INTERNATIONAL STANDARD

ISO 7478

First edition
1987-07-01

INTERNATIONAL ORGANIZATION FOR STANDARDIZATION
ORGANISATION INTERNATIONALE DE NORMALISATION
МЕЖДУНАРОДНАЯ ОРГАНИЗАЦИЯ ПО СТАНДАРТИЗАЦИИ

Information processing systems — Data communication — Multilink procedures

Systèmes de traitement de l'information — Communication de données — Procédures multiliaison

Reference number
ISO 7478 : 1987 (E)

Foreword

ISO (the International Organization for Standardization) is a worldwide federation of national standards bodies (ISO member bodies). The work of preparing International Standards is normally carried out through ISO technical committees. Each member body interested in a subject for which a technical committee has been established has the right to be represented on that committee. International organizations, governmental and non-governmental, in liaison with ISO, also take part in the work.

Draft International Standards adopted by the technical committees are circulated to the member bodies for approval before their acceptance as International Standards by the ISO Council. They are approved in accordance with ISO procedures requiring at least 75 % approval by the member bodies voting.

International Standard ISO 7478 was prepared by Technical Committee ISO/TC 97, *Information processing systems.*

Users should note that all International Standards undergo revision from time to time and that any reference made herein to any other International Standard implies its latest edition, unless otherwise stated.

Printed in Switzerland

Information processing systems — Data communication — Multilink procedures

0 Introduction

Multilink procedures reference the layers of the ISO Open Systems Interconnection (OSI) reference model; specifically the physical, data link, and network layers. The multilink procedures (MLP) reside in the data link layer.

Multilink procedures provide the means for accepting data units from the network layer, scheduling data units for transmission and retransmission over a group of parallel data links, and reordering the received data units prior to delivering them to the network layer. Multilink procedures provide the following general features :

a) achieve economy and reliability of service by providing multiple connections between data stations;

b) permit addition and deletion of connections without interrupting the service provided by the multiple connections;

c) optimize bandwidth utilization of a group of connections through load sharing;

d) achieve graceful degradation of service when a connection(s) fails;

e) provide each multiple connection group with a single logical data link appearance to the network layer; and

f) provide, when required, resequencing of the received data units prior to delivering them to the network layer.

1 Scope and field of application

This International Standard specifies multilink procedures where a multiplicity of parallel data links at the data link layer are used to provide a variable bandwidth data link between network layer entities. The multilink procedures (MLP) exist as a new upper sublayer of the data link layer, operating between the network layer and a multiplicity of single data link protocol functions (SLPs) in the data link layer (see figure 1).

This International Standard does not specify the way in which the SLPs indicate to the MLP that the transmission of a multilink frame has successfully been completed.

These multilink procedures do not preclude the use of different single link procedures, each with differing delay characteristics and/or line speeds to form one multilink group.

When the procedures defined by this International Standard are to be used on one or more parallel data links, both ends of the data link must know that these procedures are to be used before the first multilink frame is sent. This could be achieved by a prior agreement that all communications on this data link will use these procedures, or by one of the SLPs negotiating the use of these procedures, or by some other means. The method by which both ends achieve a common understanding as to the use or non-use of these multilink procedures is not defined in this International Standard.

2 Definitions and parameters

2.1 MLP reset confirmation bit (C) : The MLP reset confirmation bit is used in reply to the R bit set to 1 to confirm that all of the MLP state variables have been reset. The C bit set to 0 is used in normal communication; i.e., no reset request has been activated. The C bit set to 1 is used to reply to the multilink frame with the R bit set to 1, and indicates that the MLP state variable reset process has been completed. In this C = 1 case, the multilink frame is used without a data unit field.

2.2 multilink procedures (MLP) : The protocols employed to transfer data over one or more SLPs which form a group.

2.3 multilink send sequence number [MN(S)] : Value of the multilink sequence number assigned to a multilink frame.

NOTE — The number assigned lies in the range of 0 to 4095 inclusive and is used to resequence, when resequencing is required, and to detect missing or duplicate multilink frames at the receiving multilink procedure data station (MLP) before data units are delivered to the network layer.

2.4 lost multilink frame timer (MT1) : During low traffic periods, the expiration of this timer provides the means to determine that the multilink frame with MN(S) equal to MV(R) is lost.

2.5 group busy timer (MT2) : Optional. At the receiving MLP, the expiration of this timer provides the means to determine that a "blocked" condition exists because buffers were exhausted before resequencing could be accomplished.

2.6 MLP reset confirmation timer (MT3) : Timer MT3 is used to provide a means of identifying that the multilink frame with the C bit set to 1 that is expected following the

