ECMA

EUROPEAN COMPUTER MANUFACTURERS ASSOCIATION

STANDARD ECMA-49

HDLC
ELEMENTS OF PROCEDURE

2nd Edition-August 1979

Free copies of this document are available from ECMA,
European Computer Manufacturers Association
114 Rue du Rhône - 1204 Geneva (Switzerland)

ECMA

EUROPEAN COMPUTER MANUFACTURERS ASSOCIATION

STANDARD ECMA-49

HDLC
ELEMENTS OF PROCEDURE

2nd Edition-August 1979

BRIEF HISTORY

In June 1967 ECMA TC9 received a proposal for a character-oriented two-way simultaneous Data Link Control procedure which was intended to improve the performance of existing procedures specifically for point-to-point data links.

In March 1970 ECMA TC9 received a proposal for a bit-oriented Data Link Control procedure which could be applied for point-to-point data links as well as for multipoint data links. It had a numbering system which was known as "single numbering".

In September 1971 ECMA TC9 received another proposal which had a numbering system referred to as "double numbering" where information transfer in opposite directions is using independent numbers and which was believed to be applicable to a very wide range of applications. To avoid substantial delay ECMA TC9 decided to split the standardization of data link procedures in two steps:

Frame Structure, and
Elements of Procedure.

A separate Standard for High Level Data Link Control Frame Structure was adopted by the General Assembly of ECMA on December 13, 1973 (ECMA-40). A second edition of this Standard has been issued in September 1976.

In May 1974 ECMA TC9 supported the elements of High Level Data Link Control Procedure with double (independent) numbering in producing a working paper to ISO/TC97/SC6. It was then also recognized and agreed that further details of the procedure should be dealt with in additional Standards called "Classes of Procedure".

In September 1976 ECMA TC9 produced Standard ECMA-49 on Elements of Procedure. Following work in ISO and CCITT and the preparation of the first Standards on Classes of Procedure a Second Edition was prepared and presented to the General Assembly of ECMA in June 1979.

Adopted as 2nd Edition of Standard ECMA-49 by the General Assembly on June 21, 1979.

TABLE OF CONTENTS

1. SCOPE

This Standard ECMA-49 defines in detail the Elements of Data Link Control Procedure for bit sequence independent data transmission using the Standard ECMA-40, Frame Structure, and frame numbering in both directions.

These HDLC Elements of Procedure are defined specifically in terms of the actions that occur on receipt of commands at a Secondary or Combined.

This Standard is intended to cover a wide range of applications, for instance, one way, two way alternate or two way simultaneous information transfer between stations which are usually buffered, including operations on different types of data circuits, e.g. multipoint/point-to-point configurations, duplex/half duplex facilities, switched/non-switched services.

In HDLC procedures, the normal cycle of the code transparent data communication between two data stations consists of the transfer of frames containing information from the data source to the data sink, acknowledged by frames in the opposite direction. Until the station receives an acknowledgement, it must hold the original information in memory in case a retransmission is necessary.

The basic configurations of stations are specified indicating how Secondaries are responsible to a Primary and how Combined's are responsible to each other.

Implementation of a particular system will normally require only a limited selection of commands and responses. Equipment designed for particular systems may not implement all the commands and responses defined in this Standard.

2. REFERENCES

ECMA-40 HDLC Frame Structure

ECMA-60 HDLC Unbalanced Class of Procedure

ECMA-61 HDLC Balanced Class of Procedure

3. DEFINITIONS

For the purpose of this Standard the following terms have the meaning indicated.

3.1 Data Link

An assembly of two or more terminal installations and the interconnecting line operating according to a particular method that permits information to be exchanged; in this context the term "terminal installation" does not include the data source and the data sink.

3.2 Station

An association of logical elements from which or to which information is transferred over the data link, including those elements which control the flow by link control procedures. It does not include data source(s) and data sink(s).

3.3 Primary

That part of a station that supports the primary control functions of the data link. It is permanently responsible for the control of the data link. The Primary generates commands for transmission and interprets the responses which are received.

3.4 Secondary

That part of a station that executes data link control functions as instructed by the Primary. A Secondary interprets received commands and generates responses for transmission.

3.5 Combined

That part of a station that includes the functions of both a Primary and a Secondary. It generates and interprets both commands and responses.

NOTE 1: In the remainder of this Standard the terms "Primary" and "Secondary" also refer to the primary and secondary functions of the Combined, unless otherwise stated.

3.6 Command

An instruction represented in the control field of a frame and transmitted by the Primary to the Secondary identified in the address field. It causes the addressed Secondary to execute a specific data link control function.

