Reprinted from NATIONAL TELECOMMUNICATIONS CONFERENCE,
November 30 - December 4, 1980

HDLC PERFORMANCE: COMPARISON OF NORMAL RESPONSE MODE AND ASYNCHRONOUS BALANCED MODE OF OPERATION

Werner Bux and Karl Kümmerle
IBM Zurich Research Laboratory, 8803 Rüschlikon, Switzerland

Hong Linh Truong
Institute of Switching and Data Techniques, University of Stuttgart, 7000 Stuttgart, W. Germany

ABSTRACT: The objective of this paper is to compare the performance of the different operational modes of the standard data link-control procedure HDLC: The Normal Response Mode (NRM), an unbalanced procedure with primary and secondary stations which have different functional capabilities, and the Asynchronous Balanced Mode (ABM), a balanced procedure with functionally equivalent, so-called combined stations. Performance is measured in terms of: 1) the maximum throughput to be achieved over a point-to-point link, and 2) the mean transfer time of messages as a function of the link utilization. To obtain an accurate account of the performance of these procedures, the approach taken is to implement the two protocols in a simulation model of a data link, i.e., including all details of the sequencing mechanism and of the error-recovery procedures. The general conclusion from the study is that throughput and delay performance of ABM is better than, or at least equal to NRM. However, there exists a broad range of parameters where the performance of both modes is similar.

INTRODUCTION

The increasing need for reliable and efficient data communication has forced computer manufacturers and standardization bodies to develop a number of data link-control procedures. The majority of these procedures are contained in the link-control procedures standardized by ISO under the name of HDLC (High Level Data Link Control) [1]-[3]. HDLC represents a data link-control architecture applicable to a wide variety of links: half-duplex, full-duplex, point-to-point, and multipoint.

In this paper, we restrict ourselves to the consideration of full-duplex, point-to-point links. For this kind of operation, we investigate the two operational modes which seem to be the most frequently applied in practice: i) Normal Response Mode (NRM), a procedure based on the master-slave principle using two types of stations with different functional capabilities, primary and secondary, and ii) Asynchronous Balanced Mode (ABM), a symmetrical procedure with fully equivalent, so-called combined stations.

The main differences between NRM and ABM are:

In *NRM*, the primary station is responsible for link initialization and disconnection and for error recovery. The secondary station is only capable of responding to commands from the primary. In particular, a secondary is not allowed to transmit without having received explicit permission from the primary in the form of a polling command. This strict master-slave policy is mainly motivated by the operation of half-duplex links and of multipoint configurations.

For point-to-point, full-duplex links, such an unbalanced operation is possible but not necessary: *ABM* is a fully symmetrical procedure in the sense that the combined stations are equally responsible for link initialization and disconnection, and for error recovery. Furthermore, the type of transmission is asynchronous, meaning that both stations can transmit without permission from the other station.

This paper is intended to shed more light onto the performance aspect of both modes. To obtain accurate, unbiased, and detailed results, our approach was to implement both procedures within the simulation model of a full-duplex, point-to-point link. More precisely, we have implemented the *information transfer part* of both procedures including all details of the sequencing and error-recovery mechanisms. The transmission medium is determined through transmission rate, propagation delay, and error characteristic.

In the next section, we briefly describe the major characteristics of both HDLC procedures. This is followed by a short outline of the simulation model, while the main part of the paper is a discussion and comparison of the typical performance results for both modes. Finally, we present the general conclusions.

HDLC - NORMAL RESPONSE MODE AND ASYNCHRONOUS BALANCED MODE

In this section, we briefly review the main features of HDLC and discuss the commonalities and differences of the elements of procedure for the two modes of operation subsequently investigated. We assume that the reader is familiar with the basics of HDLC, namely frame format and the double-numbering scheme, see e.g., [4].

The following commands and responses from the basic repertoire have been implemented in the simulation model: Information (I), Receive Ready (RR), and Receive Not Ready (RNR). In addition, we use the command/response Reject (REJ) being offered under optional functions for efficient error recovery.

CH1539-6/80/0000-0070 $00.75 © 1980 IEEE

Error-Free Operation

There are two properties of the HDLC procedure which mainly influence the performance in the error-free case: the modulus value of the sequence numbers and the usage of the Poll/Final (P/F) bit which is different in the two modes considered:

1. Modulus of Sequence Numbers: The sequence numbers of the I-frames cycle through the set of numbers 0, 1, 2, ..., M-1, where M is the modulus value. The modulus equals eight for the unextended and 128 for the extended format. Under certain conditions, e.g., long propagation delays, it may influence the data flow over the link, because a station must stop transmitting I-frames if it has M-1 unacknowledged I-frames simultaneously outstanding.

