78

RESULTS ON THE PERFORMANCE OF BALANCED HDLC PROCEDURES

W. Bux,
Institute of Switching and Data
Technics, University of
Stuttgart, Stuttgart, W. Germany

K. Kümmerle,
IBM Zurich Research Laboratory,
8803 Rüschlikon, Switzerland

H. L. Truong,
Institute of Switching and Data
Technics, University of
Stuttgart, Stuttgart, W. Germany

ABSTRACT

The High-Level Data-Link Control (HDLC) Balanced Class of Procedure is intended: (i) for use on links carrying heavy traffic volumes where delays and link efficiency are important factors, and (ii) for situations which require equal control capability at both ends of the link. The prime objective of the present paper is to analyze the performance of this class of procedure, i.e., to quantitatively study the interaction among a multiplicity of parameters which are procedure specific, characterize the properties of the transmission medium, and identify the operational characteristics and requirements. The approach taken is to consider two kinds of operation, a saturated case characterized by optimum throughput as the most suitable measure of performance, and a non-saturated situation for which waiting times and transfer times are the appropriate measures. The results provide a fundamental insight into how the most relevant parameters interact and effect these performance measures. In particular it was shown that an optimum selection of parameters strongly depends on the anticipated mode of operation. Due to the complexity of the problem the analysis was performed by means of simulation techniques.

1. INTRODUCTION

Both international and national standardization bodies and computer manufacturers have defined data-link control procedures, e.g., ISO High-Level Data-Link Control Procedure (HDLC) [1-3], ANSI Advanced Data Communication Control Procedures (ADCCP) [4], Digital Equipment Digital Data Communications Message Protocol (DDCMP) [5], IBM Synchronous Data-Link Control Procedure (SDLC) [6]. All these procedures serve the purpose stated above and contain the same or similar functional elements. In the subsequent discussions we concentrate on HDLC, in particular on the HDLC Balanced Class of Procedure. This standard which is applicable to point-to-point configurations over either dedicated or switched data transmission facilities was developed by ISO and accepted by CCITT. Hence, though not discussed here, the Balanced Mode of HDLC is fully compatible with LAP B, the second level of CCITT Recommendation X.25.

The objective of the present investigation is to study the performance of HDLC when it is operated in balanced mode over a point-to-point configuration. We are interested in: (i) identifying the essential parameters which determine protocol performance, and (ii) analyzing their impact on performance under various conditions. Two different and important traffic situations will be treated. First, the saturated operation where a station always has information to be sent, and, second, the non-saturated operation where the amount of information to be transmitted varies statistically. In the first case, the maximum throughput is the essential measure of performance, whereas in the second case, average delays or delay distribution functions characterize the performance of an HDLC-controlled communication link. Existing studies on the performance of data-link control procedures have mostly addressed the saturated traffic case [7-9]. An analysis of a balanced class of procedure has not yet been published.

In the next section we briefly review the main features of HDLC balanced mode, explain the simulation model, and discuss the major parameters; finally, Section 3 contains a discussion of results.

2. HDLC BALANCED CLASS OF PROCEDURE - PERFORMANCE MODEL

2.1 Background Information on the Procedure

Balanced operation is intended for use on links carrying heavy traffic volumes where delays and link efficiency are important factors. Furthermore, it is intended for situations which require equal control capability at both ends of the link. Therefore for balanced operation the stations at both sides of the link are so-called Combined Stations which means that each station can send both commands and responses and receive both commands and responses, see Fig.1. Since the type of transmission response is autonomous, the class of procedure considered here can be designated as operating in Asynchronous Balanced Mode (ABM). (The classes of procedures for unbalanced operation, contrariwise, have two types of stations, primary and secondary, and can use two types of transmission response, normal and asynchronous.)

