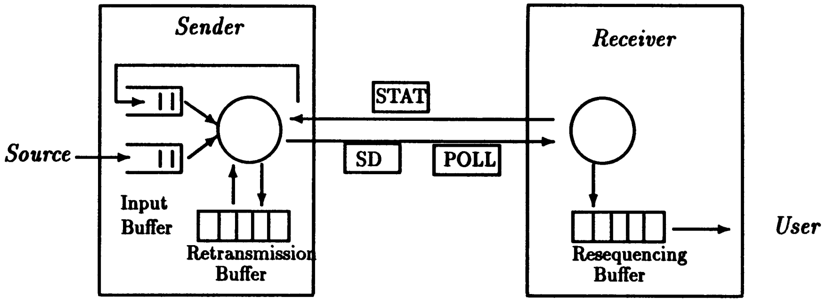

Figure 1: The elementary data link

3. Upon receipt of a POLL PDU, the Receiver sends back to the Sender a "Sollicited STAT PDU", which acknowledges correctly received frames and tells the Sender a list of missing ones.

4. Correctly received SD PDUs are immediately delivered to the user, provided all SD PDUs of lower number have been already received. If this is not the case, the SD PDU is stored until missing PDUs are correctly received.

   If the number of a received SD PDU is strictly greater than the expected one, the Receiver detects loss. It then sends back to the Sender a "unsollicited STAT PDU", reporting the loss.

5. Upon receipt of a STAT PDU, the Sender erases all correctly received SD PDUs. If missing SD PDUs are reported, the Sender retransmits them immediately if the STAT PDU is unsollicited. In case of a "sollicited STA PDU", it compares the number of the POLL PDU given in the STAT with the one stored in the buffer along with the SD PDU. Only if this last is lower will the retransmission be triggered. This avoids redundant retransmissions while allowing for detection of multiple losses (loss of a resent frame).

   Each time a SD PDU is retransmitted, the Poll number associated with it in the retransmission buffer is updated.

Remark: The "Flow Control" Procedure
The receiving end transmits in the STAT PDUs the amount of credits the Sender is allowed to use. This corresponds to the memory available in the resequencing buffer, and can serve two different purposes: either the memory is of limited and variable size (a common memory dynamically shared among different "Receivers", for instance), or the end user may control the speed at which the Sender is working.
In what follows, we assume no restriction due to credit management.

2.2 Parameters, Measures of Performance and Assumptions

The following variables denote the parameters used in the study, and the measures of performance:

- λ is the arrival rate of new frames in the input queue (number of frames per time unit). The arrival process is assumed to be Poisson.
- T^{prop} is the *propagation delay* (time needed for a bit to travel from the Sender to the Receiver). It depends on the distance between end users and on the underlying network (e.g. ATM connections probably give rise to variable propagation delays). One must also take into account processing and switching delays.
- T^{RT} is the *Round-Trip delay*: the time needed for a single bit to travel to the Receiver and to come back to the Sender. If processing delays are neglected, this corresponds to twice the propagation delay. The correspondance between delays and distances has to be made for each particular network
- T^{poll} is the time interval between two successive Poll PDUs. They are assumed to be sent with a fixed periodicity.
- p is the probability for a frame to be lost or errored. It is assumed independent of its length, for simplicity. It is also independent of the distance.
- The *End-to-End Transfer Delay* is defined as the time elapsed between the instant the first bit of the frame enters the input queue, and the instant the frame is made available to the User.
- The *End-to-End Additional Delay* (additional delay for short) is defined as the Transfer Delay decreased by the propagation delay and by the emission time; it accounts for any supplementary waiting a frame incurs due to retransmissions (of itself, or of previous frames, since the frames are to be delivered in sequence).

See figure 2 for the definition of delays: the transfer delay is denoted as w' and the additional delay as w.

3 The End-to-End Delays

3.1 Derivation of the End-to-End Additional Delay

For the frame numbered n, let w'_n be the transfer delay: see Figure 2. We assume that frames in sequence are immediately read by the user, so that:

- if the frame is sent in a "no-error period" (e.g. frame #1 in the figure), w'_n is simply the sum of the emission time and the propagation time;
- if the frame is sent in an "error period" (e.g. frame #3), w'_n ends as soon as the error is corrected: in Figure 2, a "batch" is observed, composed with frames #2, 3, etc. when frame #2 is received.