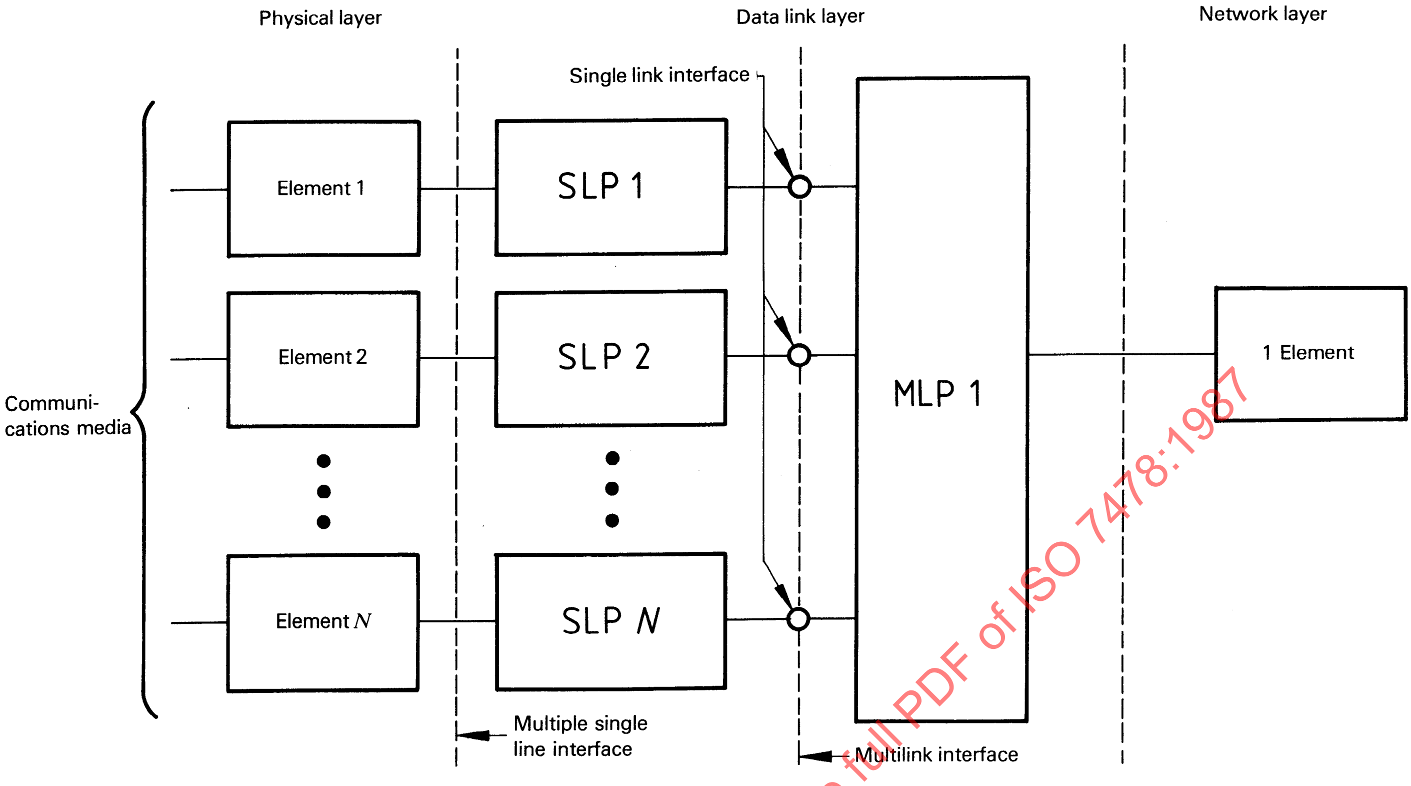

Figure 1 — Multilink functional organization

transmission of the multilink frame with the R bit set to 1 has not been received.

2.7 multilink receive state variable [MV(R)] : Denotes the MN(S) of the next in-sequence multilink frame to be transferred to the network layer.

NOTE — The MV(R) identifies the lower edge of the receive window.

2.8 multilink send state variable [MV(S)] : Denotes the next in-sequence MN(S) to be assigned to a multilink frame to be sent over this multilink group.

2.9 multilink frame acknowledgment state variable [MV(T)] : Denotes the MN(S) of the oldest multilink frame awaiting an indication of acknowledgment from the local SLP.

NOTE — The MV(T) identifies the lower edge of the transmit window.

2.10 multilink window size (MW) : For a given direction of transmission, both the transmitting MLP and the receiving MLP shall use the same value of MW.

NOTE — The transmit window contains the sequence numbers MV(T) to MV(T) + MW − 1 inclusive. The receive window contains the sequence numbers MV(R) to MV(R) + MW − 1 inclusive.[1)]

2.11 receive MLP window guard (MX) : The range of sequence numbers of fixed size, of higher value than those in the receive window, that, if received as MN(S), indicates that any multilink frames in the range MV(R) to MN(S)-received − MW that have not been received have been lost.

2.12 range of abnormal multilink frames (MZ) : The range of sequence numbers in which MN(S) should not be received during normal operation.

NOTE — Multilink frames received in range MZ should be discarded.

2.13 number of SLP retransmission attempts (N) : The number of times that a single link procedure data station (SLP) attempts retransmission of a multilink frame before notifying the MLP of the situation.

NOTE — Subsequent SLP action should be dependent upon SLP design. Subsequent MLP action should involve assigning the multilink frame to the same or one or more other SLPs for transmission.

1) MW is a system parameter which can never exceed 4095 − MX. Factors which will affect the value of parameter MW include, but are not limited to, link transmission and propagation delays, the number of links, the range of multilink frame lengths, and single link parameters (N retransmission attempts, response times, and outstanding number of unacknowledged I frames).

2.14 MLP reset request bit (R) : The MLP reset request bit is used to request a reset of the MLP state variables. The R bit set to 0 is used in normal communication; i.e., no request for a multilink reset. The R bit set to 1 is used to request the reset of the receiving MLP state variables.

In the R = 1 case, the multilink frame data unit field does not contain higher layer information, but may contain an optional Cause Field that incorporates the reason for the reset.

NOTE — For the DTE-DCE X.25 interface application[1)], a Cause Field of 8 bits is specified. The encoding of that 8-bit field is a subject for further study.