2. P/F-Bit Usage: The P-bit is generally used to solicit a response from the other station. Apart from time-out recovery situations, a station may have only one P-bit outstanding at a given time. Before it can issue another frame with the P-bit set to 1, it must receive a response from the other station with the F-bit set to 1. The particular meaning of the P- and F-bit is different in NRM and ABM.

In *NRM* the secondary cannot transmit until a command with the P-bit set to 1 is received. The secondary must set the F-bit to 1 in the last frame of its response. Following transmission of the frame with the F-bit set to 1, the secondary must stop transmitting until another frame with the P-bit set to 1 is received. The strategy in our implementation is to poll the secondary station as often as possible to give it the opportunity to transmit at the earliest point in time.

In *ABM*, each combined station can send both commands with or without P-bit and responses with or without F-bit. In contrast to NRM, each station transmits on an asynchronous basis, i.e., without having to wait for any permission from the other station. The P-bit is used to solicit a response at the earliest opportunity with the F-bit set to 1. For example, if a station wants to get positive acknowledgment that a particular command has been received, it may set the P-bit in that command to 1.

Error-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 P/F-bit recovery.

1. P/F-Bit Recovery: As P- and F-bits are always exchanged as a pair, they can be used for checkpointing purposes.

If in *NRM* a station receives a frame with the P/F-bit set to 1, it initiates retransmission of unacknowledged I-frames with sequence numbers less than or equal to the N(S) number of the last frame transmitted with the P/F-bit set to 1.

In *ABM*, checkpoint retransmission is only initiated based on frames received with the F-bit set to 1, to avoid possible interference with other recovery possibilities. HDLC does not specify for ABM under which conditions the P-bit has to be set. We adopted the strategy of setting the P-bit only in those cases where it is necessary to query the status of the other station, e.g., after the time-out has expired (see section on Time-Out Recovery below).

2. REJ Recovery: The REJ command/response is used for a more timely reporting of sequence errors. Its use is explained in the following example: Assume that the I-frame with N(S) = 1 of Station-A is received in error and, therefore, discarded by Station-B. When Station-B receives the next error-free I-frame, e.g., the frame with N(S) = 2, it informs Station-A of a sequence error by issuing a REJ frame with N(R) = 1. Upon receipt of this frame, Station-A retransmits the requested I-frame with N(S) = 1 plus all additional I-frames which have been subsequently transmitted. To avoid retransmission of correct I-frames which followed a disturbed I-frame, HDLC provides the optional Selective-Reject function (SREJ), where only the disturbed I-frame has to be retransmitted. Due to the greater complexity, the option SREJ has been implemented only rarely in current systems.

If a retransmitted I-frame is again received in error, then the REJ recovery can not be repeated for reasons of possible ambiguities. The error situation has then to be resolved either by checkpointing or by time-out recovery.

3. 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. To recover from such error situations, HDLC provides a time-out function. Use of the time-out also differs in NRM and ABM.

In *NRM* only the primary station has a time-out function. Its use is not specified in the Standard [2], [3]. In Annex B of [2], however, labeled "Timer Considerations" the following recommendations are given: The timer is started upon transmission of a command with the P-bit set to 1. It is restarted when an error-free frame with the F-bit set to 0 is received, and it is stopped when an error-free frame with the F-bit set to 1 is received. Upon expiration of the timer, it is recommended that the primary queries status with a supervisory frame.

In *ABM*, each combined station has a time-out function. Again, HDLC does not specify how the timer should be handled. For our implementation, we adopted the following rules: The timer is started (provided it is not running already) every time an I-frame or a frame with the P-bit set to 1 is sent. It is restarted when a frame is received which acknowledges a not-yet acknowledged I-frame. The timer is stopped when no unacknowledged I-frames and no P-bit are outstanding. Upon expiration of the timer, an RR command with the P-bit set to 1 is issued prior to retransmission to avoid duplication of I-frames already sent.

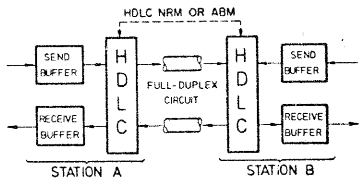

Figure 1 Structure of the data link model.