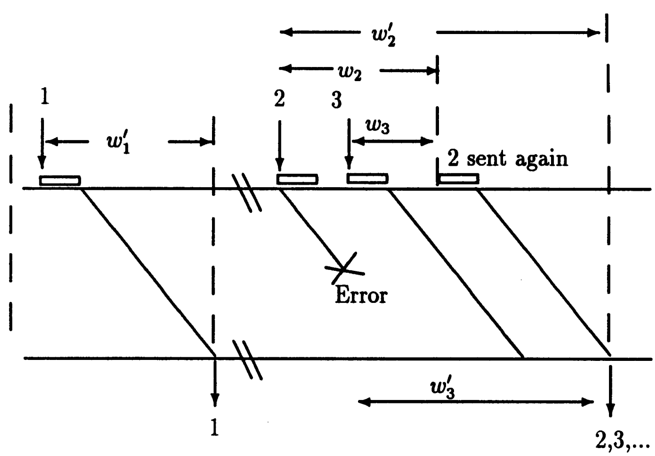

Figure 2: End-to-end transfer delay, End-to-end additional delay

The analysis considers the additional delay, defined as $w_n = w'_n - T^{\mathrm{prop}} - s_n$, with s_n representing the emission time and T^{prop} the propagation delay.
Let us define the *successfull sending* of a frame as the event: beginning of the last emission of the frame. Note that even if the frame is sent successfully it may have to wait in the resequencing queue. The "equivalent" service time of a frame, x_n, is defined as the time interval between the first sending and the successfull sending. x_n represents the contribution of the frame to the overall queueing.
Let t_n be the time elapsed between sending of $n-1$-th and n-th frame:

- if frame $\#n$ is such that $x_n = 0$ (it incurs no loss), then $w_n = w_{n-1} - t_n$, or 0 if the expression is negative.
- if frame $\#n$ is lost, and if x_n is smaller than $w_{n-1} - t_n$, then $w_n = w_{n-1} - t_n$ - that is, the frame $\#n$ is corrected before the preceding error is corrected, and it remains in the queue due to this previous error (this is the case for instance if the previous error needs more than one reemission to be corrected).
- if frame $\#n$ is in error, and if x_n is larger than $w_{n-1} - t_n$, then $w_n = x_n$ (this is the case for Frame #2 in the Figure).

Finally, one has the following relation between the w_n, t_n and x_n's:

$$w_n = \max\{x_n, w_{n-1} - t_n\} \tag{1}$$

From the recurrence relation, one can derive an equation for $W(t)$, the probability distribution function of w. Let:

- $W_n(t) = P\{w_n \leq t\}$, and $W(t) = \lim_{n\to\infty} W_n(t)$;

- $F_n(t) = P\{x_n \leq t\}$, and $F(t) = \lim_{n\to\infty} F_n(t)$;
- t_n is distributed according to a Poisson process with rate λ.

$$\begin{aligned} Pr\{w_n \leq t\} &= Pr\{x_{n-1} \leq t \ \underline{and}\ w_{n-1} - t_n \leq t\} \\ &= Pr\{x_{n-1} \leq t\}.Pr\{w_{n-1} \leq t_n + t\} \\ &= Pr\{x_{n-1} \leq t\} \int_u Pr\{w_{n-1} \leq t + u\}.Pr\{t_n \in [u, u + du[\} \end{aligned}$$

The second equation follows from independence assumptions. The third equation is obtained by conditioning on u, the value of t_n. Assuming the existence of a stationary limit:

$$W(t) = F(t). \int_u \lambda e^{-\lambda u} W(t+u) dx \tag{2}$$

The equation is solved by putting first $Q(t) = e^{-\lambda t} W(t)$, yielding

$$Q(t) = \lambda F(t) \int_{u \geq t} Q(u) du$$

One then introduce the function $H(x) = \int_{u \geq x} Q(u)du$, that is $Q(x) = -H'(x)$; the equation becomes then $H'(x) = -\lambda F(x).H(x)$, the solution of which is immediate.
The normalisation condition ($W(\infty) = 1$) gives finally:

$$\begin{aligned} W(x) &= F(x) e^{-\lambda G(x)} \\ \text{where} \quad & G(x) = \int_x^{\infty} [1 - F(u)] du \end{aligned} \tag{3}$$

Actually, Equation 3 gives only the "transmission part" of the delay. Frames may be delayed in the input buffer, due to other frames being sent. In the case where no credit management can block the Sender, one may assume independence between the two components of the end-to-end delay; the delay in the Input Buffer is then estimated as the sojourn time in a M/GI/1 queue (the service time being the transmission time of the frame). Preliminary simulation studies have shown [3] this delay to be negligible, and we omit it in the following. Note that with a 10 Mbit/s link rate, it takes $100\mu s$ to send a 1 kbit-packet.
The independence assumption implies also that the total delay is the ***convolution*** of the two distributions.

