Service Specific Connection Oriented Protocol for the ATM Adaptation Layer

Naotaka MORITA, Ichiro INOUE, Michiharu MITO, and Takeshi AKAIKE
NTT Communication Switching Laboratories

Abstract

Service Specific Connection Oriented Protocol is now being studied for use as the ATM Adaptation Layer Protocol for connection oriented data communication and UNI/NNI signalling in the B-ISDN. The key feature of this protocol is a numbered poll mechanism that enables status information to be exchanged frequently and that is helpful for selective retransmission for frame loss. This paper examines this mechanism from the viewpoints of polling interval and frame length. Because this protocol was designed to perform selective retransmission for error recovery, some requirements - such as saving buffer - are not met. Two enhancements are therefore proposed: one consists of a Go-back-N-like procedure and nonassured credit, and the other is an SD PDU with a P bit. Guidelines to determine the POLL PDU interval are also given.

1. Introduction

Assured mode high-speed data communication by means of the B-ISDN will be possible within a few years. An appropriate protocol suitable for high-speed ATM connection as well as for signalling will be required. The ATM Adaptation Layer (AAL) [1],[2] has been discussed with these applications in mind and ITU-TS (the former CCITT) is trying to adopt a new protocol called the Service Specific Connection Oriented Protocol (SSCOP) [3].

This paper, which examines the advantage of frequently exchanging status information in SSCOP [4],[5],[6] and which proposes some enhancements of key functions, is organized as follows. Section 2 briefly introduces SSCOP procedures from such viewpoints as error recovery control, flow control, and keep alive control, and Section 3 uses simulation results to evaluate the numbered poll mechanism. Section 4 proposes an enhancement of SSCOP to simplify implementation and reduce buffer requirements, and Section 5 concludes by discussing the POLL PDU interval and presenting guidelines for selecting this interval.

2. SSCOP procedures

This section describes the basic architecture and procedures providing SSCOP functions. SSCOP procedures are described for three important functions: error recovery control, flow control, and keep alive control.

2.1. Definitions

•PDU types

SD (Sequenced Data) PDU: carries user information.

POLL PDU: issued by a transmitter to request receiver's state information, data acknowledgement, and new credit.

S-STAT (Solicited STATus) PDU: issued by a receiver in response to a POLL PDU and informs the transmitter about the receiver's state, data acknowledgement, and new credit.

U-STAT (Unsolicited STATus) PDU: issued by a receiver autonomously; and informs the transmitter about the receiver's condition, data acknowledgement, and new credit.

•Variables

-At the transmitter

V(S): The sequence number of the next SD PDU to be transmitted.

V(A): The sequence number of the next in-sequence SD PDU expected to be acknowledged.

V(P): The current value of the poll sequence number, which is incremented before the next POLL PDU is transmitted. This variable is associated with each transmitted or retransmitted SD PDU.

-At the receiver

V(R): The sequence number of the next in-sequence SD PDU expected to be received.

V(H): The sequence number of the next highest SD PDU to be received.

-In the frame

N(S): Sequence number in the SD PDU.

0-8186-3670-X/93 $03.00 © 1993 IEEE

N(R): SD PDU's number acknowledged by the receiver.
N(P): Poll sequence number used in POLL and S-STAT PDUs. N(P) in a STAT-PDU is copied from N(P) in a POLL PDU.

2.2. Error recovery control

Since the SSCOP is expected to be an end-to-end protocol over a high-speed network, selective frame retransmission is required to prevent unnecessary frame retransmission while keeping the required throughput. This selective retransmission can be achieved by using two kinds of retransmission requests: an initial retransmission request and a full retransmission request.

•Initial Retransmission Request

When the receiver detects frame loss by finding that the N(S) in a received SD PDU is higher than the V(H) stored in the receiver, the receiver requests selective frame retransmission by U-STAT PDU. More than one lost frame can be requested, but the frames requested must be successive. If the receiver detects another frame loss, another U-STAT PDU can be issued without any restriction from the protocol point of view. This simple retransmission request provides a quick selective retransmission mechanism but is not sufficient for all cases.