2.15 sequence check option bit (S) : The S bit is only significant when V = 1 (indicating that sequencing of received multilink frames shall not be required). S = 1 shall mean no MN(S) number has been assigned. S = 0 shall mean an MN(S) number has been assigned so, although sequencing shall not be required, a duplicate multilink frame check may be made, as well as a missing multilink frame identified.

2.16 single link procedures (SLP) : The data link protocols employed to establish, maintain, transfer data, and terminate a logical data link over a single data circuit.

2.17 void sequencing bit (V) : The V bit indicates if a received multilink frame shall be subjected to sequencing constraints. V = 1 means sequencing shall not be required. V = 0 means sequencing shall be required.

3 Multilink frame format

To permit resequencing of data units (for example, packets), a multilink control (MLC) field shall be provided. This field shall be transmitted as the first two octets in the information field of an SLP transmission unit. Figure 2 shows the MLC field and its relationship to the data unit and the SLP header/trailer.

NOTE — Some format modification may be required if code dependent procedures are used by the SLP.

The two octet MLC field shall immediately follow the SLP header and shall extend the data link layer envelope used for transporting data units between network layer entities. The single link and multilink control fields shall be generated and used at the data link layer only and shall not be forwarded to the network layer.

The multilink sequence number MN(S) shall consist of twelve bits, providing a modulo count of 4096. MN(S) shall be split into two fields, MNH(S) and MNL(S), as shown in figure 2. Four control bits shall be provided in the MLC field for use in providing additional multilink control features.

The void sequencing bit V shall be used to indicate when resequencing of received data units into the sequence received from the network layer shall not be required prior to release to the network layer. When V = 1, the data unit shall not have to be resequenced and the receiving MLP shall immediately deliver the data unit to the network layer, regardless of its order in the received sequence. When V = 0, the data units shall be kept in sequence prior to release to the network layer. V = 0 shall indicate an MN(S) number is present.

The sequence check option bit S shall have significance only when V = 1 (indicating that resequencing of multilink frames is not required). The S bit shall have no significance when V = 0. When V = 1 and S = 1, the transmitting MLP shall not have assigned an MN(S) value to the multilink frame. Consequently, the data unit contained in that multilink frame shall be forwarded to the network layer without checking for sequence integrity or duplicates. When V = 1 and S = 0, the transmitting MLP shall have assigned an MN(S) value to the multilink frame to facilitate checking for duplicate or missing multilink frames, even though sequence integrity is not observed. Duplicate multilink frames shall not have their data units delivered to the network layer.

The MLP reset request bit R shall be used to initiate the multilink resetting procedure. The MLP reset confirmation bit C shall be used to indicate completion of the multilink resetting procedure. During normal data interchange, the MLP reset request bit and the MLP reset confirmation bit shall be set to 0.

4 Transmitter operation

4.1 General

The transmitting MLP shall be responsible for controlling the flow of data units from the higher layer (for example, the network layer) into multilink frames and then to the SLPs for transmission to the remote receiving MLP.

The functions of the transmitting MLP shall be to

a) accept data units from the network layer;

b) allocate multilink control fields, containing the appropriate sequence number MN(S), to the data units;

c) assure that MN(S) is not assigned outside the MLP transmit window (MW);

d) pass the resultant multilink frames to the SLPs for transmission;

e) accept indications of successful transmission acknowledgments from the SLPs;

f) monitor and recover from transmission failures or difficulties that occur at the SLP sublayer; and

g) accept flow control indications from the SLPs and take appropriate actions.

4.2 Transmission of multilink frames

When the transmitting MLP accepts a data unit from the network layer, it shall place the data unit in a multilink frame,

1) See ISO 7776, *Information processing systems — Data communication — High-level data link control procedures — Description of the X.25 LAPB-compatible DTE data link procedures.*