3.2 The components of the equivalent service time

The End-to-End Waiting Time is null if the packet is successfully received at the first sending. Figure 3 shows the components of the delay in case of a single error and in case of multiple errors.

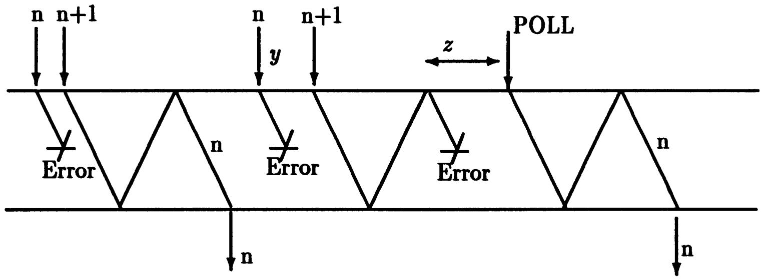

Figure 3: End-to-end delay : one or two successive errors

Assume the packet numbered n is lost. The next packet $n + 1$ to be sent, or the next POLL PDU, allows the receiver to detect the loss. Either the U-Stat (as in the left part of Figure 3) or the S-Stat reports the event to the sender and the missing packet is then sent again.
Let us denote as y the random variable which measures the time elapsed between the packet sending and the next successfull sending which signals the previous loss (either the data or POLL). One sees that y is the minimum of the two variables: remaining time until next POLL, and time to the next successfull STAT. In fact, the next packet or the STAT could be lost, too. In this case, packet $n+2$ detects the event, etc. It is possible to show that this is equivalent to assume that the interdepartures between successfully sent packets is exponentially distributed with parameter $\lambda(1-p)$ instead of λ. The distribution of y is easily obtained, taking into account the independence between data packets and POLL PDUs.
In this single-retransmission configuration, the waiting time is the sum $y + T^{\mathrm{RT}}$.
Assume now that the second sending of packet n is lost too. In this case, only the POLL PDU which follows the second sending allows the detection of the loss. The detection scheme conforms with the right hand part of Figure 3. The additional waiting time becomes $y + 2T^{\mathrm{RT}} + z$, where z stands for the delay between the S-STAT PDU is received and the next POLL is sent.
Should a third error occur, the next POLL detects it, with this time an additional waiting time of $y + 2T^{\mathrm{RT}} + z + \tau$, where τ stands for the delay between the S-STAT and the next POLL. The general case (although rather theoretical) follows easily.
From these relations, the distribution of the variable X can be derived. The complete expression is not given here. In fact, the actual estimation of the distribution depends on the relative values of T^{RT}, T^{prop}, etc. See [2] for more details about the actual calculation process.

4 Application to Traffic Engineering

The following results show the influence of the control variables. Of special concern is the value to be given to T^{poll}.