3.7 Response

A reply represented in the control field of a frame transmitted to the Primary by the Secondary identified in the address field. It reports to the Primary the action taken by the Secondary on one or more commands.

3.8 Exception

A condition in a station caused by an event it cannot handle in the normal way. This may be due to transmission errors or a station malfunction.

3.9 Sequence Exception

I frames are sequentially numbered for the purposes of error detection. A Sequence Exception is a condition in a station caused by the reception of an I frame whose sequence number is not the expected one.

4. BASIC CONFIGURATION

There are two basic configurations where the elements of procedure apply:

- Unbalanced configuration which has the Primary and one or more Secondaries,
- Balanced configuration which has two Combined.

4.1 Unbalanced Configuration

An Unbalanced Configuration has one Primary and one or more Secondaries. The configurations may be point-to-point or multipoint. The transfer may be two way alternate or two way simultaneous. The service may be switched or non-switched. The Primary is responsible for setting each Secondary in an operational mode as appropriate.
Secondaries are addressed and identified by unique addresses in the related commands and responses.

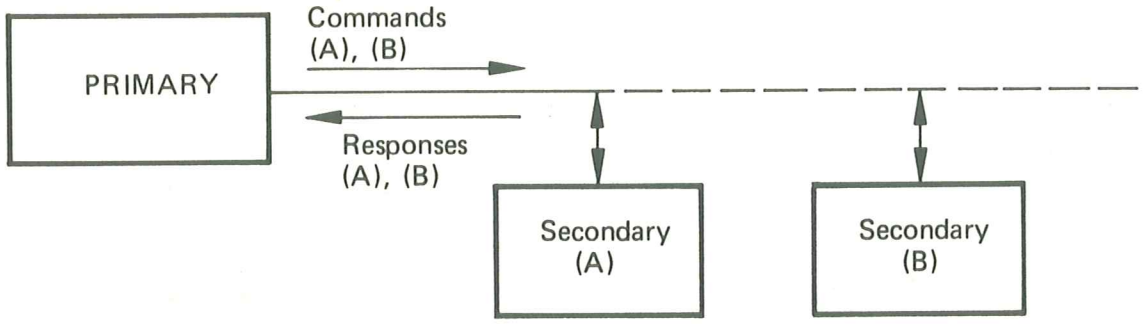

Unbalanced Configuration

4.2 Balanced Configuration

A Balanced Configuration has two Combined. The configuration is point-to-point. The transfer may be two way alternate or two way simultaneous. The service may be switched or non-switched. Both Combined have equal link control capabilities.

Each Combined has a specific address. Commands are sent with the address of the receiving Combined, Responses are sent with the address of the transmitting Combined.

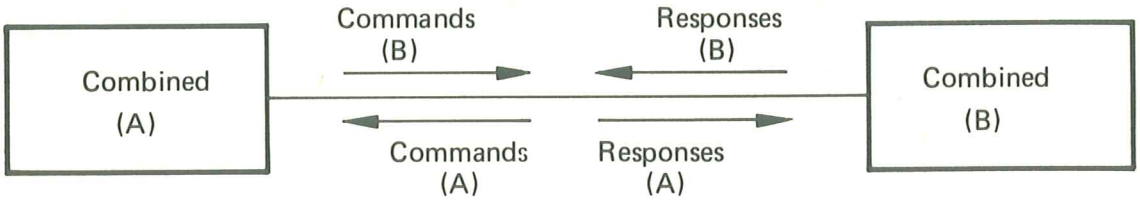

Balanced Configuration

5. MODES

5.1 Operational Modes

5.1.1 Normal Response Modes (NRM, NRME)

NRM is a data link operational mode in an unbalanced configuration in which the Secondary may initiate transmission only as the result of receiving explicit permission to do so from the Primary. After receiving permission, the Secondary shall initiate a reply transmission. This is the Secondary opportunity to respond.

The reply transmission may consist of one or more frames while maintaining the active data channel state. The last frame of the reply transmission will be explicitly indicated by the Secondary. Following the indication of the last frame the Secondary will stop transmitting until explicit permission is again received from the Primary.

NRME is an operational mode identical to the normal response mode, except that the control field in all formats will be extended by the addition of a second contiguous octet (see 6.2).

5.1.2 Asynchronous Response Modes (ARM, ARME)

ARM is a data link operational mode in an unbalanced configuration in which the Secondary may initiate transmission without receiving explicit permission from the Primary. Such an asynchronous transmission may consist of one or more frames and is used for information field transfer and/or to indicate status changes in the Secondary (e.g.: the number of the expected frame, transition from ready to busy or vice versa, occurrence of an exception).