•Full Retransmission Request

Full Retransmission Request is needed when an Initial Retransmission Request (i.e., U-STAT PDU) is lost or when the frames being retransmitted are lost again. The transmitter sends a POLL PDU to the receiver in order to solicit full status information, and the receiver informs the transmitter about the results of frame reception from V(R) to V(H) by using a S-STAT PDU. After receiving the S-STAT PDU, the transmitter checks whether indicated frames have been retransmitted after the POLL PDU soliciting the S-STAT PDU was issued. The transmitter retransmits the indicated frames only if it has not retransmitted them. The S-STAT PDU can even indicate unreceived frames that are not successive.

•Numbered status enquiry and notification

POLL and S-STAT PDUs are associated by a poll sequence number that makes it possible to check pointing frequently, since the transmitter easily finds out the timing relationship

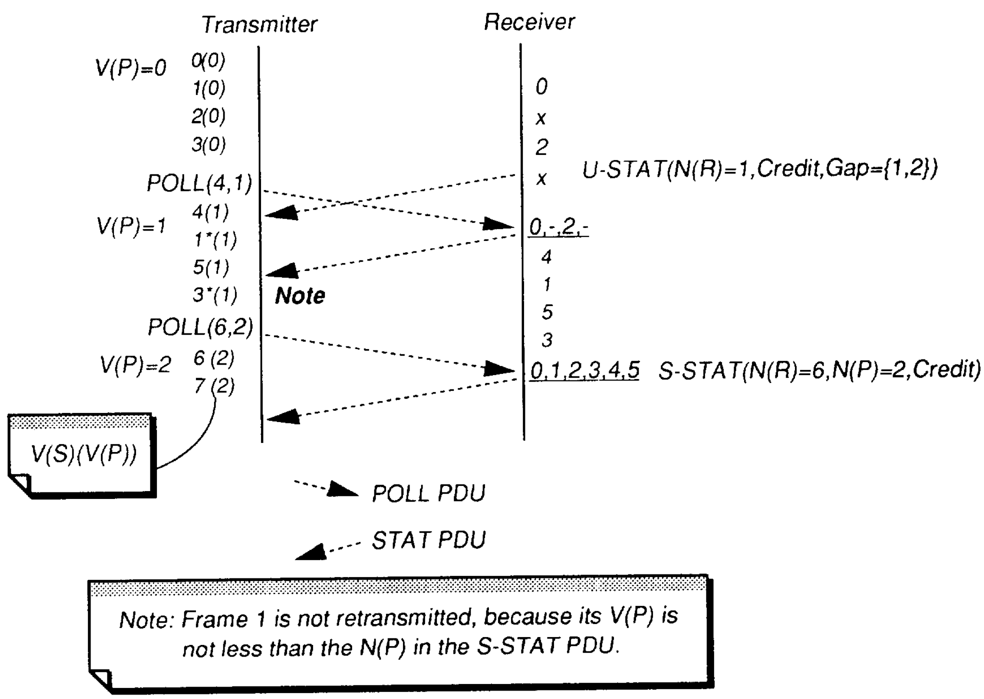

Figure 1. Collision resolution between U-STAT and S-STAT PDUs

between transmitted frames and their reception status notification. Figures 1 and 2 shows how to resolve collisions between U-STAT and S-STAT or between multiple S-STAT PDUs. This technique is quite original in the SSCOP.

The poll number V(P) is incremented before the next POLL PDU is transmitted, and the SD PDU is stored in the retransmitting buffer with the current V(P). When an S-STAT PDU is received, the N(P) in the S-STAT PDU is compared with the V(P) stored with the requested SD PDU. If N(P) is greater than the V(P), it is interpreted that the POLL PDU having solicited the S-STAT PDU was issued after the requested SD PDU had been transmitted, and the indicated SD PDU should be retransmitted. Retransmission requested by a U-STAT PDU is always accepted, since a U-STAT PDU requests only the first retransmission.