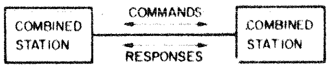

Fig. 1 Balanced configuration

CH1354-0/78/0000-0127$00.75 © 1978 IEEE

All transmissions are in frames, and each frame conforms to one of the formats shown in Fig.2 provided the unextended control-field format is used: frames transporting information, I-frames, are delimited flags, and are composed of an address field which contains either the local or remote station address, a control field which contains commands or responses, and sequence numbers, an information field which may contain any sequence of bits, and, of a frame-checking sequence. Frames containing only supervisory control sequences, S-frames, have the same structure but no information field.

a) I-FRAME

FLAG	ADDRESS	CONTROL	INFORMATION	FRAME CHECK SEQUENCE	FLAG
01111110	8 bit	8 bit	X bit	16 bit	01111110

b) S-FRAME

FLAG	ADDRESS	CONTROL	FRAME CHECK SEQUENCE	FLAG
01111110	8 bit	8 bit	16 bit	01111110

Fig. 2 Formats of I- and S-frames
(unextended control field)

For the particular purpose of this performance study, we assume: (i) that both setting up and disconnecting the link have been appropriately handled, i.e., these actions are considered as being outside the link level of this analysis, and (ii) that the established link is available and operational for the time interval required. Therefore, only the following commands and responses from the basic repertoire are used:

• I:	Information,	Information-transfer format commands/responses
• RR:	Receive ready,	Supervisory format commands/responses
• RNR:	Receive not ready.	

In addition, we use the command/response REJ (Reject) being offered under optional functions for improved performance.

In order to illustrate how these commands/responses are being used in the model to be discussed in Section 2.2, we subsequently outline various characteristic patterns of operations under the assumption that the reader is familiar with the HDLC double numbering scheme.

a) Error-Free Operation

The simplest case occurs when both stations have I-frames ready for transmission all the time. Then I-frames can be used to acknowledge reception. Should transmission and propagation delays be such that an acknowledgement has not been received after having sent MODULO minus one I-frames (where MODULO is the modulus of the sequence numbering scheme and the numbers cycle through the entire range), then sending further I-frames is held until an acknowledgement arrives.

In the case where a station has no more I-frames to be sent, it can acknowledge incoming frames by sending supervisory frames either with the command RR if further I-frames can be accepted, or with the command RNR if no further I-frames can be accepted.

b) Operation with Errors

Frames with a Frame-Check Sequence (FCS) error are discarded and have to be retransmitted. The error will manifest itself later in the form of a sequence error. We can differentiate among three mechanisms to recover errors.

b1) *Checkpoint Retransmission (P/F-Bit Recovery).* As the P and F bits are always exchanged as a pair (for every P there is one F, and the P cannot be issued until the previous P has been matched with an F), the N(R) count contained in a frame with the F bit set to "1" can be used to detect I-frame sequence errors. A combined station will initiate error recovery if the received N(R) does not acknowledge all I-frames transmitted by the combined station prior to and including the last command frame sent with the P bit set to "1". Retransmission starts with the lowest numbered unacknowledged I-frame. Since as a general feature, we would like to have fast error recovery, we have adopted the strategy of setting the P bit, thus enforcing responses with the F bit set to "1", as often as possible.

b2) *REJ Recovery.* The REJ command/response is used to initiate an earlier retransmission following the detection of a sequence error than is possible by checkpointing, e.g., if REJ is immediately transmitted upon detection of a sequence error there is no need to wait for a frame with the F bit set to "1". Retransmission starts with the I-frame indicated by the N(R) count contained in the REJ frame.

b3) *Time-Out Recovery.* A single I-frame or the last I-frame in a sequence of I-frames cannot be recovered by REJ. Also, a frame with the P bit set to "1" may be lost. Therefore, each combined station has a timer which starts when a frame is sent with the P bit set to "1". Upon expiration of the time-out, recovery is initiated. Since it can happen that an I-frame has been correctly received, but the acknowledgement has either not been sent or lost, an RR command with the P bit set to "1" is issued prior to retransmission in order to avoid a duplication of the I-frame already sent.

c) Busy Condition

The Receive-Not-Ready, RNR, is used by a station to indicate a busy condition, i.e., temporary inability to accept additional incoming I-frames. Such a situation may occur if buffer space is close to saturation. I-frames numbered up to N(R)-1 are acknowledged. I-frame N(R) and any subsequent I-frames received, if any, are not acknowledged. To indicate that the busy condition has been cleared, the receive-ready, RR or any supervisory frame, is sent with the P bit set to "1" if no P bit is outstanding at this time.