4.1 Influence of the Loss Probability

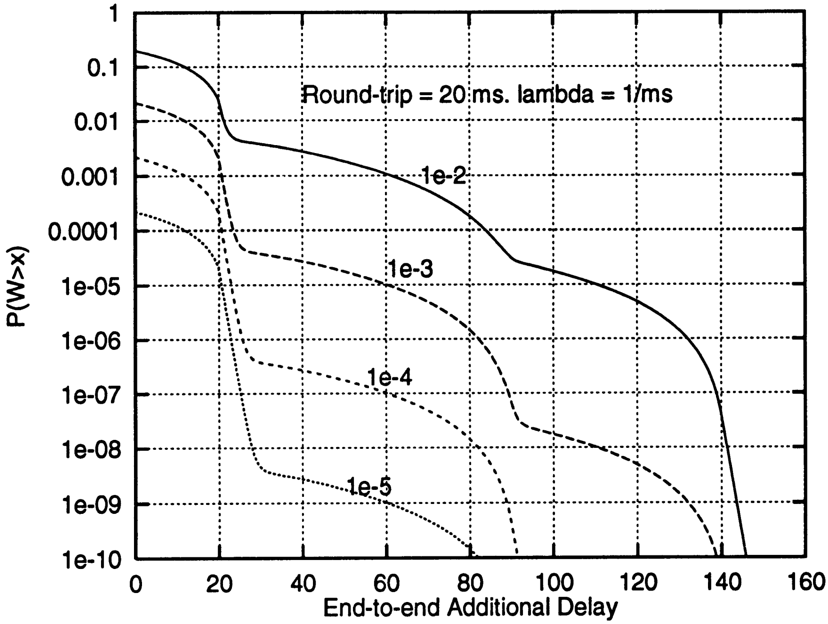

Figure 4: Delay distribution, for various Loss Probabilities

The first curves give the distribution of the end-to-end additional delay: probability that the delay is larger than t, for various packet loss ratios. As expected, the higher the probability, the longer the additional delay.
Figure 4 is drawn for a round-trip delay of 20 ms, and for a POLL time of 50 ms. The arrival rate is 1/ms.
Each "knee" of the curves corresponds to losses. When a data packet is lost, an "error period" begins, which ends when the packet is successfully retransmitted. All packets sent during the error period are delayed (Remind a PDU is delayed if it is errored and has to be resent, <u>or</u> if a preceding PDU has been lost and must be sent again.
PDUs involved in a single-error period have a delay responsible of the first knee, while PDUs sent in a 2-errors period give rise to the second knee, etc. The "knees" appear such due to the log scale on the vertical axis. Were the scale be linear, each of the domains would approximately look like a piece of straight line.
In a first approximation, the limits of the domains are as follows: the k-th error (i.e., the domain corresponding to waiting times due to frames incurring exactly k consecutive transmission errors) produces waiting time probability distributions ranging from $\lambda T^{\mathrm{RT}} p^k$ to $\lambda T^{\mathrm{RT}} p^{k+1}$. Recall that λ is the frame arrival rate in the Sender, T^{RT} is the round trip delay (approximately 2 propagation delays) and p the frame loss probability.

4.2 Influence of POLL periodicity

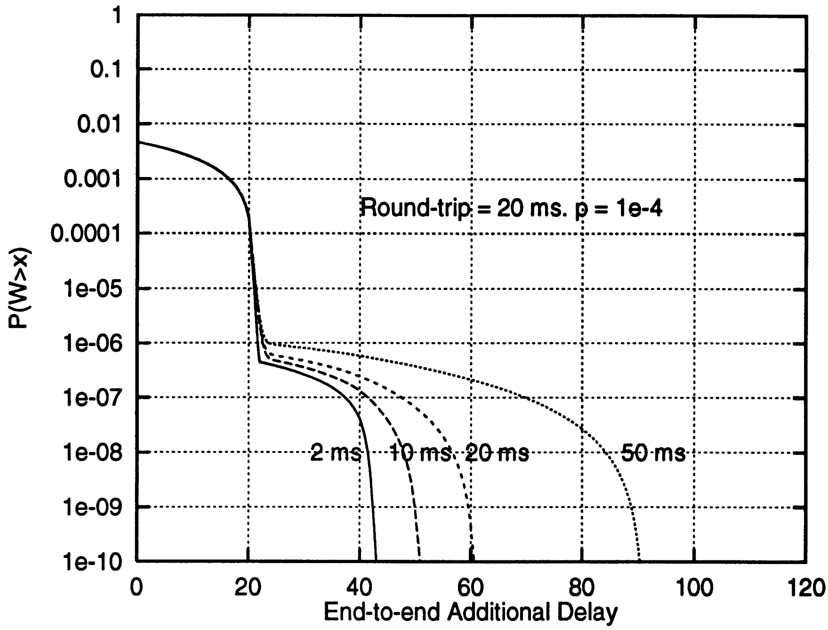

Figure 5: Delay distribution, for various POLL timers

The POLL period has no effect on PDUs involved in a single error: their additional delay is around T^{RT} (see Section 3.2). On the other hand, packets incurring two retransmissions (or more) have an additional delay which can be as large as $2T^{\mathrm{RT}} + T^{\mathrm{poll}}$, as shown on Figure 3. This can be verified on Figure 5.
The abscisses of the knees are around T^{RT} for the first one, and $2T^{\mathrm{RT}} + T^{\mathrm{poll}}$ for the following one (at least, as long as λ is large enough, see the following Figure).

4.3 Influence of the connection rate

Figure 6 illustrates the effect of the PDU rate. For low data rates, the additional delay increases. For instance, the delay at which the first knee ends is around 20 ms for $\lambda = 1$ or 3 /ms, while it goes to around 30 for 0.5/ms, and to 70 ms for lower rates.
The point is in the delay to detect the PDU error. Usually, the next packet detects the gap in numbering sequence. However, if the input rate is too low, the next POLL is likely to be sent first, and the detection delay increases up to T^{poll}, so that the first knee ends at $T^{\mathrm{RT}} + T^{\mathrm{poll}}$. This is what happens here with $\lambda = 0.02$/ms. At the same time, the level of each knee decreases, since each packet is much less likely to be involved in an error recovery period (since less packets are sent). In the limit, the additional delay a packet

incurs is only due to its own losses, so that the probability levels at which the successive steps begin are around p^k instead of $\lambda p^k T^{\mathrm{RT}}$ as for higher rates.
Note that with 1 kbit-packets, a link rate of 0.02/ms corresponds to a 20 kbit/s connection.

