A SIMULATION STUDY OF HDLC-ABM WITH SELECTIVE AND NONSELECTIVE REJECT

Goeldner, E.-H., Truong, H.L.+)

Institute of Switching and Data Technics, University of Stuttgart
Fed. Rep. of Germany

ABSTRACT

The objective of this paper is to compare the performance of the two recommended recovery options of the data link control procedure HDLC, balanced mode of operation (ABM). The reject (REJ) -recovery solicits the transmitter to go back to an already transmitted frame (which has been disturbed) and to start retransmitting sequentially at the requested I-frame. The selective reject (SREJ) -recovery, on the contrary, solicits the transmitter to retransmit only the disturbed I-frame. To obtain an accurate account of the performance of the REJ- and SREJ-option, the approach taken is a detailed simulation of the data link. The performance of the full duplex channel is measured in terms of

- the maximum throughput to be achieved over a point-to-point link and
- the mean transfer time for messages

both as a function of several essential parameters. Therefore, some conclusions can be drawn from this study in which cases the SREJ-recovery improves the performance of a HDLC-link.

1. INTRODUCTION

Within the last years the High-level Data Link Control procedure (HDLC) has been subject of several studies and a number of performance investigations have been published /1 - 3/. The HDLC documents provide several options to improve the performance of a so controlled data link. For an advanced retransmission of an erroneous I-frame the reject (REJ)- and selective reject (SREJ)- option are proposed. Almost all investigations and implementations are based on the REJ-recovery, but the SREJ-recovery, due to the greater complexity, only emerges rarely. At the Institute of Switching and Data Technics of the University of Stuttgart a detailed simulation-program has been developed to simulate a data link, controlled by HDLC - ABM (Asynchronous Balanced Mode) and using REJ- or SREJ-recovery.
Our descriptions are confined to the topics that are important for our simulation and the interpretation of the results. The reader who is not familiar with the details of HDLC is referred to the HDLC-documents, see /4 - 5/, resp. sooner publications /1 - 3/.

In this section, we briefly describe the operation of a full-duplex data link under HDLC, balanced class of procedures. For an example of operation, the sequences are illustrated in Fig.1 and 2 with the aid of commonly used diagrams, e.g., the line labelled A to B describes frames transmitted from station A to station B. The sequence is marked by the send sequence number N(S) and the receive sequence number N(R), respectively.

1.1 REJ-Recovery

A frame received in error is simply discarded by the receiver without any further action. If the frame is an I-frame, the error will manifest itself later in the form of a sequence-error or it will be detected by means of time-out or checkpointing. Upon detection of a sequence-error, the REJ command/ response is used to initiate the retransmission of the erroneous I-frame and all subsequently transmitted frames. Fig.1 shows an example of REJ-re-

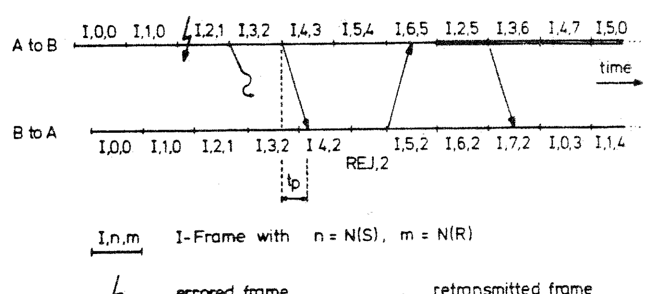

Fig.1: Example of HDLC-balanced operation - transmission error recovered by REJ-function (t_p = propagation and processing delay)

covery: The I-frame with N(S)=2 is received in error and, therefore, discarded by station B. When station B receives the next errorless I-frame, e.g. the frame with N(S)=3, it informs station A of the sequence-error by using a REJ-frame with N(R)=2. Upon receipt of the REJ-frame, station A retransmits the requested I-frame with N(S)=2 plus all additional I-frames which have been subsequently transmitted (I-frames with N(S)=3, 4, 5, 6). Therefore, station B accepts only I-frames received in sequence.

+) H.L. Truong is now with Standard Elektrik Lorenz AG Stuttgart, Fed. Rep. of Germany

1.3 SREJ-Recovery

To avoid retransmission of correct I-frames following a disturbed one, HDLC provides the optional selective-reject (SREJ) function, where only the disturbed I-frame has to be transmitted. The subsequently transmitted I-frames, received without an error, but not in sequence, are not discarded. They queue up at the receiver, and they are sequentially ordered together with the single retransmitted frame. The subsequently transmitted frames cannot be acknowledged before the selectively retransmitted frame has been received correctly.