In this manner described above, the transmitter can perform error recovery without unnecessary retransmission while S-STAT PDUs are frequent. The transmitter therefore does not need to care about the round-trip delay when determining the status enquiry (i.e., S-STAT PDU) interval. In HDLC, neither RR,P (corresponding to the POLL PDU) nor SREJ,F (corresponding to the S-STAT PDU) has a number associating them with each other. It is therefore necessary to wait for at least one round-trip delay in order to avoid confusing SREJ causing unnecessary retransmission.

•POLL PDU interval

The transmitter can control the frequency of S-STAT PDUs by setting the POLL PDU interval. If this interval is longer than the round-trip delay (RTD), the error recovery performance is quite similar to that when using unnumbered status exchanges, such as with HDLC. Differences between the performances of these two protocols can therefore easily be demonstrated by using the SSCOP and appropriately adjusting this interval. The error recovery performance related to the POLL PDU interval is shown this way in Section 6.1.

2.3. Flow control

Flow control is discussed here from two viewpoints: for a

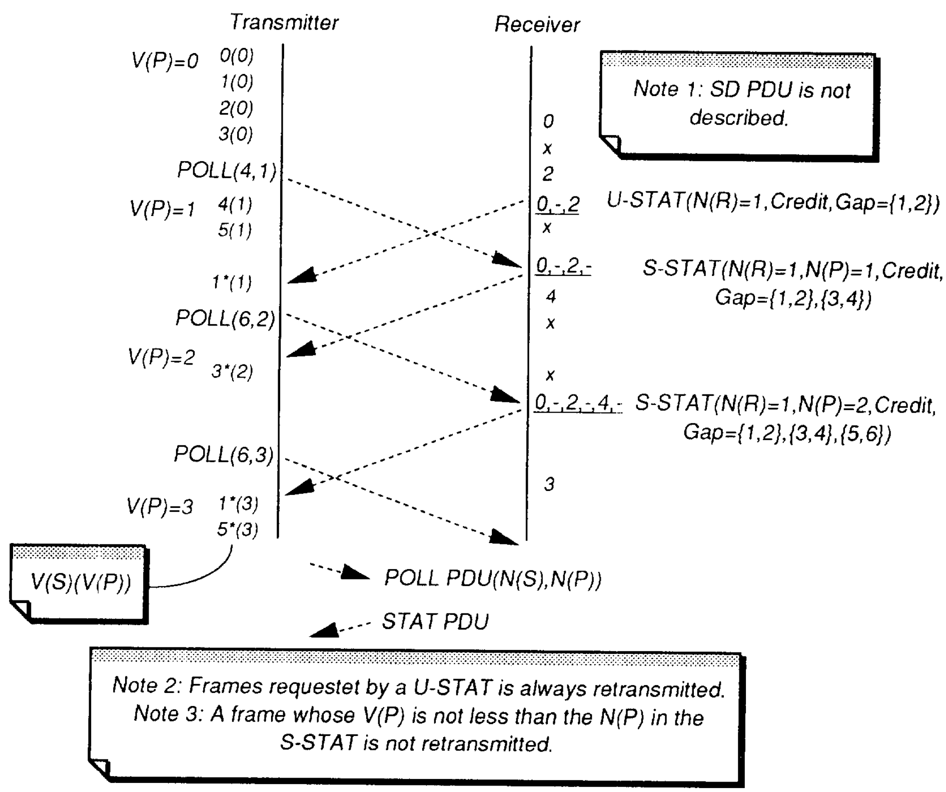

Figure 2. Collision resolution between multiple S-STAT PDUs

receiving buffer and for a retransmitting buffer. The former control for a receiving buffer prevents frame loss due to the overflow at the receiving buffer, and the latter needs an acknowledgement to release outstanding SD PDUs from the retransmitting buffer. In SSCOP, both of them uses POLL and S-STAT PDUs.

•Flow control for receiving buffer

A credit mechanism is used for this purpose. Credit values can be changed by the receiver according to the reception conditions, such as actual buffer capacity and processor utilization. The unit of credit is the maximum length of an SD PDU.

•Flow control for retransmitting buffer

Data acknowledgement is required to release outstanding SD PDUs from a retransmitting buffer. In SSCOP, the N(R) in either an S-STAT or U-STAT PDU represents the data acknowledgement. The interval between the events of updating N(R) is normally controlled by the POLL PDU interval.