PERFORMANCE MODEL

Figure 1 shows the data link model used for our simulation study. It consists of two data stations connected by a full-duplex circuit. The link is controlled by HDLC, operating either in NRM or in ABM. In NRM, Station-A is designated as primary, Station-B as secondary. In ABM, both stations are combined stations.

Messages to be transmitted are stored in the send buffer of the sending station where they have to wait for transmission. Messages are transmitted according to first-come, first-served, one message per I-frame. Throughout this paper, we assume that the size of the message buffers is unlimited and that the messages are either of constant length ℓ or have exponentially distributed lengths with mean ℓ. The transmission channels are characterized by their transmission rate v, their bit-error probability p_{bit} (independent bit errors), and their (one-way) progagation delay t_{prop}. Furthermore, we assume that processing of a received frame requires the constant time t_{proc}. For this investigation we combine propagation and processing delay in a constant but arbitrarily selectable delay $t_p = t_{proc} + t_{prop}$.

As already pointed out, the assumption for the simulation model is that the link has already been initialized. Apart from that, all mechanisms of the procedures have been implemented in full detail to obtain accurate and reliable performance results.

RESULTS

This is the main section of the paper in which typical performance results for the two HDLC modes, NRM and ABM will be discussed.

We distinguish between two categories of results: maximum throughput and mean transfer time. In the first case, we determine the maximum number of information bits which can be transmitted in one direction of the link per unit time. This means, we assume that for the direction considered, information to be sent is available at any point in time. In the second case, we are interested in the transfer time of messages, defined as time interval from the arrival (or generation) of a message at one station until its successful, i.e., error-free reception at the other station. Since the transfer time includes queueing delays, the underlying assumption is that the channel load in the direction considered is only a fraction of the full channel capacity.

Two different link types are considered: A terrestrial link with processing plus propagation delay of 50 msec and a satellite link with processing plus propagation delay of 350 msec. The above values of t_p are the same as in [5]. All simulation results are shown with their 95% confidence intervals.

Throughput Results

Figures 2 to 4 show typical results for the maximum throughput in one direction of the link as a function of message length, i.e., the length of the information field of I-frames. With respect to the traffic load on the reverse channels, we assume that the reverse channel is fully loaded, too ("two-way flow").

The throughput results are in consonance with the general behavior of link-control procedures employing an error-detection and retransmission scheme: (i) The 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, where the latter contains a fixed amount of overhead bits for flags, address, control, and FCS fields (c. f. [1]). This explains why maximum throughput must decrease for very short messages.

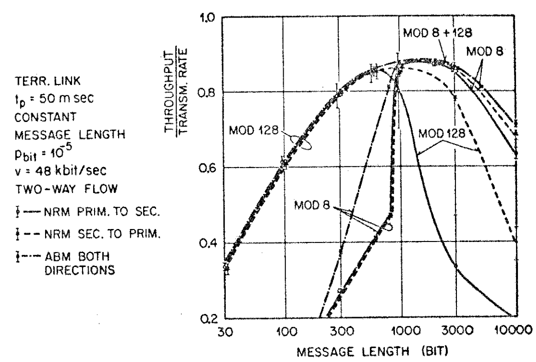

Figure 2 Maximum throughput relative to transmission rate vs. message length. Two-way flow. Constant message length.

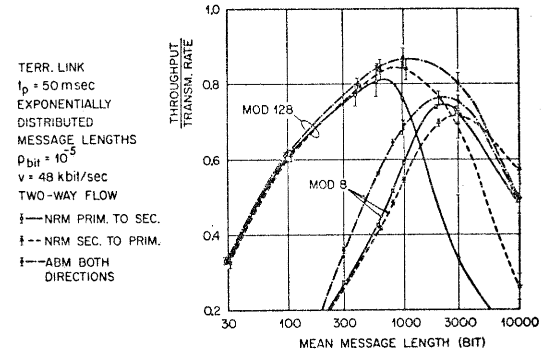

Figure 3 Maximum throughput relative to transmission rate vs. mean message length. Two-way flow. Exponentially distributed message lengths.