2.2 The Model

a) Configuration

A schematic representation of the model underlying the performance study is shown in Fig.3. The two stations are connected by a full-duplex transmission link which is characterized by transmission speed, propagation delay, and a bit-error probability. The bit errors are assumed to be statistically independent; any other error model, however, could also be implemented. The link, for instance, can be conceived as trunk line between switching nodes in a data network. Message blocks to be transmitted from A to B and vice versa are stored in message buffers. Each message buffer is controlled by a buffer manager who has the following main responsibilities. First, it denies acceptance of new message blocks to be transmitted in case of buffer saturation or when buffer saturation is imminent. Second, the buffer manager is responsible for keeping message blocks and control information stored until they are successfully transmitted across the link. Third, on the one hand, it has to notify the HDLC procedure whether buffer saturation is imminent, which then causes HDLC to signal a busy condition, and, on the other hand, to indicate whether the buffer is again available, which causes HDLC to signal a clear busy condition. Fourth, it has to remove message blocks received from the other station after a specified time, e.g., when these blocks have been delivered to a local station or to an adjacent node. Finally, the buffer manager is responsible for avoiding deadlocks, where the particular deadlock prevention scheme, which has been implemented in our model, is similar to the one suggested in Ref. [10]. For the present investigation, we have assumed that the buffer manager operates in zero time, e.g., it is not slowed-down by the processor in which the algorithms are implemented.

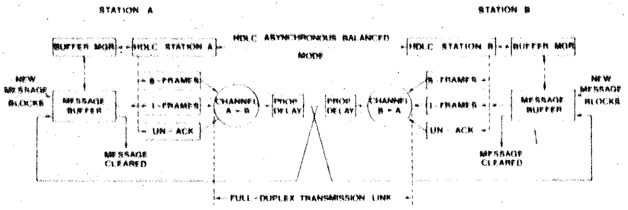

Fig. 3 Structure of performance model

b) Flow of Control and Data

The transfer of information across the full-duplex link is controlled by HDLC operating in the Asynchronous Balanced Mode (Fig.3) as explained in Section 2.1. Message blocks are assumed to be of constant length. An I-frame can contain several message blocks up to a prespecified maximum number, provided several blocks are ready for transmission at the time when the frame is assembled. Fig. 3, furthermore, shows three "logical" queues controlled by the procedure: the I-frame queue containing those information frames which have not yet been sent; the S-frame queue containing those supervisory frames which have not yet been sent; and the UN-ACK queue containing I-frames which have been sent but have not yet been acknowledged. The elements of the latter are potential candidates for retransmission. The priority scheme is determined by the HDLC procedure. HDLC also has to notify the buffer manager when frames can be removed from the UN-ACK queue, i.e., physically speaking from the message buffer when they have been acknowledged.

c) Simulation Model and Parameters

The model depicted in Fig.3 has been implemented in ALGOL and employs event-by-event simulation. It can handle symmetrical and asymmetrical traffic flows and equal or different transmission speeds in both directions of the full-duplex link. Apart from the HDLC balanced class of procedure as discussed in Section 2.1, the model is determined by the following parameters:

- Maximum length of information field in I-frames,
- MODULO value of the HDLC frame-numbering scheme,
- Transmission speed in each direction,
- Propagation delay in each direction,
- Bit-error probability,
- Length of message blocks to be transmitted in I-frames,
- Arrival rate of new message blocks,
- Distribution of interarrival times of new message blocks,
- Size of the message buffer.