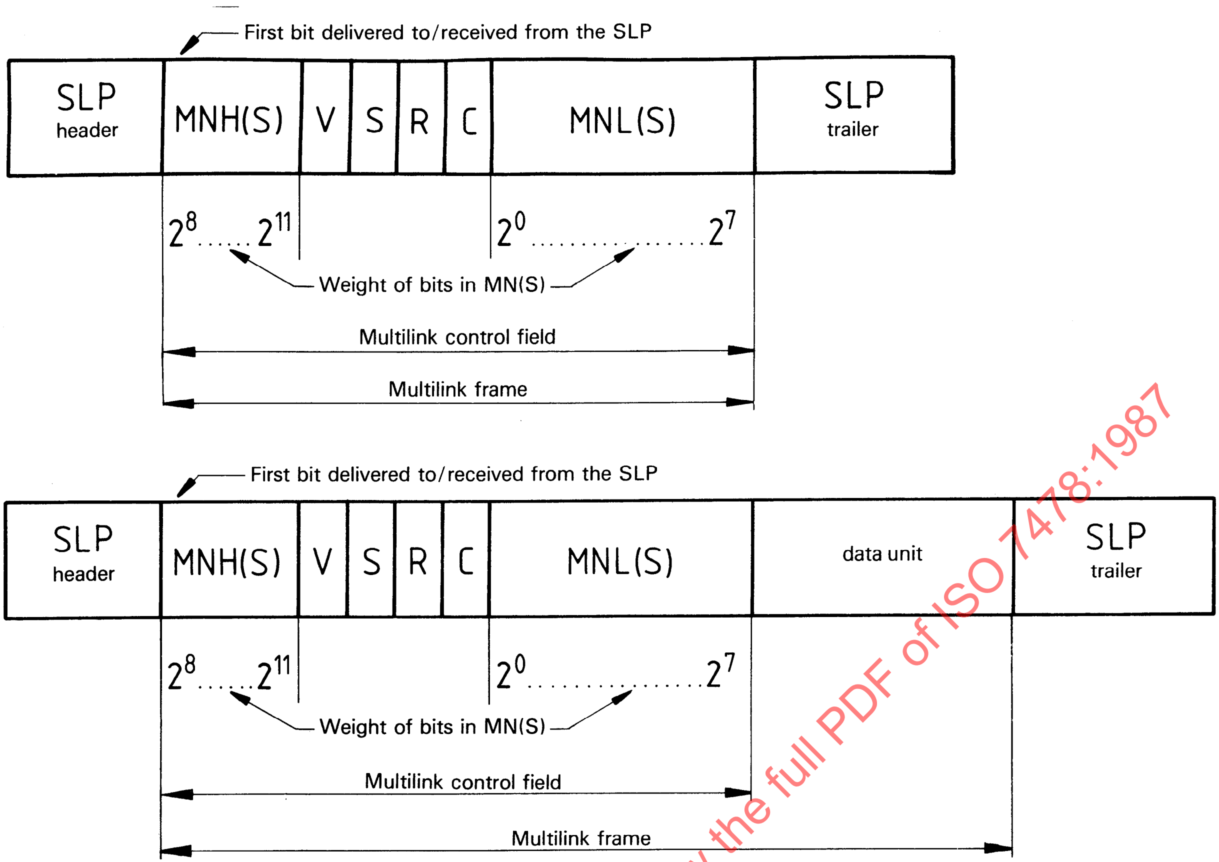

Key

MNH(S) = bits 9 to 12 of 12-bit multilink send sequence number MN(S)

MNL(S) = bits 1 to 8 of 12-bit multilink send sequence number MN(S)

V = void sequencing bit

S = sequence check option bit

R = MLP reset request bit

C = MLP reset confirmation bit

SLP = single link procedure

Figure 2 — Multilink frame formats

and, when required (V = 0, or V = 1 and S = 0), shall set the MN(S) = MV(S) and increment MV(S) by one.

In the following, incrementing send and receive state variables is in reference to a continuously repeated sequence series; i.e., 4095 is 1 higher than 4094 and 0 is 1 higher than 4095 for modulo 4096 series.

If the MN(S) is less than MV(T) + MW, and the remote data station has not indicated a busy condition on all available data links, the transmitting MLP shall assign the lowest numbered MN(S) unassigned multilink frame to an available SLP.

When the SLP successfully completes the transmission of a multilink frame(s) by receiving an acknowledgment from the remote SLP, it shall indicate this to the transmitting MLP. The transmitting MLP may then delete the acknowledged multilink frame(s). As the transmitting MLP receives indications of acknowledgments from the SLPs, MV(T) shall be advanced to denote the lowest numbered multilink frame not yet acknowledged.

The transmitting MLP shall always assign the lowest MN(S) multilink frame first. Also, the transmitting MLP may assign a multilink frame to more than one SLP.

If a multilink frame is transmitted on more than one data link (for example, to increase the probability of successful delivery), there is a possibility that one of these multilink frames (i.e., a duplicate) may be delivered to the remote MLP after an earlier one has been acknowledged. The earlier multilink frame would have resulted in the receiving MLP incrementing its MV(R) and the transmitting MLP incrementing its MV(T). To ensure that an old duplicate multilink frame is not mistaken for a new frame by the receiving MLP, the transmitting MLP shall never send a new multilink frame with an MN(S) equal to or greater than MN(S)'—MW—MX, where MN(S)' is associated with a duplicate multilink frame that is being transmitted on other SLPs, until all SLPs have either successfully transferred the multilink frame or retransmitted the frame their maximum number of times. Alternatively, the incrementing of MV(T) may be withheld until all SLPs have either successfully transferred the multilink frame or retransmitted the frame their maximum number of times.

4.3 Transmitter flow control

Flow control shall be achieved by the window size parameter MW, and through busy conditions being indicated by the remote SLPs.

The transmitting MLP shall not assign a multilink frame with an MN(S) greater than MV(T) + MW − 1. At the point when the next multilink frame to be assigned has an MN(S) = MV(T) + MW, the transmitting MLP shall hold this and subsequent multilink frames until an indication of acknowledgment advancing MV(T) is received.

In figure 3, MV(S) has become equal to MV(T) + MW. At this point the transmitting MLP shall not assign any further multilink frames until MV(T) advances.

The receiving MLP shall exercise flow control of the transmitting MLP by indicating a busy condition over one or more SLPs. The number of SLPs made busy determines the degree of transmitter flow control realized. When the transmitting MLP receives an indication of a busy condition from one or more of its SLPs, it shall reassign any unacknowledged multilink frames that were assigned to those SLPs. It shall assign the multilink frame containing the lowest MN(S) to an available SLP as specified above.

4.4 Retransmission

When a transmitting SLP does not receive an acknowledgment for a frame containing a multilink frame after N retransmissions, the transmitting MLP shall reassign this unacknowledged multilink frame for transmission by the same or other SLP(s), unless an acknowledgment from another SLP has been indicated for this multilink frame.

4.5 Link failure

If the transmitting SLP detects data link failure, the SLP shall indicate to the transmitting MLP that the data link is out of service. All unacknowledged multilink frames shall be reassigned by the transmitting MLP. The transmitting MLP shall not assign any further multilink frames to this SLP until the data link has been re-established.