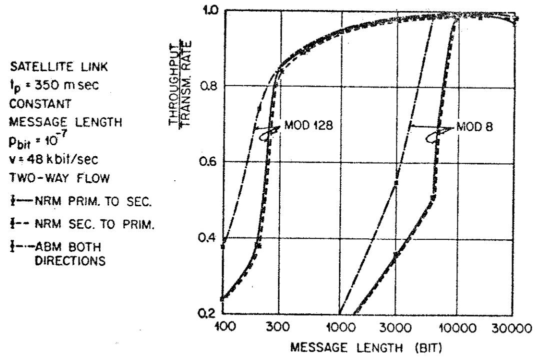

Figure 4 Maximum throughput relative to transmission rate vs. message length. Two-way flow. Satellite link.

(ii) Since the block-error probability of the I-frames increases with growing frame length, the fraction of error-free, i.e., useful transmissions, decreases. These two effects explain why maximum information throughput must decrease for very short and also for very long messages. However, they are absolutely insufficient to explain the actual throughput characteristic of a particular data link-control procedure. In the case of HDLC-controlled links, maximum throughput depends on the following parameters/mechanisms: (a) transmission channel-and application-specific parameters: transmission rate; propagation and processing delays; error-characteristic of the channels; I-field lengths; and reverse channel traffic. (b) protocol-specific parameters: framing overhead; modulus value; handling of acknowledgments; error-recovery procedures; and polling mechanism.

As expected, the following results indicate that the throughput behavior is a complex function of the above parameters, and that the results can be explained only if we take into account all effects and their interaction.

Figure 2 shows that, in the region to the left of the maximum of the throughput curves, both directions of the link under NRM are almost equivalent. However, in the case of a modulus value of eight, we observe a difference in throughput between ABM and NRM for the following reason: In *ABM*, the only reason why a station does not make full use of its send channel is that the limit of M-1 unacknowledged I-frames has been reached. As soon as a station receives an acknowledgment for one of these frames, it resumes I-frame transmission. In *NRM*, the same holds true for the primary. However, the mean acknowledgment delay for its I-frames is longer as compared to ABM, because there are time intervals where the secondary has sent a frame with the F-bit set to 1, and is therefore momentarily unable to respond. This longer acknowledgment delay is the source of the throughput degradation for the primary as compared to a combined station.

For the secondary, in NRM the acknowledgment delay is roughly equal to the one obtained in ABM because the primary is always capable of acknowledging a received I-frame. However, additional intervals of wasted time are caused — compared to ABM — when a secondary has received new acknowledgments, so that the modulus rule no longer prevents it from sending further I-frames but it may previously have sent a frame with the F-bit set to 1 and not yet received a new P-bit. This additional time causes the throughput degradation for the secondary as compared to a combined station. An interesting result of our simulation is that the — *per se* different — sources of performance degradation for primary and secondary result in roughly the same throughput value for both directions.

For the modulus-8 case, the throughput results for both directions of the NRM link and for ABM are almost equivalent in the region of the optimum and within a wide range of longer message lengths.

The direction secondary to primary of the NRM link exhibits rapid throughput degradation for longer message lengths if a modulus value of 128 is used. This phenomenon can be explained as follows: if REJ recovery fails for some reason, an error is not recovered before the secondary: (i) has reached the barrier of M-1 unacknowledged I-frames; (ii) has set the F-bit; (iii) has received the next P-bit, and (iv) has performed checkpoint retransmission on the basis of this P-bit. The direction from primary to secondary exhibits an even worse behavior than the other direction. This performance degradation is also caused by an inefficient recovery procedure in case REJ recovery fails. This can occur either when a retransmitted I-frame is garbled or because a REJ has been received in error. In either situation the primary continues to transmit I-frames until it has M-1 I-frames outstanding. This, however, does not yet recover the error situation because the secondary may continue to transmit for a very long time until it also reaches the limit of M-1 unacknowledged I-frames, and therefore sets the F-bit to 1. The primary station cannot recover from this error situation before this F-bit is received. To shorten the recovery delay in this situation, we implemented the strategy that the primary stops acknowledging received I-frames when it has M-1 I-frames outstanding.

The ABM link does not show the described effect because if REJ recovery fails, time-out recovery becomes active after a relatively short delay.

The results of Figure 2 have been obtained under the assumption of constant message length. To ensure that our general results are not restricted to this particular case, we also investigated the throughput efficiency for different message-length distributions. As an example, Figure 3 shows the throughput efficiency of the same configuration as Figure 2, however, with exponentially distributed message lengths. As long as neither modulus nor transmission errors have a significant impact, the same throughput results are obtained independent of the message-length distribution. However, both effects — modulus value and transmission errors — for the following reasons lead to stronger performance degradation if message lengths are not constant: (i) The modulus rule affects the throughput over a broader range of parameters in the case of variable message lengths, because the limit of

M-1 unacknowledged I-frames is more frequently reached than in the constant-length case. (ii) Since long I-frames are more error-prone than short ones, the retransmission frequency of long frames is higher. This explains the lower throughput efficiency for the exponentially distributed message lengths in the case of longer frames. It should be noted, however, that in Figures 2 and 3, the general relation among the results for the different HDLC modes is nearly the same, independent of the message-length distribution.