The quantities to be obtained from the model will be explained in the context of the discussion of results in Section 3.

3. NUMERICAL RESULTS

In this section, we present and discuss results obtained from simulation runs. Before the results proper are addressed, some general statements concerning the measurements are necessary.

3.1 Outline of Measurements

In order to obtain a broad view of the performance characteristics, we distinguish between two traffic situations: saturation and non-saturation. The assumption for the saturated case is that at both stations at any time message blocks are available for transmission. This implies that arrival rates of new message blocks and interarrival times are irrelevant. The measure of performance which most appropriately characterizes this mode of operation is throughput. The throughput which can be achieved under this condition is the maximum to be obtained. In the non-saturated case, the arrival rates of new messages are such that the assumption for the previous case does not hold. This represents the more realistic situation and means that the transmission channels are not fully loaded for delay reasons. As a consequence, the most relevant performance measures for this case are average delays and delay distribution functions. Throughput by definition is less than the maximum achievable throughput and, if at all, of secondary importance.

In consonance with the above considerations there are basically two categories of results: maximum throughput, or more precisely, maximum throughput of information bits, and average delay both as functions of various parameters. The maximum information throughput has been determined as the function of the following parameters: maximum length of the information field in I-frames, bit-error rate of the transmission link, modulo value used in the HDLC numbering scheme, and propagation delay which is almost zero for terrestrial links and substantially greater than zero for satellite links (Figs. 4,5). As far as average delays are concerned, we distinguish between mean waiting time and mean transfer time where the first one is defined as the time interval from the (new) arrival of a message block until the beginning of its first transmission, the second is specified through the time interval between the (new) arrival of a message block and its correct reception (including potential retransmissions) at the other station. These delays have been studied as functions of the length of the information field in I-frames, the message block length, the interarrival distribution function, and the maximum number of message blocks per I-frame (Figs. 6,7).

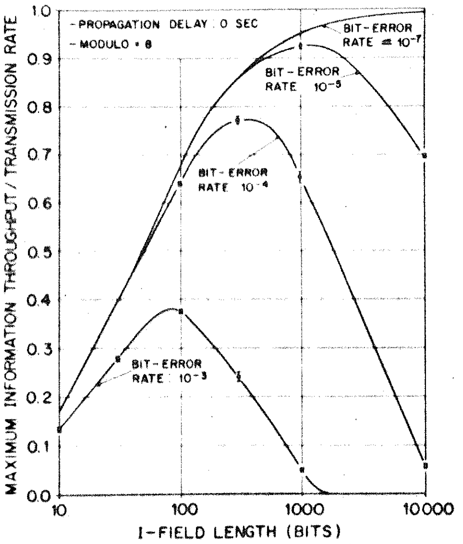

Fig. 4 Maximum information throughput/transmission rate vs. I-field length for terrestrial links

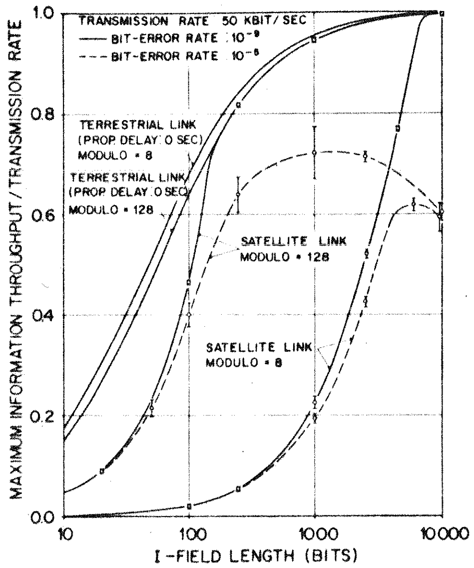

Fig. 5 Maximum information throughput/transmission rate vs. I-field length for terrestrial and satellite links