2.4. Keep alive control

POLL and S-STAT PDUs can easily provide keep alive control. When there is no user information to be transmitted, a POLL PDU is issued to determine whether the receiver is normal. In general, the POLL PDU interval for keep alive would be longer than that for error recovery or flow control.

3. Evaluation of SSCOP

The numbered poll mechanism provides frequent status enquiry and notification, but its effect depends on appropriate selection of the parameters involved: POLL PDU interval, length of SD PDUs, round-trip delay, and transmission speed. The effects of the numbered poll mechanism relating to the first two parameters are demonstrated here.

3.1. Simulation model

The simulation model is shown in Figure 3 and its parameters are listed in Table 1.

3.2. Response time

Response time is defined as the time between a frame being transmitted and the corresponding acknowledgement being received. This value is useful for appropriately sizing a retransmitting buffer. Since SSCOP adopts selective retransmission, the S-STAT PDU can selectively acknowledge frames successfully received by the receiver.

Figure 4 shows the effect of SD PDU loss in estimating response time. For simplicity, it is assumed that at most one frame is lost during Tpoll. The situation without loss is as shown in Figure 4 a): acknowledgement for all SD PDUs during Tpoll is performed by an S-STAT PDU, and the average response time RT_a is

$$RT\_a = 0.5 \times Tpoll + RTD.$$

When there is a single loss, the SD PDUs received before the loss are acknowledged by a U-STAT PDU and the SD PDUs received after the loss are acknowledged by the S-STAT PDU issued in response to the next POLL PDU. The lost SD PDU is acknowledged by the S-STAT PDU after the reception of the retransmission. Except for the lost SD PDU, the average response time, which is independent of the relationship between Tpoll and RTD, is

$$RT\_o = 0.25 \times Tpoll + RTD$$

The response time for the lost SD PDU depends on the relationship between Tpoll and RTD.

Table 1. Simulation Parameters

Parameters	Abbreviation	Variation
Length of SD PDU	L	4096, 64K [B]
Offered traffic	λ	0.5 V [bps]
Distribution of SD PDU generation		Poisson
Link speed	V	156 [Mbps]
Round-trip delay	RTD	20 [msec]
Cell Loss Rate	CLR	$10^{-3}, \ldots, 10^{-6}$
POLL PDU interval	Tpoll	10, 50 [msec]
ATM cell size	Z	53 [B]
User information capacity per cell	C	48 [B]

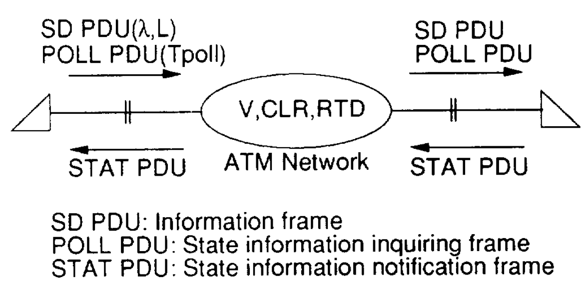

Figure 3. Simulation Model

When Tpoll < RTD as is shown in Figure 4 b), the response time for SD PDU is

$$RT\_b\_sd = 0.5 \times Tpoll + 2 \times RTD$$

RTD is assumed as an integral multiple of Tpoll to simplify the estimation. The average response time for all SD PDUs is

$$RT\_b = (1-1/F) \times RT\_o + (1/F) \times RT\_b\_sd$$

where $F = \lambda Tpoll/8L$.