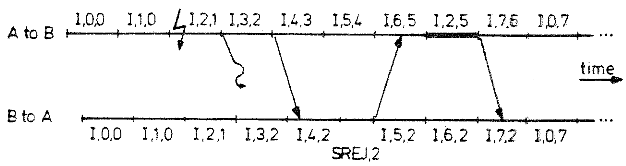

Fig.2: Example of HDLC-balanced operation-transmission error recovered by SREJ-function

Fig.2 shows the same example of information exchange. Now, however, station B uses a SREJ-frame to initiate the retransmission only for the disturbed I-frame with N(S)=2. The I-frames with N(S)= 3, 4, 5 and 6 queue up, and after receiving the correct I-frame with N(S)=2 all frames up to N(S)=6 are acknowledged (with the frame I,0,7 from station B.

As mentioned above, the performance of the SREJ-recovery option has been investigated only a few times. J.Peters demonstrates calculation and simulation of the SREJ-recovery in /6/ under the simplifying assumption that no restrictions are imposed on the modulus numbering (window-mechanism). M.Easton's results in /7/ depend on similar assumptions: no limitation by the modulus-numbering, no errors in the I-frame requested with a SREJ- or REJ-frame and, therefore, no checkpointing, resp. P/F-recovery is necessary.

In order to obtain accurate and reliable performance results, all mechanisms of the procedure (apart from the initialization) have been implemented in full detail in an event-by-event simulation-program. Such details, for example, are modulo-numbering, REJ-, respectively SREJ-recovery or checkpointing with P/F-bit recovery as defined in the HDLC-documents.

In the next chapter we present our queueing model and the assumptions for implementing the procedures. In chapter 3 several simulation results are shown, all comparing REJ-recovery with SREJ-recovery on different links and as a function of several parameters.

2. MODELLING AND SIMULATION

The presupposition for simulating the HDLC-data link is a detailed model. Fig.3 shows the structure of this queueing model. Two identical stations are connected by a full-duplex circuit. This link is controlled by HDLC, balanced class of procedures, including either the optional function REJ or the function SREJ.

Messages to be transmitted from station A to station B or vice versa are stored in the send-queue of the sending station. The message will be lost if every place in the queue is occupied. The messages are transmitted according to first-come, first-served, one message per I-frame. A copy of every message sent will be stored in the repeat-queue for possible retransmission. They are cancelled as soon as the acknowledgement is received. The arrival-process of the messages is defined by the inter-arrival-time and its distribution-function. Similarly, the message length l may be constant or randomly distributed by a given function.

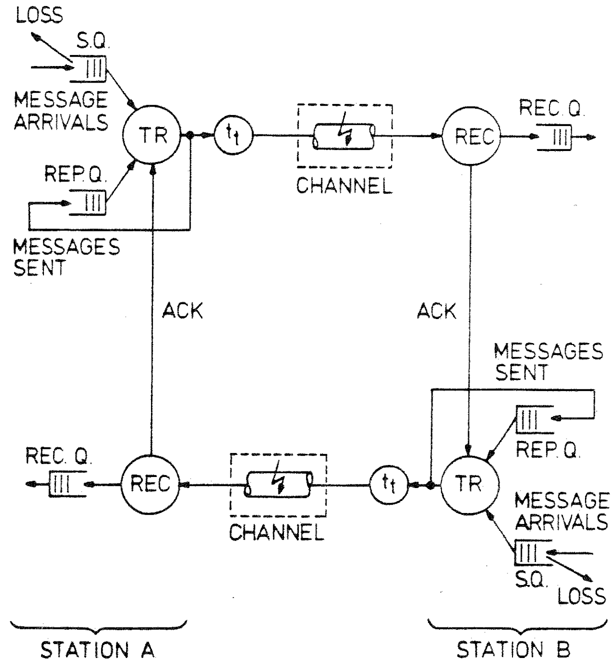

TR	Transmitter
REC	Receiver
S.Q.	Send-queue
REP.Q.	Repeat-queue
REC.Q.	Receive-queue
ACK	Acknowledgements
t_t	Transmission time = l/v (I-frame = t_I; S-frame = t_S
CHANNEL	incl. propagation delay, errors

Fig.3: Model of the FDX-data-link operating under HDLC-ABM

The transmission channels are characterized by their transmission rate v, the one-way propagation delay t_p (including the processing delay of transmitter and receiver) and the probability of errors. We assume that every error in a frame is detected by the frame-checking-sequence, and thus, the frame is discarded by the receiver. Every correctly received I-frame is acknowledged as soon as possible either by an I-frame or (if there are no messages to transmit in this direction) by a receive-ready (RR) frame.