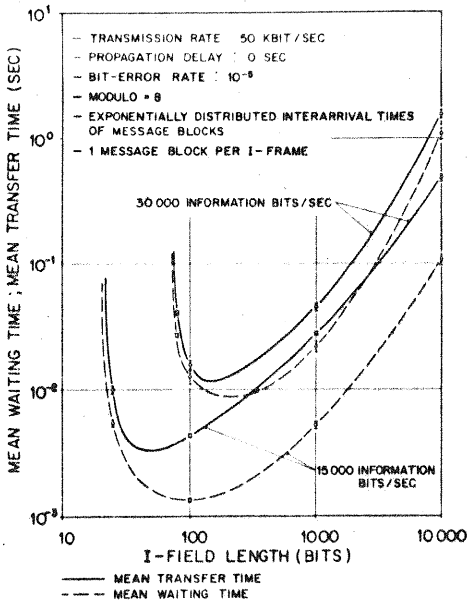

Fig. 6 Mean waiting time and mean transfer time vs. I-field length

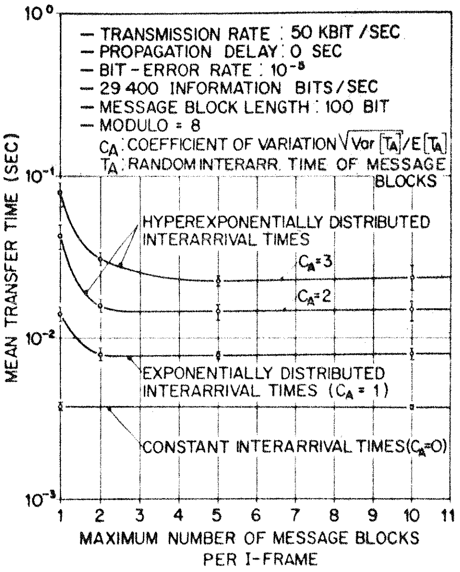

Fig. 7 Mean transfer time vs. maximum number of message blocks per I-frame

Further quantities which will be discussed for the non-saturated traffic case comprise channel load and mean number of retransmissions per I-frame (Figs. 8,9).

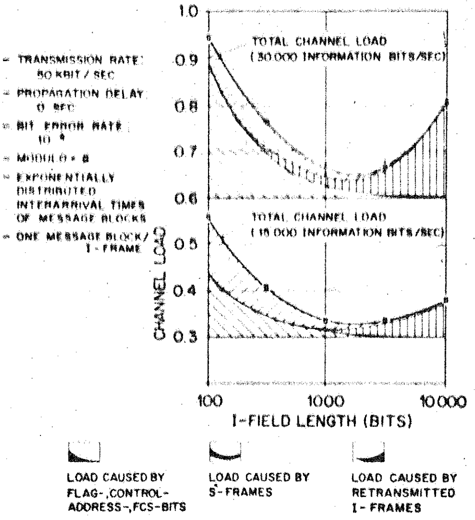

Fig. 8 Channel load vs. I-field length

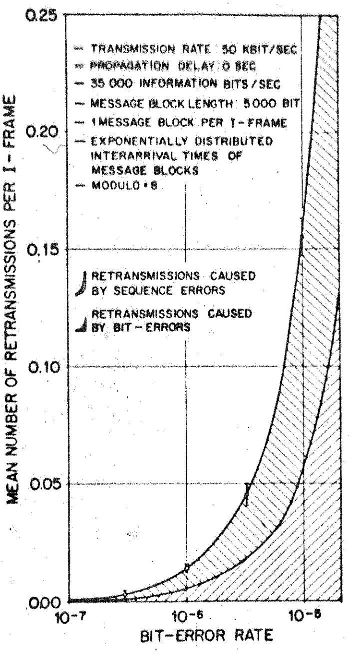

Fig. 9 Mean number of retransmissions vs. bit-error rate

3.2 Discussion of Results

Two general remarks hold for the following discussion:

1) The simulation results in Figs.4-9 are shown with the 95% confidence intervals.
2) The quantities in all diagrams are related to one direction of transmission on the transmission link.

a) The Saturated Case

Figure 4 shows the maximum throughput of information bits (message bits) relative to the transmission rate (speed) as a function of the length of the information field in I-frames. Since a terrestrial link has been assumed, the propagation delay is practically zero; the modulo value is eight, and the bit-error rate (bit-error probability) varies from 10^{-3} to 10^{-7}. The ideal throughput characteristic in case of no errors follows the ratio of I-field length to I-frame length where an I-frame carries a constant overhead of 48 bit for the unextended control field (Fig. 2) and 56 bit for the extended control-field format:

$$\frac{ITHmax}{TR} = \frac{I}{S + I} \qquad (1)$$

ITHmax: maximum information throughput
TR: transmission rate
I: I-field length in I-frames
S: overhead per I-frame = length of supervisory frame.

For bit-error probabilities equal to or less than 10^{-7} the throughput characteristic starts to follow this ideal line, i.e., the longer the I-field the better the throughput. For bit-error rates $\geq 10^{-5}$ the curves show distinct maxima in the range considered here. This behavior is typical for transmission strategies employing retransmission for error recovery [7]. The physical explanation is as follows. For short I-field length the probability that an I-frame gets disturbed (block error probability) is low and at the same time the overhead relative to the number of message bits carried in a frame is high. For longer I-field lengths the relative overhead decreases but the block error probability, and thus retransmission activity, increases. Therefore the curves necessarily show a maximum. An additional interesting feature is that improving the bit-error rate from 10^{-3} to 10^{-4} yields a gain in throughput of almost 100%, whereas the step from 10^{-4} to 10^{-5} yields approximately another 20%.

The impact of the modulo value and propagation delay of the parameter is shown in Fig.5 by comparing terrestrial and satellite links. First, the terrestrial case is repeated for a bit-error rate of 10^{-9} and MODULO = 8, 128. Both curves follow very closely the ideal throughput characteristic according to equation (1); the impact of the modulo value is not significant. Second, we have a satellite link with a one-way propagation delay P_D = 270 ms. If we assume the same bit-error rate 10^{-9}, which is a realistic value for satellite channels and MODULO = 8, the maximum information throughput is significantly reduced. This is caused by the rule that transmission of I-frames has to be discontinued if MODULO-1 I-frames are unacknowledged. For a significantly higher modulo value, 128, the throughput is substantially improved and coincides with the values of a terrestrial link for I-field lengths ≥ 160 bits. In case of an error-free channel and a constant I-field length one can calculate upper and lower bounds for the throughput by determining both the minimum and maximum times required for acknowledging an I-frame:

$$g_1 \cdot \frac{I}{S + I} < \frac{\mathrm{ITHmax}}{\mathrm{TR}} \leq g_2 \cdot \frac{I}{S + I} \tag{2a}$$

$$g_1 = \mathrm{Min} \quad 1, \frac{(\mathrm{MODULO} - 1)\ (S + I)}{2 \cdot P \cdot \mathrm{TR} \ + \ 3(S + I)} \tag{2b}$$

$$g_2 = \mathrm{Min} \quad 1, \frac{(\mathrm{MODULO} - 1)\ (S + I)}{2 \cdot P \cdot \mathrm{TR} + 2 \cdot S + I} \tag{2c}$$

P = Propagation delay in one direction.