When Tpoll > RTD, as is shown in Figure 4 c), two cases should be considered: one is that the retransmission of the lost SD PDU is performed before the end of the Tpoll, and the other is that the retransmission is performed after the end of Tpoll. For the condition shown in Figure 4 c-1), the response time for the lost SD PDU is

$$RT\_c1\_sd = 0.5 \times Tpoll + 1.5 \times RTD$$

For the condition shown in Figure 4 c-2), the response time for the lost SD PDU is

$$RT\_c2\_sd = Tpoll + 1.5 \times RTD$$

The average response time for the lost SD PDU is

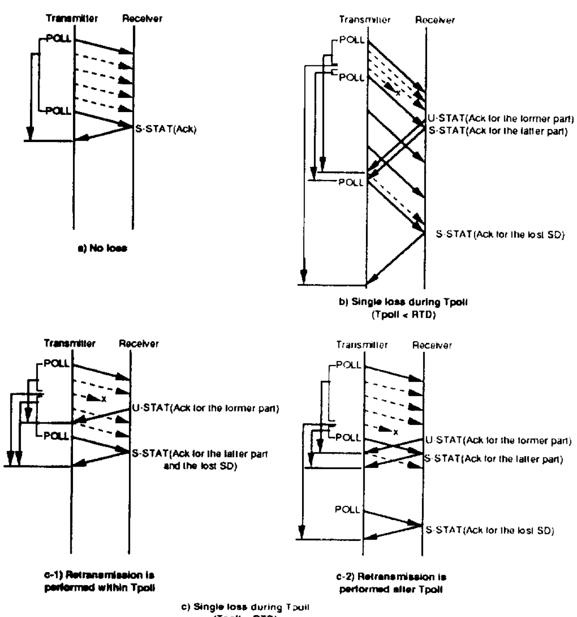

Figure 4. Time chart for estimating the average response time

$$RT\_c\_sd = \{(Tpoll - RTD) \times RT\_c1\_sd + RTD \times RT\_c2\_sd\}/Tpoll = 0.5 \times Tpoll + 2 \times RTD$$

As the result, RT_b_sd = RT_c_sd, and the response time for the lost SD PDU as well as for the other SD PDUs is independent of the relationship between Tpoll and RTD. Therefore, the average response time in case c) is

$$RT\_c = RT\_b = (1-1/F) \times RT\_o + (1/F) \times RT\_b\_sd$$

If we assumed that Pok is the probability of no loss during Tpoll, we can derive the average response time RT

$$RT = Pok \times RT\_a + (1 - Pok) \times RT\_b,$$

where $Pok = (1-CLR)^{\lambda \times Tpoll/8C}$.

Assuming that CLR is small enough for us to ignore the second part of the formula,

$$RT = RT\_a = 0.5 \times Tpoll + RTD.$$

This means that response time can be reduced according to the frequency of POLL PDU. If CLR is large,

$$RT = RT\_b = (1+1/F) \times (0.25 \times Tpoll + RTD)$$

Figure 5 shows the simulation results of this consideration. When one loss occurs, the preceding frames are acknowledged by a U-STAT PDU and the following frames are acknowledged by the next S-STAT PDU. This implies that issuing U-STAT PDUs more frequently - whenever frame loss is detected - will result in earlier acknowledgement. This is shown in the case that L is 4096 B, Tpoll is 50 ms, and CLR is 10^{-4}.

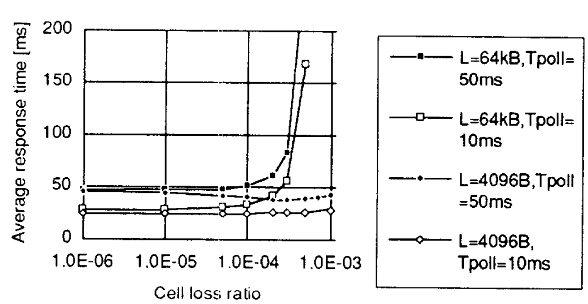

Figure 5. Simulated average response time RTD=20ms

3.3. End-to-end transmission delay

End-to-end transmission delay is defined as the interval between a frame being transmitted from the upper layer to the AAL at the transmitter and the frame being delivered to the upper layer at the receiver. Figure 6 shows the time chart when frame loss occurs. Once a frame is lost, successive frames are stored at the receiving buffer to wait for the retransmission of the lost frames and then reordered for delivery to the upper layer in sequence.