4.6 Transmitter reconfiguration

A data link may be taken out of service by disconnection at the physical layer or the data link layer. The transmitting SLP shall indicate to the transmitting MLP that the data link is out of service. All unacknowledged multilink frames that are assigned to that SLP shall be reassigned by the transmitting MLP.

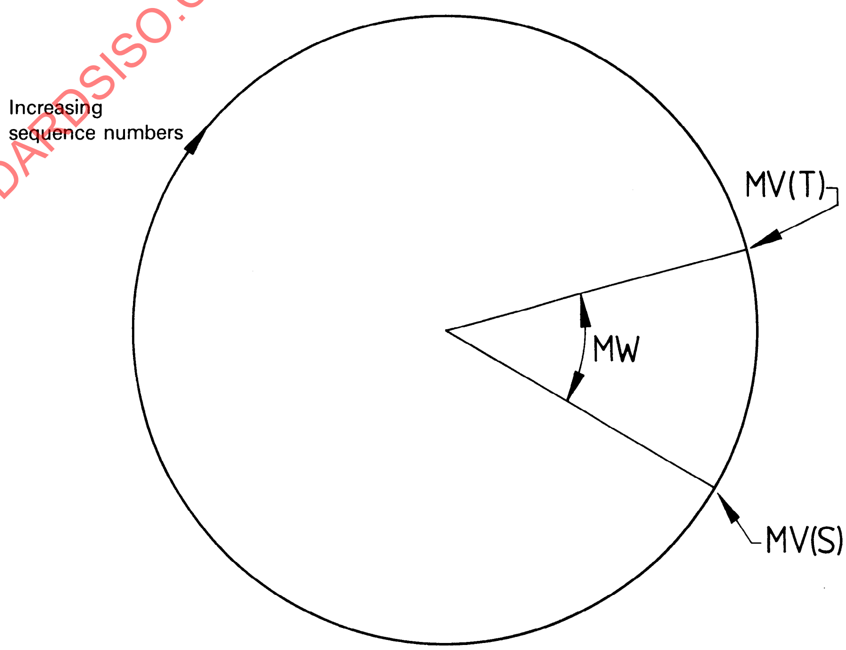

Figure 3 — Transmitter flow control

A data link returned to service, or put in service for the first time, shall be indicated to the transmitting MLP by the SLP. The transmitting MLP may then assign multilink frames to this SLP.

5 Receiver operation

5.1 General

As multilink frames are received from the receiving SLP, the contents of the multilink control field shall be examined. When V = 0, the multilink frame shall be resequenced before the data unit is delivered to the network layer, as described below. In the unlikely event a transmission error occurs which is undetected by the receiving SLP, or an internal machine error occurs, a multilink frame may be lost. The receiving MLP can detect such lost multilink frames as described below. Higher level recovery mechanisms are assumed to be capable of recovering from such conditions. Should the receiving MLP begin to run out of buffering resources, a method of flow control described below may be used.

5.2 Reception of multilink frames

Any multilink frame less than two octets in length shall be discarded by the receiving MLP.

If V = 0, the multilink frame shall be resequenced as described below. If V = 1 and S = 0, the data unit contained in the multilink frame shall be immediately delivered to the network layer if MV(R) is less than or equal to MN(S) and MN(S) is less than or equal to MV(R) + MW + MX − 1 and is not a duplicate MN(S). If V = 1 and S = 1, the data unit contained in the multilink frame shall be immediately delivered to the network layer.

The receiving MLP sequence numbers used to control the reception of multilink frames are divided into three regions as shown in figure 4.

MV(R) is the MN(S) of the next in-sequence multilink frame whose data unit is to be passed to the network layer. Multilink frames with higher MN(S) may already have been received and may be held waiting for the expected multilink frame before being passed to the network layer. The sequence numbers of the three regions relative to MV(R) shall be as follows :

a) the receive window MW contains the sequence numbers MV(R) to MV(R) + MW − 1 inclusive;

b) the guard region MX contains the sequence numbers MV(R) + MW to MV(R) + MW + MX − 1 inclusive;

c) the abnormal region MZ contains the sequence numbers MV(R) + MW + MX to MV(R) − 1 inclusive.

When a multilink frame with MN(S) = MV(R) has been received, whether it has been held according to 5.3 or not, its data unit shall be delivered to the higher layer (for example, the network layer), and the receiving MLP shall increment MV(R).

5.3 Resequencing

The receiving MLP shall resequence multilink frames by holding all multilink frames requiring resequencing and with MN(S) in the range MV(R) + 1 to MV(R) + MW − 1 prior to delivering their data units to the network layer. All multilink frames with sequence numbers less than MV(R) have been resequenced, if necessary, and their data units delivered to the network layer.