The queueing time of the receive-queue has been assumed zero, and, therefore, the receive-queue does not have any impact on throughput and delay.

The two stations are identical and the link has a symmetrical structure.

Besides the general description of the handling of the P-bit in case of balanced mode of operation (ABM) there is still no recommendation about the definite time schedule it has to be sent. Furthermore, HDLC does not specify how the timer should be handled. For our implementation we adopted the following rules:

1) P- and F-bit are not transmitted with I-frames. Only S-frames are used to set P- or F-bit to one. S-frames envelope 48 or 56 bit and therefore, the frame-error-probability is naturally low.
2) The timer is started every time an S-frame with P-bit set to one is sent (still waiting for the F-bit).
3) The timer is stopped when the F-bit is received.
4) If the timer is stopped or expirated, as soon as possible a S-frame with the P-bit set to one is issued and accordingly checkpointing is initiated.
5) The minimal duration of a time-out has been evaluated by

   $t_{out,min} > 2\ t_p + t_I + t_S$

   t_p = propagation and processing delay
   t_I = maximum transmission time of an I-frame
   t_S = transmission time of a S-frame.

For best studying the influence of REJ- and SREJ-option we simulated the HDLC-ABM controlled link using either only REJ-recovery or only SREJ-recovery. While using the SREJ-option, only selected retransmission (the single, disturbed I-frame) was executed, even as a result of checkpointing.
All the following results have been simulated due to these rules.
The assumption for our simulation model is that the link has already been initialized.
In all simulations presented in this paper we have assumed independent bit-errors only.

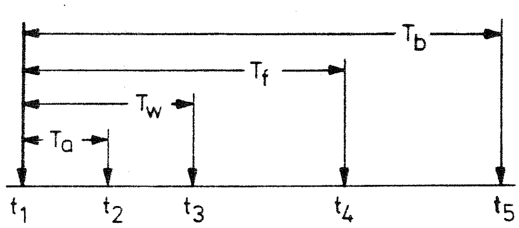

EVENTS
t_1 Message arrived at send-queue
t_2 Arrival of the next message
t_3 I-frame sent for the first time
t_4 I-frame received correctly
t_5 I-frame acknowledged at the transmitter.

MEASURED TIMES
t_a Inter-arrival time
t_w Waiting time for first service
t_f Transfer time
t_b Buffer-holding time

Fig.4: Definition of the measured times

As already pointed out, the parameters describing the stations, the channel and the messages can be chosen arbitrarily. Message length, the time between two message arrivals and the time interval of the errors may be constant or randomly distributed according to the following options: negative exponential Erlangian with order K or hyperexponential distribution function with a given coefficient of variation.

The quantities measured during the simulation are:
- probability of loss (arriving messages)
- throughput of the information
- total channel load
- transfer time
- queue length at send-queue
- buffer-holding time at send-queue

Fig.4 shows the definition of some interesting measured times.

3. RESULTS

This is the main section of the paper in which typical performance results for the two HDLC recovery options, REJ and SREJ will be discussed. All results showing the simulation of REJ-recovery are marked by a triangle and bold lines, the results of SREJ-recovery simulations by a cross and dashed line. In case of no visible distinction between SREJ- and REJ-recovery results, only the triangle has been drawn.
Since the link configuration in ABM and the offered load are symmetrical, we need to consider one direction of the link, only. Here, we evaluated the direction from station A to station B. The results for the other direction are quite similar.
We distinguish between two main categories of results: maximum throughput and mean transfer time.

3.1 Throughput Results

In this case we determine the maximum number of information bits transmitted in one direction of the link, presuming that both stations have information to send any time (saturated traffic/saturated message queues).