Figure 6 also shows the probability of multiple losses of a frame, the number of frames to be waited for, and average increase in transmission delay. According to the results in the table, average end-to-end transmission delay TD is

$$TD = TDok + Perr1 \times Nerr1 \times \Delta\, TDerr1 + Perr2 \times Nerr2 \times \Delta\, TDerr2.$$

If the CLR is zero or small,

$$TDok = 0.5 \times RTD,$$

or

$$TD = TDok + Perr1 \times Nerr1 \times \Delta\, TDerr1 = 0.5 \times RTD + \{1-(1-CLR)^{L/C}\} \times \lambda\, RTD/8L \times 0.5 \times RTD.$$

As long as the error probability is small, TD is independent of Tpoll. The advantage of the numbered poll appears when a frame is lost more than once.

Figure 7 shows the simulation results. The advantage of the numbered poll exchange appears when the CLR is greater than 10^{-3} in case of L is 4096 B, and when the CLR is greater than 10^{-4} in case of L is 64 kB.

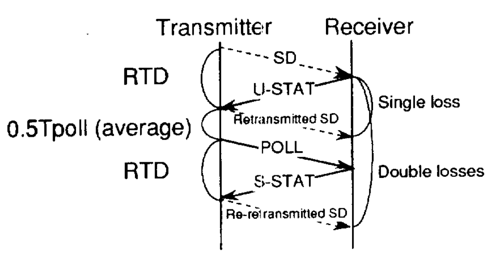

Number of losses	Probability	Number of frames involved	Δ Delay
0(No error)	$Pokf=(1-CLR)^{L/C}$	1	TDok=0.5RTD
1(Single loss)	$Perr1=(1-Pokf)$ $=1-(1-CLR)^{L/C}$	$Nerr1=\lambda RTD/8L$	ΔTDerr1=0.5RTD
2(Double losses)	$Perr2=Perr1^2$ $=\{1-(1-CLR)^{L/C}\}^2$	$Nerr2=\lambda(0.5Tpoll+2RTD)/8L$	ΔTDerr2=0.25Tpoll+RTD

Figure 6. Time chart for estimation of end-to-end transmission delay

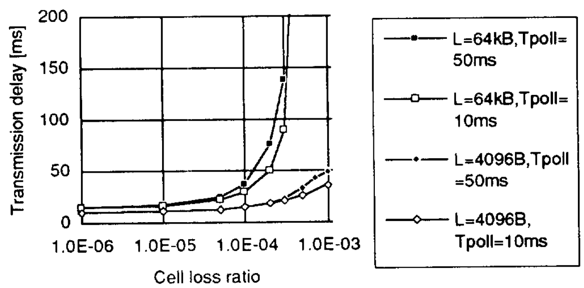

Figure 7. Simulated end-to-end transmission delay
RTD=20ms

3.4. Occupancy of receiving buffer

The sufficient receiving buffer is necessary for making good use of selective retransmission. As illustrated in the previous section, the buffers required when one error occurs, Buf1, and additional buffers required when the same frame suffers another error, ΔBuf2, are

$$Buf1 = RTD \times \lambda$$

and

$$\Delta Buf2 = (RTD + 0.5 \times Tpoll) \times \lambda.$$

The probability of each being needed, Perr1 and Perr2, are

$$Perr1 = Perr$$

and

$$Perr2 = Perr^2,$$

where $Perr = 1 - (1 - CLR)^{L/C}$

Thus the average receiving buffer required when a frame is lost no more than two times is

$$Buf = Perr1 \times Buf1 + Perr2 \times \Delta Buf2 = Perr \times RTD \times \lambda + Perr^2 \times (RTD + 0.5 \times Tpoll) \times \lambda$$

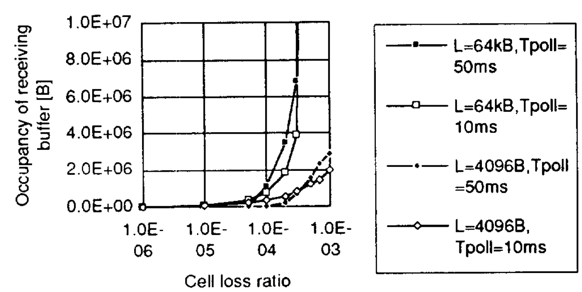

Figure 8. Simulated occupancy of receiving buffer
(Fixed frame length equal to L)
RTD=20ms

The occupancy of the receiving buffer is also independent of Tpoll when CLR is small.