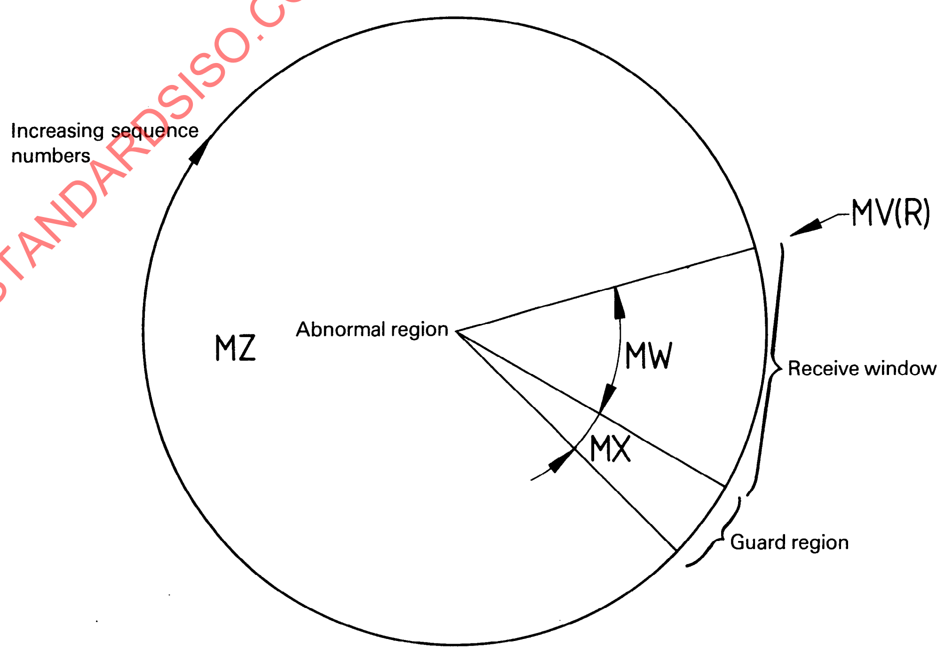

Figure 4 — Reception of multilink frames

5.4 Detecting lost multilink frames

Because

— MW is a fixed parameter known to both the receiving MLP and the transmitting MLP; and

— the transmitting MLP's window starts at the smallest (modulo 4096) multilink sequence number whose acknowledgment has not yet been indicated by an SLP,

the receiving MLP should not receive multilink frames, other than duplicate multilink frames, with MN(S) outside the receive window.

When the receiving SLP accepts and acknowledges a transmission unit containing a multilink frame and the transmitting SLP indicates this acknowledgement to the transmitting MLP, the transmitting MLP may advance its window, permitting it to assign multilink frames with higher MN(S)s to SLPs. If the received multilink frame is subsequently lost, the receiving MLP's receive window will not be advanced. The receiving MLP shall wait for the missing multilink frame until

a) it receives a multilink frame with MN(S) in the guard region; or

b) timer MT1 expires.

When a lost multilink frame is detected, the network layer may be notified.

When a multilink frame with MN(S) in the guard region is received, it shall be considered a valid multilink frame. Multilink frames MV(R) to MN(S)-received minus MW which have not been received shall be declared lost. Multilink frames MV(R) to MN(S)-received minus MW which have been received shall be passed to the network layer. The receive window shall be rotated such that MV(R) shall be MN(S)-received minus MW + 1. Normal handling of multilink frames, as defined in 5.2, shall now be resumed (see figure 5).

The guard region shall be large enough for the receiving MLP to recognize the highest MN(S) outside of its receive window that it may legitimately receive after a multilink frame loss has occurred. MX should be no larger than MW and, for example :

a) in a system where the transmitting MLP assigns h_i in-sequence, contiguous multilink frames at a time to the ith SLP, then MX should be greater than or equal to the sum of all the $h_i + 1 - h_{min}$ where h_{min} equals the smallest h_i encountered. When there are L SLPs in the multilink group, MX should be greater than or equal to

$$\sum_{i=1}^{L} h_i + 1 - h_{min};\ \text{or}$$

b) in a system where the transmitting MLP assigns on a rotation basis h in-sequence, contiguous multilink frames at a time to each SLP, MX at the receiving MLP should be greater than or equal to $h(L-1) + 1$, where L is the number of SLPs in the multilink group.

When the timer MT1 runs out, the multilink frame with MN(S) = MV(R) is declared lost. The receive window is rotated, as before. Timer MT1 shall be started when a multilink frame with MN(S) higher than MV(R) and requiring resequencing is received. When a multilink frame with MN(S) = MV(R) is received, and additional multilink frames are awaiting delivery pending receipt of a multilink frame with MN(S) equal to the new MV(R), timer MT1 shall be restarted; otherwise, timer MT1 shall be reset. Timer MT1 shall be disabled when all SLPs are indicating a busy condition.

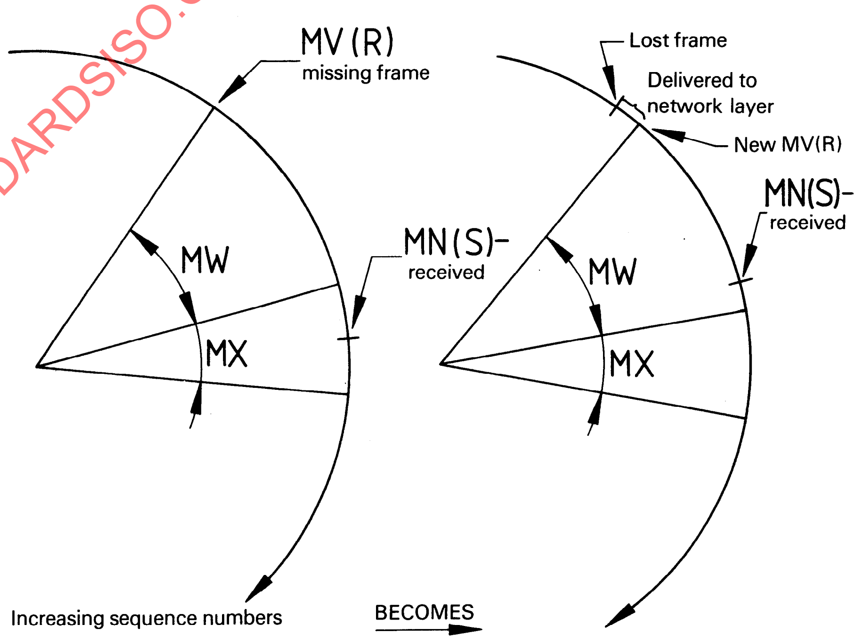

Figure 5 — Detecting lost multilink frames

5.5 Receiver flow control

When flow control of the transmitting MLP is desired, one or more receiving SLPs may indicate a busy condition.