Figure 5 in addition shows how the maximum throughput decreases in case of a higher bit-error rate which may be caused by a poor quality terrestrial extension of the satellite channel.

b) The Non-Saturated Case

Figure 6 represents the average waiting times and average transfer times defined at the end of Section 3.1 as functions of the I-field length. The parameters have been chosen as follows: transmission rate in each direction 50 kbit/s, terrestrial link, i.e., propagation delay practically zero, bit-error rate 10^{-5}, modulo value = 8, exponentially distributed interarrival times of message blocks, each I-frame carries one message block (of constant) length. Two actual throughput values (information bits) are considered: 15000 bit/s and 30000 bit/s corresponding to a link utilization due to information of 30% to 60%, respectively. Both average waiting time and average transfer time show distinct minima. The steep rise of each curve to the left of its minimum, i.e., for small values of the I-field length, is due to the relatively large overhead for flags, address control, and FCS fields. In other words, since the useful throughput is maintained at the same value, an appreciable number of overhead (framing) bits is required to achieve this throughput.

With growing I-field lengths the framing overhead decreases, but there are three other effects causing a rise of the curves to the right of the minimum: (i) the block error probability grows, (ii) error recovery takes longer, and (iii) waiting and transfer times, even if there are no errors, grow proportionally with the length of I-frames. A comparison of Fig. 4 with Fig. 6 shows that the shortest delays occur at significantly smaller I-field lengths than the maximum throughput values: the increase of transmission time and block error probability with growing I-frame lengths outweighs the impact of the decreasing overhead at smaller I-field lengths. Furthermore, we can observe that the average transfer times reach their minimum for smaller values of the I-field length than the average waiting times because the transfer time includes the transmission time which itself grows with the I-field length.

Finally, we conclude the discussion of this figure with the statement that it is possible to calculate lower bounds for the average waiting and transfer times by means of results readily available from M/D/1 queueing systems.

So far we have assumed that the information field of an I-frame contains a single message block. We now drop this assumption and investigate the potential reduction of the mean transfer time particularly in the light of interarrival processes with clustered arrivals of message blocks by allowing more than one message block in an I-frame. Figure 7 depicts the mean transfer time as the function of the maximum number of message blocks to be carried in an I-frame. The parameters are as follows: transmission rate 50 kbit/s, propagation delay practically zero, bit-error rate 10^{-5}, message block length 100 bit, modulo value = 8, actual throughput (information bits) 29400 bit/s. Three types of interarrival processes are compared, each characterized by its coefficient of variation C_A: (i) constant interarrival times ($C_A = 0$), exponential interarrival times ($C_A = 1$), and hyperexponentially distributed interarrival times ($C_A = 2,3$). As expected, the transfer time is independent of the maximum number of message blocks per I-frame for $C_A = 0$. In case of exponentially distributed interarrival times, the reduction in transfer time relative to the leftmost abscissa value amounts to roughly 40%. If the interarrival time distribution has a greater variance but the same mean value, $C_A = 2,3$, the gain is more pronounced and reaches approximately 75% for $C_A = 3$. It should be noted, however, that this gain strongly depends on the lengths of message blocks. Another potential advantage, though not investigated here, may consist in saving processor cycles in the processor driving the procedure.

The next question we address is the channel load or channel utilization as a function of the I-field length where the total load is broken down into its constituent components, Fig.8. This will provide an additional view from a different angle and help to gain a clear picture about the interplay of the relevant mechanisms of the procedure. The parameters are the same as in Fig.6, and again two actual throughput values (information bits) are considered: 15000 bit/s and 30000 bit/s. First, it is obvious and has been explained before that the load due to flag, address, control, and FCS bits decreases with increasing I-field length. Second, the load due to retransmitted I-frames increases with the length of the I-field since the block-error probability increases. Third, the load caused by the transmission of S-frames decreases with a growing I-field length, since the actual information throughput is kept constant which in turn reduces the number of I-frames and, consequently, the load due to supervisory frames. Finally, the total channel load indicates a minimum at a certain I-field length which, if compared with the optimum length of the I-field, in the saturated case (Fig.4) for maximum throughput is slightly shifted to the right.