Figure 8 shows the simulation results. The advantage of the numbered poll exchange appears when the CLR is greater than 10^{-3} in case of L is 4096 B, and when the CLR is greater than 10^{-4} in case of L is 64 kB.

4. Proposal of Go-back-N retransmission and nonassured credit

4.1. Unit of credit

An important factor in a credit mechanism is the unit of credit: that is whether credit is handled on a frame basis or an information basis. Credit precisely reflecting occupancy of the receiving buffer uses the buffer efficiently, but the additional complexity of frame processing may decrease throughput. In practice, the receiving buffer is managed according to a certain amount of information; e.g., 512 or 1024 octets. This means that a frame shorter than this buffer unit occupies one buffer and that so many frames cannot be received - even if each frame is shorter than the buffer unit and the total amount of information they contain is less than the nominal buffer capacity. When the frame length is not fixed, the most efficient way to use a receiving buffer is therefore to negotiate the buffer unit as well as the maximum frame size. In this way, however, credit calculation gets more complex. Furthermore, issues related to buffer usage must be considered from two aspects: normal procedure and error recovery procedure. The procedure proposed here increases buffer utilization by stressing the normal procedure even when credit is handled on a frame basis.

4.2. Advantage of Go-back-N

Selective retransmission is effective when the product of round-trip delay and transmission speed represented here by K, is large, the error pattern is not bursty, and buffer resources are sufficient. It requires additional processing, however, for reordering the retransmitted frames, and it requires a large buffer for storing frames in order to keep the requested throughput. The Go-back-N mechanism, on the other hand, is preferable when K is small, the distribution of frame loss is bursty, and buffer resources are not sufficient. This section proposes an enhancement of SSCOP that provides an effect similar to that provided by Go-back-N. This proposal is related to the policy of giving credit described below.

4.3. Nonassured credit

In selective retransmission, the receiving buffer is used mainly for error recovery. Most of the buffer is usually not used. Furthermore, actual traffic may not be constant. This means that even when an error occurs, the full capacity of the receiving buffer may not be exhausted. When there is no error, on the other hand, a certain amount of credit is required for the transmitter to transmit frames meeting the throughput requirement and round-trip delay condition. In the current SSCOP specification, a credit value guarantees the reception of frames and this imposes a large buffer requirement on the receiver. To relax this restriction, we propose nonassured credit by which the receiver can give credit to the transmitter beyond the buffer capacity required during error-free operation. This idea is more effective for multiple incoming connection using shared receiving buffers. This mechanism requires SSCOP enhancement to provide a Go-back-N-like procedure as is described below.

4.4. Proposed procedure

The proposed procedure modifies only the receiver's action. Once the receiver detects frame loss by receiving an SD PDU out of sequence and there is no more room to receive all frames corresponding to the detected sequence gap, it uses a U-STAT PDU to request retransmission going back to the V(R) frame. Successive SD PDUs are discarded until an SD PDU with N(S) equal to V(R) is received. If the receiver has a buffer too small to store additional frames out of sequence, this mechanism is similar to the Go-back-N mechanism. One issue of this proposal is how to set the upper bound of requested frames and credit in the U-STAT PDU. We propose that the actual buffer occupancy be indicated when determining the upper bound and new credit. If the actual space in the receiving buffer is smaller than the space corresponding to the previous credit, the credit is set less than the previously given credit. The throughput restriction due to the receiving buffer appears only if the error occurs. Frames received after issuing the U-STAT PDU and discarded are requested by the future S-STAT PDU. This enhancement lets the receiving buffer be used efficiently even if the unit of credit is frame basis.

5. The POLL PDU interval

This section discusses the POLL PDU interval and presents guidelines for determining this parameter. An additional mechanism providing quicker acknowledgement is also proposed.