3.1.1 Throughput vs Message Length

The diagrams below show the typical results for the maximum throughput of information as a function of message length, i.e. the length of the information field of the I-frames.
The throughput results are in agreement with the general behavior of a link-controlled procedure employing an error-detection and retransmission scheme:

- maximum throughput of information bits in any case has an upper bound given by the ratio of I-field length to total I-frame length. Each I-frame carries a fixed amount of overhead bits for flags, address, control, and FCS-fields (the same number of bits as a S-frame). Hence, the throughput must decrease for very short messages.
- On the other side, the block error probability P_B of the I-frame increases with growing frame length according to:

  $$P_B = 1-(1-(P_{bit})^{l_I + l_S}$$

  with l_I = number of information bits
  l_S = number of supervisory bits (48 if MOD=8, 56 if MOD=128)

Therefore, the fraction of error-free, i.e. useful transmission, decreases.

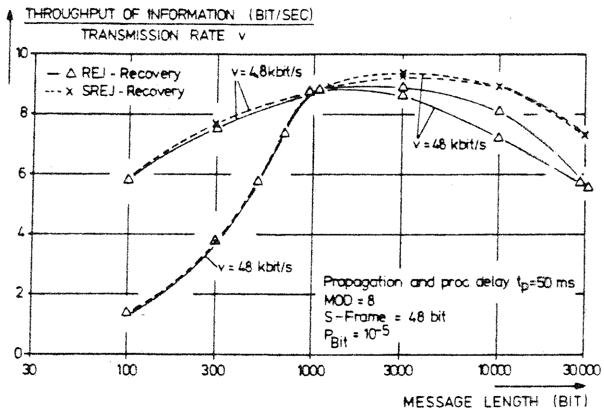

Fig.5: Throughput versus message length transmission rate v = 4.8kbit/s; 48kbit/s

Fig.5 shows the throughput efficiency on a terrestrial link (t_p = 50 ms). The behavior of the curves for v = 4.8 kbit/s is according to the two rules mentioned above; increasing overhead for shorter messages, increasing block error probability for longer messages. An additional characteristic - distinct in particular for v = 48 kbit/s - is the impact of the modulus-numbering, causing a drastic throughput degradation for shorter message lengths. Each station has to stop sending further I-frames, if modulus -1 I-frames are outstanding simultaneously. Up to 1000 bit message lenghts this limitation by modulo-numbering is the dominating factor. There is no difference between REJ-or SREJ-recovery. Using REJ-recovery, the station is allowed to repeat up to modulo -1 unacknowledged frames. On the other hand, using SREJ-recovery only the disturbed frame is retransmitted, but the station has to wait for the acknowledgement of this retransmitted frame. Therefore, while the impact of modulus-numbering is dominant no improvement is possible by using the SREJ-option!
For message lengths greater than 1000 bit the use of SREJ-recovery may be appropriate, especially while using v = 48 kbit/s.

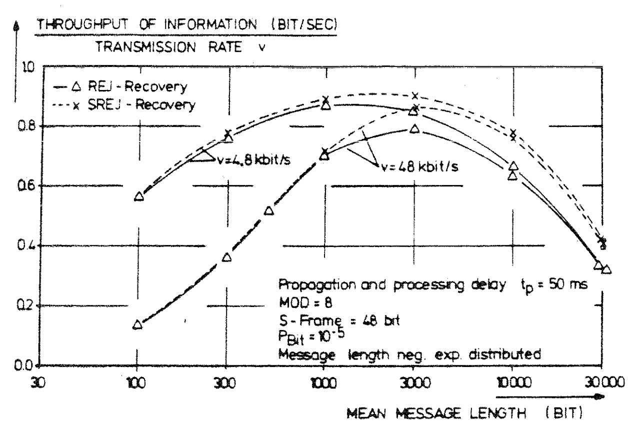

Fig.6: Throughput versus message length transmission rate v = 4.8 kbit/s; 48 kbit/s message length neg.exp.distributed (truncated) max.message length = 10 mean message length

Fig.6 shows simulation results of the same link as before, but with truncated neg.exp. distributed message lengths. The behavior of the curves is quite similar to Fig.5, but now the fluctuating message length does not cause so distinct effects as in Fig.5.