If the receiving MLP can exhaust its buffer capacity before resequencing can be completed, timer MT2 may be implemented. Whenever a busy condition is indicated by the receiving MLP on all SLPs and multilink frames at the receiving MLP are awaiting resequencing, timer MT2 shall be started. When the busy condition is cleared on one or more of the SLPs, timer MT2 shall be reset.

Should timer MT2 run out, the multilink frame with MN(S) = MV(R) is blocked and shall be considered lost. MV(R) shall be incremented to the next sequence number not yet received, and data units contained in multilink frames with intervening multilink sequence numbers are delivered to the network layer. Timer MT2 shall be restarted if the busy condition remains in effect on all SLPs and more multilink frames are awaiting resequencing.

The period of timer MT2 may be independently established by the implementer of each MLP.

6 Initializing the MLP

MV(S), MV(T) and MV(R) shall be initialized to zero by applying the resetting procedure described in clause 7.

7 Resetting the MLP

The multilink resetting procedure provides the mechanism for synchronizing the multilink sequence numbers for the interconnected MLPs. Following a successful multilink resetting procedure, the multilink sequence numbering in each direction begins with the value 0. A multilink frame with R = 1 is used to request multilink reset, and a multilink frame with C = 1 confirms that the multilink reset process has been completed. An MLP resets MV(S) and MV(T) to zero on transfer of a multilink frame with R = 1; and resets the MV(R) to zero on receipt of a multilink frame with R = 1.

When an MLP initiates the resetting procedure, it removes all of the unacknowledged multilink frames that are held in that MLP and its associated SLPs, and retains control of those frames. Hereafter, the initiating MLP does not transfer a multilink frame with R = C = 0 until the reset process is completed. (One method to remove multilink frames in the SLP is to disconnect the data link of the SLP.) The initiating MLP then resets its multilink send state variable MV(S) and its transmitted multilink frame acknowledged state variable MV(T) to zero. The initiating MLP then transfers a multilink frame with R = 1 as a reset request to one of its SLPs and starts timer MT3. The value of the MN(S) field in the R = 1 frame may be any value, since when R = 1 the MN(S) field is ignored by the receiving MLP. The initiating MLP continues to receive and process multilink frames from the remote MLP, in accordance with the procedures as described in 5.2, until it receives a multilink frame with R = 1 from the remote MLP.

An MLP which has received a multilink frame with R = 1 (reset request) in the normal communication status from an initiating MLP starts the operation as described above; that MLP should receive no multilink frame with R = C = 0 from the other MLP until the reset process is completed. Any such multilink frame received is discarded. When an MLP has already initiated its own multilink resetting procedure and has transferred the multilink frame with R = 1 to one of its SLPs for transmission, that MLP does not repeat the above operation upon receipt of a multilink frame with R = 1 from the other MLP.

Receipt of a multilink frame with R = 1 (reset request) causes the receiving MLP to deliver to the higher level those multilink frame data units already received and to identify those frames assigned to SLP's but unacknowledged.

The higher level may be informed of the multilink frame loss at the original value of MV(R) and at any subsequent value(s) of MV(R) for which there has been no multilink frame received up to and including the highest numbered multilink frame received. The receiving MLP then resets its multilink receive state variable MV(R) to zero.

After an MLP assigns a multilink frame with R = 1 to one of its SLPs, it shall receive indication of successful or unsuccessful transmission from that SLP as one of the conditions before transferring a multilink frame with C = 1; when the initiating MLP then receives a multilink frame with R = 1, and has completed the multilink state variable resetting operation described above, the initiating MLP transfers a multilink frame with C = 1 (reset confirmation) to the other MLP. When an MLP has

a) received a multilink frame with R = 1;

b) transferred a multilink frame with R = 1; and

c) completed the multilink state variable resetting operation described above,

that MLP then transfers a multilink frame with C = 1 (reset confirmation) to the other MLP as soon as possible given that indication of the successful or unsuccessful transmission of the R = 1 multilink frame has been received from that MLP's SLP. The C = 1 multilink frame is a reply to the multilink frame with R = 1. The value of the MN(S) field in the above C = 1 frame may be any value, since when C = 1 the MN(S) field is ignored by the receiving MLP. The multilink sequence number MN(S) received in each direction following multilink reset will begin with the value zero.

When an MLP uses only one SLP to transmit the multilink frame with R = 1 and the multilink frame with C = 1, the MLP can transfer the multilink frame with C = 1 immediately after the multilink frame with R = 1, without waiting for SLP indication of transmission completion. An MLP shall not retransmit a multilink frame with R = 1 or a multilink frame with C = 1, unless timer MT3 runs out. An MLP may use two different SLPs as long as one is used for transmitting the multilink frame with R = 1 and the other is used for transmitting the multilink frame with C = 1 following receipt of the SLP indication of successful or unsuccessful transmission of the R = 1 multilink frame. A multilink frame with R = C = 1 is never used.