5.1. Error recovery

As shown in the previous discussion, the numbered poll mechanism has an advantage that enables status information to be exchange frequently. The actual effect, however, only appears when the initial retransmission procedure fails (i.e., U-STAT PDU is lost or retransmitted SD PDU is lost again). This is because almost all frame losses are recovered by using a U-STAT PDU. For error recovery, therefore, the POLL PDU interval need not be shorter than the round-trip delay.

One algorithm for issuing a POLL PDU is triggered by the end of an information burst. Issuing a POLL PDU just after retransmission may be effective for bursty frame loss.

5.2. Flow control

At the receiver side, the size of the buffer is determined to utilize selective retransmission when assured credit is assumed. In this case, the POLL PDU interval is determined to maintain required throughput even if error occurs. If a user does not care about the throughput decrease caused by frame loss, the size of the receiving buffer can be reduced by the Go-back-N procedure and nonassured credit as proposed in the previous section.

At the transmitter side, on the other hand, the minimum buffer space cannot be reduced to less than the amount of outstanding information. The minimum size of retransmission buffer can be achieved by the most frequent acknowledgement, i.e., acknowledgement to every frame. In this case, a poll sequence number is not necessary, since what is required is only a new N(R) and new credit rather than the list of lost frames. But the current SSCOP uses the same pair of PDUs, POLL and STAT PDUs, to perform both error recovery and flow controls. In this mechanism, the frequency of information exchange for both controls is the same. Independent control, such as frequent flow control and infrequent error recovery, cannot be performed. Therefore, a simple solicitation of credit is proposed, such as the addition of a P bit to the SD PDU. When this P bit is set, the receiver uses a U-STAT PDU to inform the transmitter about its V(R) and credit. Since the transmitter wants V(A) updated rather than the reception information about each SD PDU, P bit and U-STAT PDU is sufficient rather than Poll sequence number and S-STAT PDU. This enhancement can save space in the retransmitting buffer without increase of PDUs to be processed at the receiver.

5.3. Keep alive

Keep alive control monitors the availability of the receiver by the transmitter even if there is no user information to be transmitted. This is done in order to resume the transmission without any connection confirmation procedure (including connection establishment). This function imposes additional information exchanges and therefore undesirable in certain environments, such as one with usage-sensitive charging. The use of POLL PDU for keep alive must therefore be optional.

5.4. Guidelines for setting the POLL PDU interval

- The POLL PDU interval for error recovery control need not be shorter than RTD.
- The POLL PDU interval for acknowledgment is determined according to the size of the retransmitting buffer. Instead of frequent polling, an SD PDU with a P bit is sufficient.
- The POLL PDU interval for keep alive should be manageable for various charging environments.

6. Conclusion

The key feature of the SSCOP is a numbered poll mechanism whose advantages - in terms of average response time, end-to-end transmission delay, and occupancy of receiving buffer - appear when the cell loss rate is greater than 10^{-4}, round-trip delay is 20 ms, requested throughput is about 80 Mbps, and frame length is either 4 or 64 kB. To reduce buffer requirements and simplify implementation, an enhancement for performing a Go-back-N-like procedure is proposed. Another proposed enhancement, an SD PDU with a P bit, provides quicker acknowledgement. Guidelines for determining the most important parameter of SSCOP, the POLL PDU interval, are also presented.

References

[1]CCITT recommendation I.362 "B-ISDN ATM Adaptation Layer Functional Description."

[2]CCITT recommendation I.363 "B-ISDN ATM Adaptation Layer Specification."

[3]CCITT draft recommendation Q.SAAL.0, Q.SAAL.1, Q.SAAL.2, Q.SAAL.3, February 1993.

[4]M. Mito, T. Akaike, "A performance analysis of AAL Service Specific Connection Oriented Protocol," SIGCOMM '93 to be submitted.

[5]A. N. Netravali, W. R. Roome, and K. Sabnani, "Design and Implementation of A High Speed Transport Protocol," IEEE Transactions on Communications, Nov. 1990.

[6]W. Bux, P. Kermani, W. Kleinoeder, "Performance of an Improved Data Link Control Protocol," Proceedings of 9th ICCC, Tel Aviv, Nov. 1988.