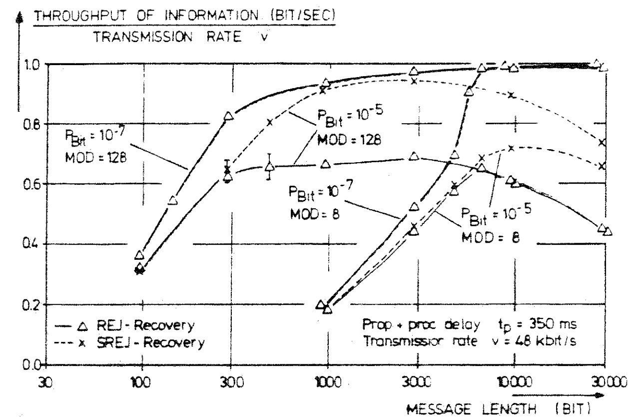

Fig.7: Throughput versus message length P_{bit} = 10^{-5}; 10^{-7} Modulus = 8; 128

The explanation of the throughput behavior of a satellite link in Fig.7 is virtually the same. Two bit-error probabilities are considered, 10^{-7} and 10^{-5}, where P_{bit} = 10^{-5} may be caused by a poor terrestrial extension of the satellite channels. The impact of modulus is, of course more distinct due to the long propagation delay of the satellite channel. Using a modulus value of 8, the impact of the window mechanism ranges up to a message length of 7000 bit. Increasing the modulus value from 8 to 128 yields a substantial improvement of the throughput. Besides this, on a channel with a bit-error probability P_{bit} = 10^{-7} the employment of SREJ-recovery does not lead to a higher throughput. On the contrary, the throughput of a "bad" link with P_{bit} = 10^{-5} increases in a significant way using SREJ-recovery. This effect is obviously to explain by means of Fig.8 and 9.

3.1.2 Throughput versus Propagation Delay

In Fig.8 the throughput is drawn as a function of propagation and processing delay relative to the I-frame transmission time (here: 22ms).
For reasons of clarity the length of the S-frame is constant at 56 bits in Fig.8 and 9, irrespective of the modulus value employed. The graph for a link without errors (P_B = 0.00) shows the impact of the window mechanism: with growing propagation delay the throughput decreases in a very evident manner. The critical delay is about 2 t_p/t_I and for longer delays there is no difference between REJ- and SREJ-recovery, not even at higher error-probability. Only for very short delays and a relatively high frame-error probability the SREJ-option increases the throughput.

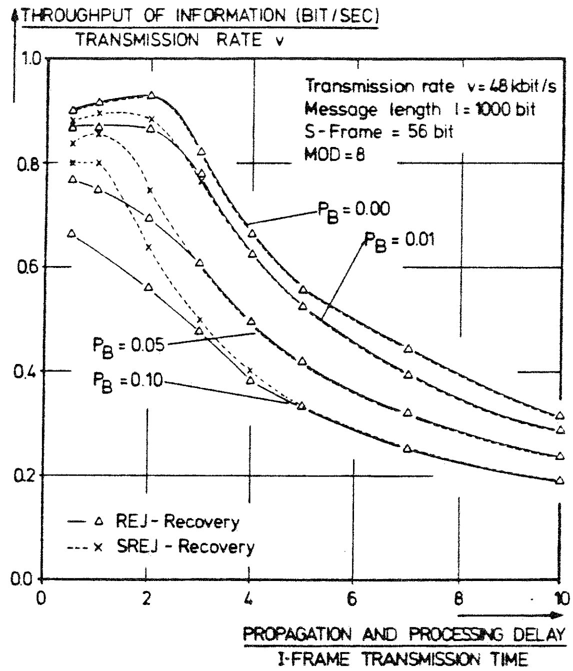

Fig.8: Throughput versus propagation and proc.delay Frame-error probability
P_B *= 0.00; 0.01; 0.05; 0.10*