When an MLP receives the multilink frame with C = 1, the MLP stops its timer MT3. The transmission of a multilink frame with C = 1 to a remote SLP and the reception of a multilink frame with C = 1 from the remote MLP completes the multilink resetting procedure for an MLP. The first multilink frame transferred with R = C = 0 shall have a multilink sequence number MN(S) value of zero. After an MLP transfers a multilink frame with C = 1 to a SLP, the MLP may receive one or more multilink frames with R = C = 0. After an MLP receives a multilink frame with C = 1, the MLP may transfer one or more multilink frames with R = C =0.

When an MLP additionally receives one or more multilink frames with R = 1 between receiving a multilink frame with R = 1 and transferring multilink frames with C = 1, the MLP shall discard the extra multilink frames with R = 1. When an MLP receives a multilink frame with C = 1, which is not a reply to a multilink frame with R = 1, the MLP shall discard the multilink frame with C = 1.

After an MLP transfers a multilink frame with C = 1 on one of its SLPs, the MLP may receive a multilink frame with R = 1 from the other MLP. The MLP shall regard the multilink frame with R = 1 as a new reset request and shall start the multilink resetting procedure from the beginning. When an MLP, which has not received a multilink frame with R = 1, transfers a multilink frame with R = 1, and therefore receives a multilink frame with C = 1, the MLP shall restart the resetting procedure from the beginning.

When timer MT3 runs out, the MLP restarts the multilink resetting procedure from the beginning. The value of timer MT3 shall be large enough to include the transmission, retransmission, and propagation delays in the SLPs, and the operation time of the MLP that receives a multilink frame with R = 1 and responds with a multilink frame with C = 1.

Annex

Examples of multilink resetting procedures

(This annex does not form part of the standard.)

A.1 Introduction

The following examples illustrate application of the multilink resetting procedures in the cases of :

a) MLP reset initiated by a single station; and

b) MLP reset initiated by both stations simultaneously.

A.2 MLP reset initiated by a single station

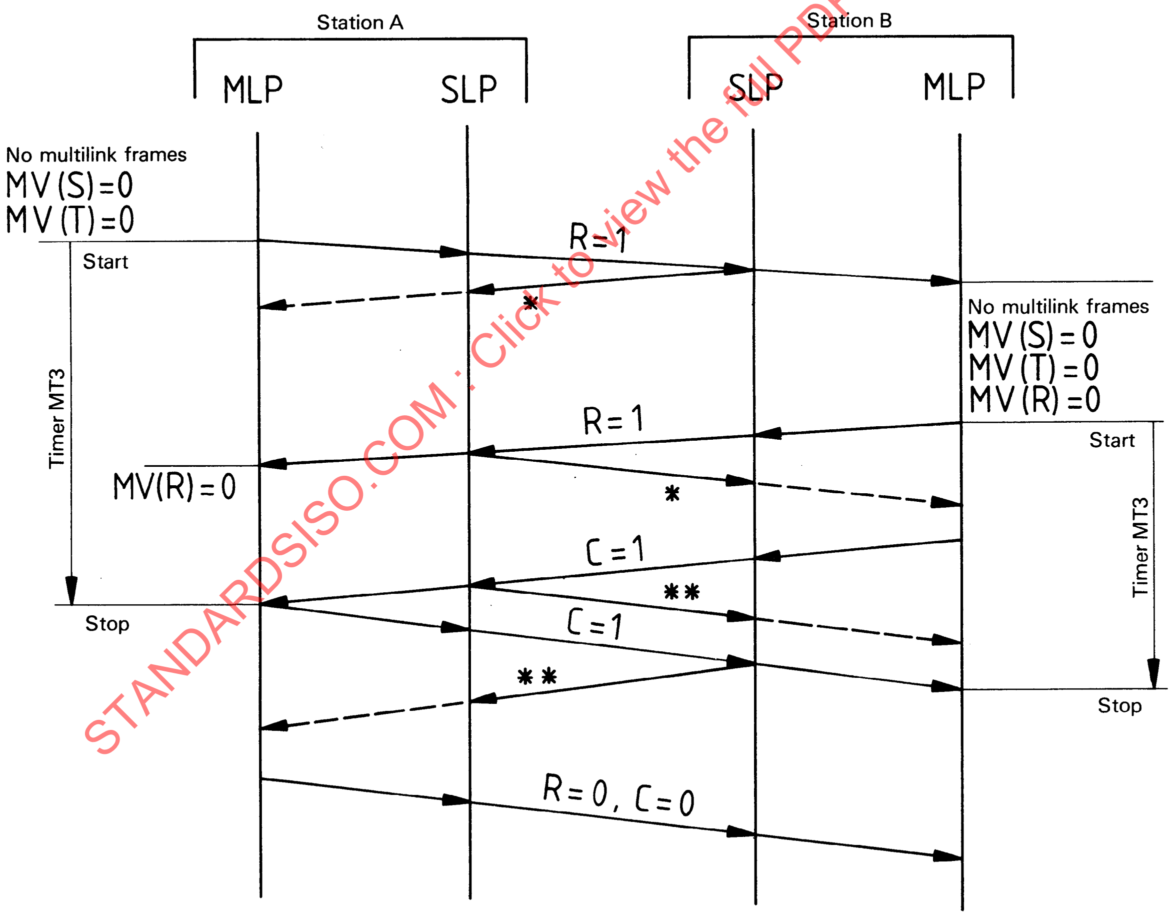

* The SLP frame that acknowledges delivery of the multilink frame with R = 1.

** The SLP frame that acknowledges delivery of the multilink frame with C = 1.