Figure 4 shows throughput results for two-way flow over a satellite link. Because of the better error characteristic, $P_{bit} = 10^{-7}$ as opposed to 10^{-5} in the previous examples, we observe over the whole range considered only the effects described above for the region to the left of the maximum. Obviously, the use of a modulo-128 numbering scheme drastically improves the throughput efficiency for short I-frames.

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. An interesting performance measure in this case is the mean transfer time of the messages. Transfer time is defined as the time interval from the arrival (or generation) of a message at one station until its successful reception at the other station. This means that the transfer time includes queueing, transmission, processing and propagation times, as well as possible additional delays caused by retransmissions.

In the following examples, we consider terrestrial links, and assume symmetrical traffic conditions, i.e., that the traffic offered to both channels is equal.

Figure 5 shows for a constant message length of 5000 bit, the mean transfer time relative to the transmission time of an I-frame as a function of the useful channel load, which is defined as the ratio of throughput of information bits per unit time and the transmission rate. Two different cases are considered: error-free transmission, $P_{bit} = 0$, and channels with a bit-error rate of $P_{bit} = 10^{-5}$. In the error-free case, ABM and the direction from primary to secondary of the NRM link show the same delay values. The messages transmitted from secondary to primary suffer an additional delay due to the polling mechanism. The mean value of this additional delay is almost exactly equal to the sum t_p of processing and propagation delay. This result can be explained as follows: Under the assumed ideal conditions, the queueing delay for the ABM link and for the primary-secondary direction of the NRM link is that of a simple M/D/1 queue, whereas for the reverse direction of the NRM link the following modification of an M/D/1 queueing model is necessary: As long as the server is busy, the model operates like an M/D/1 queue. When the server becomes idle, the next opportunity to obtain service is after a fixed time interval τ has elapsed. If no customer arrives during this

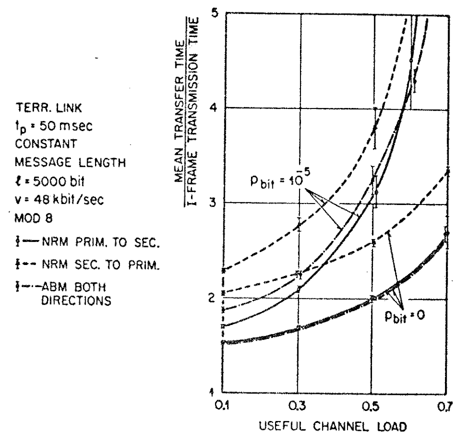

Figure 5 Mean transfer time relative to I-frame transmission time vs. useful channel load. 5000 bit constant message length.

time, a new interval τ is started etc. It can be easily shown that the mean queueing time of the modified and of the original M/D/1 queue differ by $\tau/2$. Applied to the channel from secondary to primary, the interval τ corresponds to the time interval when the secondary has sent the F-bit but has not yet received a new P-bit. This time is approximately equal to $2t_p$, which explains our simulation results.

For both links, we obtain, of course, longer transfer times if transmission errors occur on the channels. Apart from very high loads, the direction from primary to secondary of the NRM link performs slightly better than the ABM link for the following reason: In ABM, the majority of the transmission errors is recovered via REJ recovery, whereas errors occurring on the primary-to-secondary channel of the NRM link are mostly recovered via P/F-bit recovery. Because of our strategy to poll the secondary as often as possible, the latter recovery mechanism works faster under light and medium traffic. This causes the difference between both delay curves.

Whereas in Figure 5 constant length messages have been assumed, Figure 6 shows the corresponding transfer-time results in case of exponentially distributed message lengths. Obviously, this leads to longer transfer times because the queueing delay increases with the variance of the transmission times. Like in the throughput case, however, we can also observe from Figures 5 and 6 that the general relation between the performance of the two different modes is robust with respect to the message-length distribution.