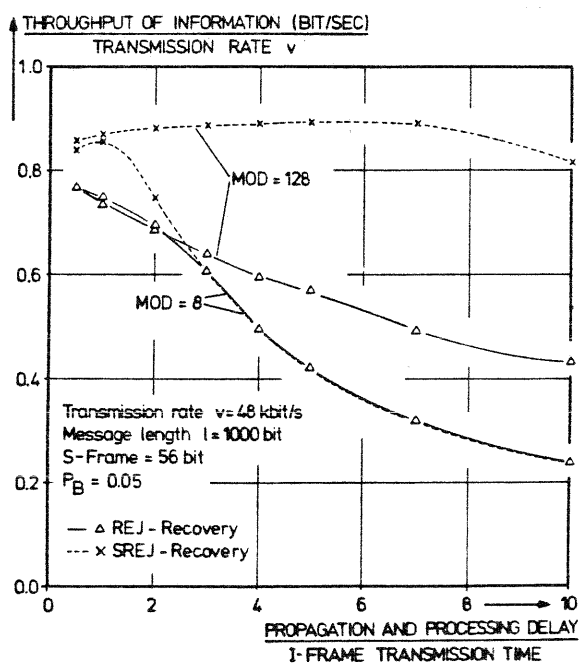

Fig.9: Throughput versus propagation and proc.delay Modulus = 8; 128

In Fig.9 the curve of Fig.8 with P_B = 0.05 and modulus = 8 appears again. But without the window limitation (modulo = 128) there is a very important difference between REJ and SREJ. Using the SREJ-option and modulus extended to 128, the throughput is at a very high level up to 10 t_p/t_I. Operating with REJ-recovery the throughput drops nearly constantly with growing delay. This effect is based on the increasing time a recovery needs on the link with increasing delay. In the case of REJ usage a large number of error-free transmitted I-frames have to be retransmitted after a single disturbed frame.

3.2 Transfer Time Results

Here we consider the case where the channels are only loaded to a fraction of their full capacity to guarantee finite delay. Interesting performance measures in this case are mean transfer time and mean buffer-holding time. The transfer time is defined as the time interval from the arrival of a message at one station until its successful reception at the other station. The buffer-holding time is the interval between the arrival of a message at a station and the reception of an acknowledgement for this message. This means that the transfer time includes queueing, transmission, processing and propagation times, as well as possible additional delays caused by retransmissions. The buffer-holding time comprises the transfer plus the acknowledgement time of a message (c.f. Fig.4).

3.2.1 Transfer Time versus Useful Channel Load

In the following examples we again assume symmetrical traffic configuration, i.e. the traffic in both directions is equal.

Fig.10 shows for a constant message length of 5000 bit the mean transfer time as a function of the useful channel load which is defined as the ratio of throughput of information bits per time-unit and the transmission rate. In case of low

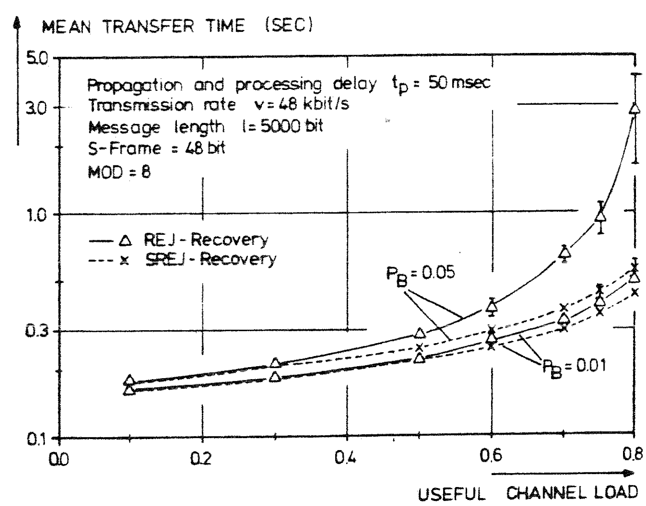

Fig.10: Mean Transfer Time versus Useful Channel Load
P_B *= 0.01; 0.05*

frame error probability (p_B = 0.01) the use of SREJ-recovery saves only little time. On the other side, the link with p_B = 0.05 and REJ-recovery has a distinct increasing transfer time with increasing channel load. The maximum throughput of this link with REJ-recovery is nearly reached, and thus, the waiting times for the first service increases, too (c.f. Fig.11, mean buffer-holding time for the same link).

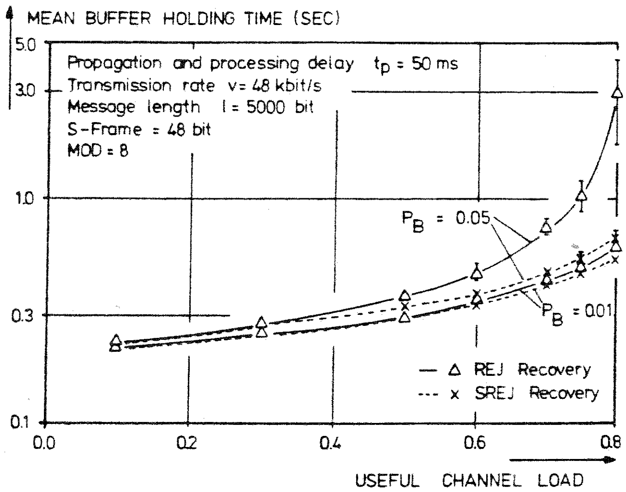

Fig.11: Mean buffer-holding time versus useful channel load
p_B = 0.01; 0.05

In the next diagram the same link is simulated again, but now the message length is randomly distributed. As the maximum throughput is lower (c.f. Fig.6) the transfer time rises in an evident manner at high channel loads, even for SREJ-recovery and p_B = 0.05.