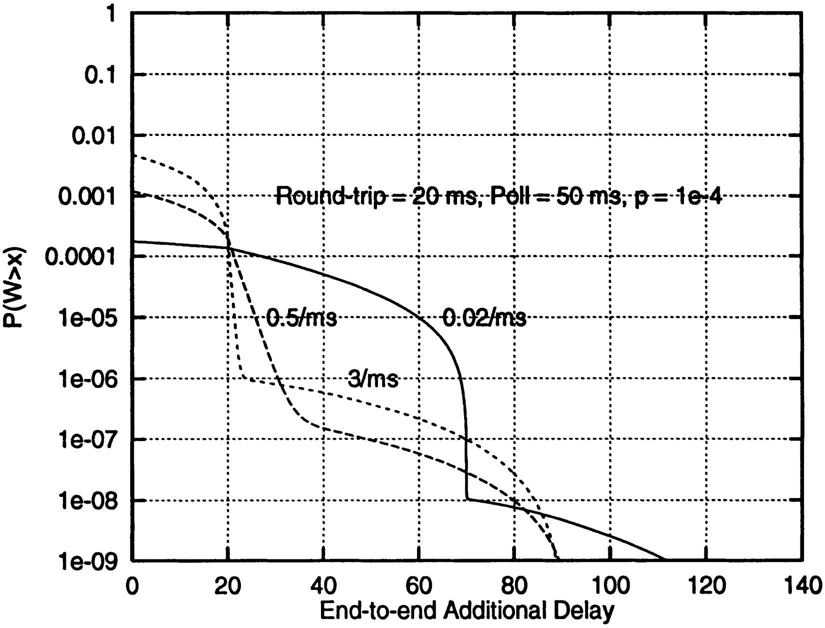

Figure 6: Delay distribution, for various packet arrival rates

5 Conclusions

We have presented an analytical derivation of the end-to-end delay to be expected with the Selective Repeat scheme used in the SSCOP. The analysis is summarized Equation (3) which can be easily programmed and which allows a study of the protocol for realistic figures.
The results show the influence of the main traffic and control parameters: the round-trip delay, the PDU error ratio, the packet load, the poll period. Figure 7 summarizes the typical behavior of the probability distribution function of the End-to-end additional delay: successive "steps" correspond to packets involved in 1, 2, etc. retransmissions. The actual distribution lays below the envelope. The results show that for typical loss levels, the additional delay, i.e. the part of the delay which is related with error recovery, is less than $\lambda p T^{\mathrm{RT}}$. For most cases, this figure is less than 1%. This may question the utility of the dimensioning criterion presently proposed by Rec. E733. Concerning the poll period, its value is of little importance as long as loss figures remain moderate <u>and</u> the load is not too low. For data links with low utilisation, T^{poll} would have to be shortened.

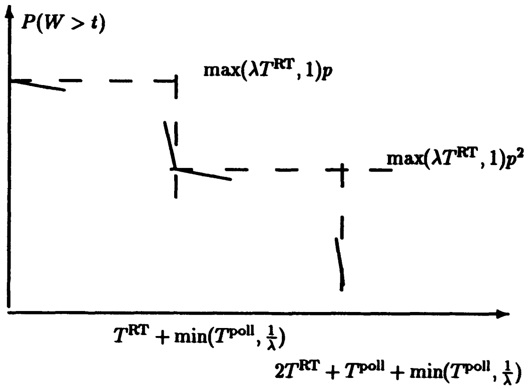

Figure 7: Schematic representation of the shape of Additional Delay

References

[1] D. Bertsekas, R. Gallager: *Data Networks.* Prentice Hall, second edition 1992.

[2] G. Hébuterne: *Service Specific Connection Oriented Protocol: delay calculation (no credit limitation.* INT Technical Report, october 1995.

[3] G. Hébuterne, W. Monin: *Preliminary performance of the SSCOP.* 14 th International Teletraffic Congress, Antibes Juan les Pins, June 1994.

[4] ITU-T Recommendation I.121: *Broadband aspects of ISDN.*

[5] ITU-T Draft Recommendation Q.2110: *Service Specific Connection Oriented Protocol (SSCOP).*