Figure 7 represents the transfer-time characteristic of NRM and ABM links in a more general manner. It shows the mean transfer time relative to the I-frame transmission times as a function of the sum of processing and propagation delay. The useful channel load Y_u is held constant at 0.5. Although obtained for particular values of I-field length ℓ and transmission rate v, this diagram is also typical of other I-frame lengths and transmission rates, provided the I-frame block-error probability is 0 or 10^{-2}.

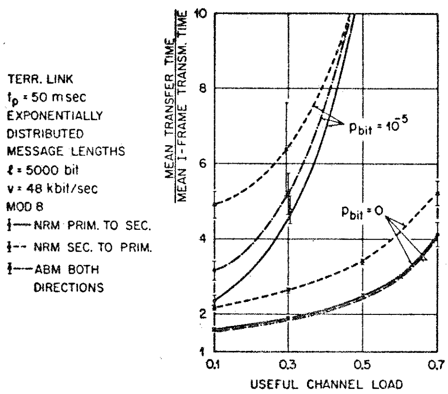

Figure 6 Mean transfer time relative to mean I-frame transmission time vs. useful channel load. Exponentially distributed message lengths (mean 5000 bit).

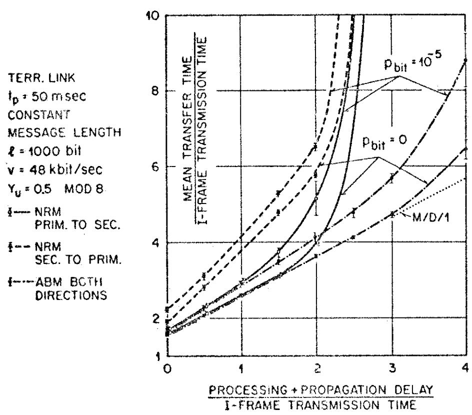

Figure 7 Mean transfer time relative to I-frame transmission time vs. processing plus propagation delay relative to I-frame transmission time.

For small values of the processing plus propagation delay, the mean transfer time grows linearly as subsequently explained: (i) In ABM and for the primary-to-secondary direction of the NRM link, the queueing delay remains constant because the useful channel load, Y_u, is held constant. For the reverse direction of the NRM link, the mean queueing time grows linearly with t_p as described above. (ii) For ABM, it has been shown analytically that the mean recovery time in case of small processing and propagation delays grows linearly with this quantity [6]. Due to the similar recovery mechanisms, it is not surprising that this result also applies to NRM. Deviation from linearity indicates the impact of the modulus. Characteristically, this impact is experienced in NRM at considerably smaller ratios of processing plus propagation delay and I-frame transmission times than in ABM.

CONCLUSIONS

The performance results presented in this paper allow the following conclusions:

(1) The link operated in ABM performs better than, or at least as well as the link under NRM. However, apart from extreme cases, there is a relatively broad range of parameters where the performance of both modes is similar.

(2) In NRM, the modulus of the sequence numbers affects performance at significantly smaller values of the processing and propagation delay than in ABM.

(3) The situations where the throughput of both modes differs significantly are: (i) In the case of long propagation and processing delays where the modulus rule in conjunction with the polling operation results in a throughput degradation of NRM due to longer and more frequent intervals of wasted time. (ii) For high block-error probabilities where in NRM a long recovery delay can occur if REJ recovery fails or is not allowed. In ABM, this error situation is efficiently resolved by time-out recovery.

(4) With respect to message transfer time, both directions of an ABM link and the direction from primary to secondary in NRM show comparable results as long as the modulus value has no impact. The direction secondary-to-primary suffers from an additional delay due to the polling operation. The mean of this additional delay is in the order of the processing plus propagation delay. Depending on the ratio of processing plus propagation delay to I-frame transmission time, the modulus value can have a much stronger delay-increasing effect in NRM than in ABM.

ACKNOWLEDGMENT

The authors would like to thank Prof. Dr. A. Lotze for his continued interest and support; D. Delestre, G. Grinenwald, and W. Schneider for writing the simulation models.

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 (independent numbering). International Standard ISO 4335.

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

[4] R.J. Cypser, "Communications architecture for distributed systems," Addison-Wesley (Reading, Massachusetts, 1978).

[5] K.C. Traynham and R.F. Steen, "SDLC and BSC on satellite links: a performance comparison," *ACM Computer Communication Review*, Vol. 7, 3-14, 1977.

[6] W. Bux, K. Kümmerle, and H.L. Truong, "Balanced HDLC procedures: A performance analysis," to appear in IEEE Trans. Commun., November 1980.