Finally, Fig.9 represents the mean number of retransmissions of I-frames relative to a successfully transmitted I-frame versus the bit-error rate. Parameters: transmission rate 50 kbit/s, propagation delay zero, actual information throughput 35000 bit/s, one message block per I-frame, message block length 3000 bit, exponentially distributed interrival times. HDLC retransmits I-frames for two reasons: (i) the frame is disturbed (FCS error) and therefore discarded; (ii) the frame is correct but out of sequence because one of the previous frames was disturbed and discarded or lost. The figure clearly indicates that the fraction of retransmissions due to sequence errors constitutes a substantial portion of the total number of retransmissions and, thus, points at potential benefits if the optional Selective REJect (SREJ) command/response is employed.

4. CONCLUSIONS

The major findings of this investigation are:

1) The distinction of two modes of operation, the saturated case and the non-saturated case, is the key for performance considerations of this type since it allows focusing on the conflicting issues of throughput and delay separately.
2) The HDLC Balanced Class of Procedure allows efficient utilization of the transmission links and low waiting and transfer times provided the relevant parameters are adjusted to meet the needs of the specific application.
3) The interaction among procedure specific parameters, parameters characterizing the transmission medium, and the required operational parameters is very complex. The results of this investigation help to clarify this issue and, thus, can be used to specify the relevant parameters as indicated in the previous point.

Some further findings follow:

a) In the saturated case and for non-zero bit-error rates, the throughput has distinct maxima determined by the length of the information field in I-frames and by the bit-error rate. For low error rates, the throughput is primarily limited by the framing overhead and approaches the theoretical limit. The modulo value significantly impacts throughput for satellite links, whereas the throughput of terrestrial links is insensitive to this quantity.
b) In the non-saturated case, waiting and transfer times show distinct minima determined by the length of the information field in I-frames and by the information throughput. The optimum I-field length, however, is considerably smaller than required if maximum throughput is the objective. Furthermore, it has been demonstrated that for clustered message-block arrivals the mean transfer time can be reduced if several message blocks are carried in a single I-frame. Finally, it is important to realize that for a given information throughput the HDLC-controlled link can be operated with a minimum amount of overhead which results in a minimum channel load.

ACKNOWLEDGEMENT

The authors would like to thank Prof. Dr. A. Lotze for his continued interest in the work, J. Grinenwald for writing the HDLC part of the model, P. Kühn for helpful discussions, and H. Rudin for reviewing the manuscript.

REFERENCES

[1] Data Communication - High Level Data Link Control Procedures - Frame Structure, International Standard ISO 3309.

[2] Data Communication - High Level Data Link Control Procedures - Elements of Procedures (independent numbering). International Standard ISO 4335

[3] HDLC - Proposed Balanced Class of Procedures, Document ISO/TC 97/5C6-No. 1444.

[4] American National Standard for Advanced Data Communication Control Procedures (ADCCP). Document ANSI/B5R 3.66.

[5] Teicholtz, N.A.: Digital Network Architecture, Proceedings of the European Conference on Communications Networks, London, 1975, pp. 13-24.

[6] IBM Synchronous Data Link Control - General Information. GA 2730930, 1974.

[7] Chu, W.W.: Optimal Message Block Size for Computer Communications with Error Detection and Retransmission Strategies. IEEE Transactions on Communications, Vol. COM-22, No. 10, 1974, pp. 1516-1525.

[8] Traynham, K.C. and Steen, R.F.: SDLC and BSC on Satellite Links: A Performance Comparison. ACM Computer Communication Review, Oct. 1977, Vol. 7, No. 4, pp. 3-14.

[9] Gelenbe, E., Labetoulle, J., and Pujolle, G.: Performance Evaluation of the Protocol HDLC. Proceedings of the Symposium on Computer Network Protocols, Liège, Belgium, February 13-15, 1978, pp. G3.1-G3.5.

[10] Guenther K.: Vermeidumg von Deadlocks im Allgemeinen und von Pufferdeadlocks bei Store-and-Forward-Communication im Besonderen. Gesellschaft für Mathematik und Datenverarbeitung, GMD Dokument No. 20, 1975.