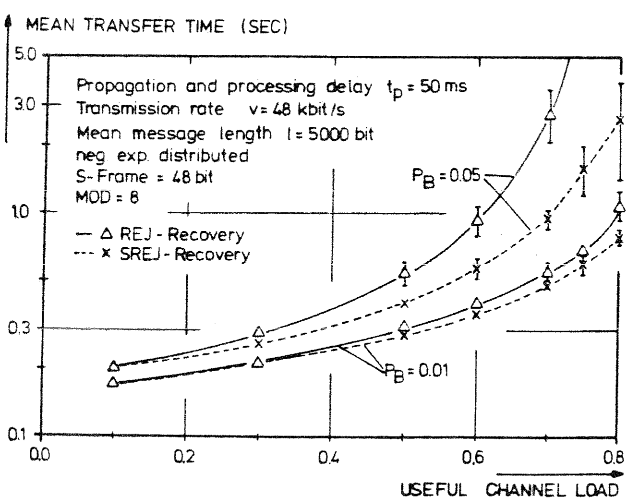

Fig.12: Mean transfer time versus useful channel load
p_B = 0.01; 0.05
message length neg. exp. distributed (truncated)
max. message length = 10 mean message length

3.2.2 Mean Transfer Time versus Message Length

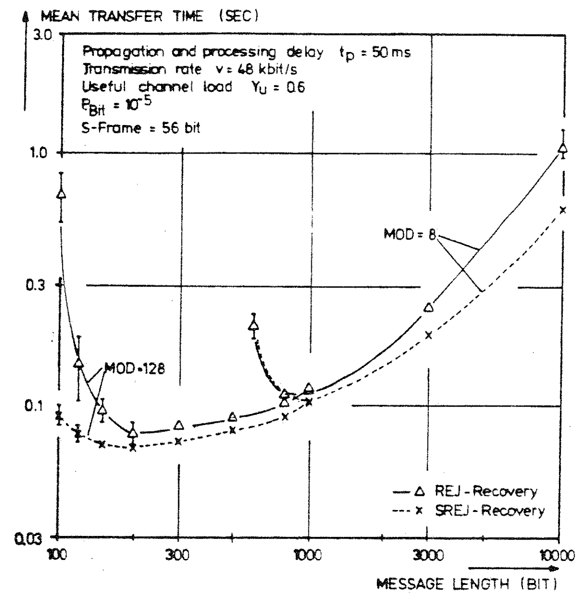

Fig.13: Mean transfer time versus message length
Mod. = 8; 128

Our next study is the influence of the message length to the transfer time. Using modulus = 8 (but S-frame = 56 bit) the transfer time rises according to the increasing message length respectively transmission time. For shorter messages (below 1000 bit/message), the transfer time rises again due to the limitation by the window-mechanism. In this case no improvement is possible using SREJ-recovery. The extended numbering format (modulus=128) allows to transmit even shorter messages. With the REJ-option the fraction of overhead-bits and retransmitted frames grow with decreasing message length and, therefore, the transfer time rises strongly for the given channel load 0.6. In case of SREJ-recovery even for 100 bit message length the transfer time remains rather low, because only the disturbed frame has to be transmitted once more.

3.2.3 Mean Transfer Time versus Propagation Delay

Fig.14 shows the mean transfer time as a function of propagation and processing delay. The mean transfer time is measured for a useful channel load of 0.6. The throughput of the same link is shown in Fig.9. At about 3 t_p/t_I the mean transfer time for modulus = 8 increases strongly. This is due to the fact that the useful channel load of 0.6 is the maximum value (c.f. Fig.9). This means a link loaded up to almost its maximum throughput bears a heavy ascent of the transmission time caused by the growing send-queue and the long waiting times.
Using the extended numbering with modulus = 128 the use of REJ-recovery allows only a maximum delay of 4 t_p/t_I. In Fig.9 this is the point the maximum throughput intersects the line of 0.6. Only the use of SREJ-recovery allows to convey a useful load of 0.6 up to very long delays. In this case the transfer time rises only linearly with growing delay.

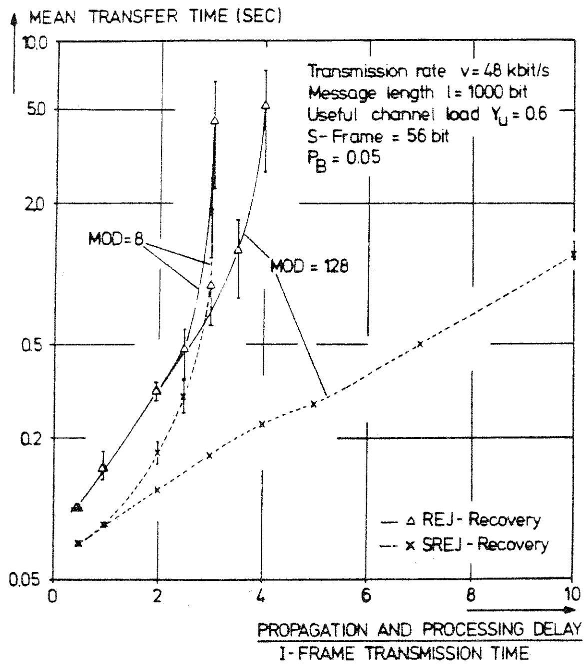

Fig.14: Mean transfer time versus propagation and processing delay Mod = 8; 128

4. CONCLUSION

The major contribution of this paper is to give some suggestions in which cases the performance of a HDLC controlled data-link (ABM, full-duplex circuit) improves by using the SREJ-recovery.

According to the simulation results mentioned above, we give the following advices for using the SREJ recovery-option on a HDLC controlled link:

i) The use of the SREJ-option to improve the performance is only significant if
   - there is no limitation by the window mechanism (modulo-numbering)
   - there is a link with rather high probability of errors.

ii) Best performance results can be achieved by the use of the SREJ-option for links with a
   - high propagation delay and a high modulus-value (e.g. satellite link)

iii) In cases the links are not saturated, there is generally no constraint to implement the SREJ-option. Nevertheless, it will be expedient to establish the SREJ-option in some special constellations of parameters as demonstrated in
   - Fig.13 for message lengths $l \leq 100$ bit
   - Fig.14 for $t_P/t_I > 4$ and modulus = 128.

In this context it should be notified that the implementation of the SREJ-option leads to higher expenses for soft- and hardware of the combined station, such as

- more intelligence due to the greater complexity of the SREJ-recovery
- buffering of the messages at the receiver
- sequencing of the messages after SREJ-recovery actions.

ACKNOWLEDGEMENT

The authors would like to express their gratitude to Prof.(em.) Dr.-Ing. Dr.-Ing.E.h. A. Lotze and Prof. Dr.-Ing. P.J. Kuehn, head of the Institute of Switching and Data Technics, University of Stuttgart, for supporting this work.

REFERENCES

/1/ Truong,H.L.; On the Performance of HDLC-Controlled Data Links.
Report on Studies in Congestion Therory. Institute of Switching and Data Technics, University of Stuttgart, Report No.33
ISBN 3-922403-43-3, 1982.

/2/ Bux,W., Truong,H.L.; High Level Data Link Control - Traffic Considerations.
Proceedings of the 9th International Teletraffic Congress, Vol.2, Session 44, Torremolinos (Spain), 1979.

/3/ Bux,W., Kuemmerle,K., Truong,H.L.; Data Link Control Performance: Results Comparing HDLC Operational Modes.
Computer Networks 6 (1982) pp.37-51
North-Holland Publishing Company.

/4/ Data Communication - High Level Data Link Control Procedures - Frame Structure.
International Standard ISO 3309 (1976).

/5/ Data Communication - High Level Data Link Control Procedures - Elemens of Procedure (independent numbering)
International Standard ISO 4335 (1976).

/6/ Peters,J.; Math and Simulation Tools for Testing HDLC protocol options.
Data Communications, April 1982, pp.111-126.

/7/ Easton,M.C.; Batch Throughput Efficiency of ADCCP/HDLC/SDLC Selective Reject Protocols.
IEEE Transactions on Communications Vol. Com.28, No.2, Feb.1980, pp.187-195.