A Secondary operating in ARM can transmit only when it detects an idle data channel state for two way alternate transfer, or at any time for two way simultaneous transfer.

ARME is an operational mode identical to the asynchronous response mode, except that the control field in all formats will be extended by a second contiguous octet (see 6.2).

5.1.3 Asynchronous Balanced Modes (ABM, ABME)

ABM is a data link operational mode in a balanced configuration in which either Combined may send commands at any time and may initiate response frame transmission without receiving explicit permission from the other Combined. Such an asynchronous transmission may consist of one or more frames and is used for information field transfer and/or to indicate status changes in the Combined (for example, the number of the expected frame, transition from ready to busy or vice versa, occurrence of an exception).

ABMEis an operational mode identical to asynchronous balanced mode, except that the control field in all formats will be extended by a second contiguous octet (see 6.2).

5.2 Disconnected Modes

Disconnected modes are provided to prevent a Secondary from appearing on the data link in an operational mode.
These modes differ from the operational modes in that the Secondary is logically disconnected from the data link, i.e. there is no information transfer.

Control field extension does not apply to the disconnected modes.

Further details on the disconnected modes can be found in 8.1.

5.2.1 Normal Disconnected Mode (NDM)

NDM is an unbalanced data link non-operational mode in which the Secondary is logically disconnected from the data link and is therefore not permitted to transmit or accept information. The Secondary has some limited capability to respond but only as a result of receiving an explicit permission to do so.

5.2.2 Asynchronous Disconnected Mode (ADM)

ADM is a data link non-operational mode, in which the Secondary is logically disconnected from the data link and is therefore not permitted to transmit or accept information. The Secondary has limited asynchronous response mode capability. It may respond in two way alternate exchange upon detection of an idle data channel state, and in two way simultaneous exchange at any time.

5.3 Initialization Mode

IM is a data link non-operational mode in which the remote Secondary link control function may be initialized by local Primary action.

Control field extension does not apply to initialization mode.

Further details on initialization modes can be found in 8.2.

6. CONTROL FIELD

6.1 Control Field Formats

The three formats defined for the control field are used to perform:

- Numbered information transfer
- Numbered supervisory functions,
- Unnumbered control functions.

The control field structure for the three formats is defined below. The least significant bits of the 3 bit sequence numbers are bits 2 and 6. Bit 1 is transmitted first.

	CONTROL FIELD BITS							
BIT ORDER	1	2	3	4	5	6	7	8
I frame format	0	N(S)			P/F	N(R)		
S frame format	1	0	S	S	P/F	N(R)		
U frame format	1	1	M	M	P/F	M	M	M

LEGEND

N(S) = Send Sequence Number

N(R) = Receive Sequence Number

P/F = Poll/Final Bit

S = Supervisory Bits

M = Modifier Bits

6.1.1 Information Frame - I frame

The I frame is used for information transfer. Unless otherwise specified, it is the only format which may contain an information field. The functions of N(S), N(R) and P/F are independent, i.e. each I frame has an N(S) and an N(R) sequence number that may or may not acknowledge additional frames at the receiving data station, and a P or F bit that may be set to ONE or ZERO.

I frames are sequentially numbered in the range zero through m-1 where m is the modulus of the sequence numbers (the value m-1 is followed by the value zero). For the unextended control field m = 8.

See 6.2 for a description of the extended control field.

The maximum number of sequentially numbered I frames that the Primary or Secondary may have outstanding (i.e. unacknowledged) at any given time shall never exceed m-1. This restriction is to prevent any ambiguity in the association of transmitted frames with sequence numbers during normal operation and/or error recovery action.

NOTE 2: The number of outstanding frames may be further restricted by the station storage capability; i.e. the number of I frames that can be stored for transmission and/or retransmission to recover from a transmission error.

6.1.2 Supervisory Frame - S frame

The S frame is used to perform data link supervisory functions such as acknowledging I frames, requesting re-transmission of I frames, or indicating temporary inability to receive I frames.

This frame contains one sequence number N(R) which specifies the expected I frame. There are two bits allocated for supervisory function codes (S bits).

6.1.3 Unnumbered Frame - U frame

The U frame is used to provide additional Primary and Secondary data link control functions. An exception is the UI frame as specified in 7.3.17. This format contains no sequence numbers. There are 5 bits available for unnumbered functions (M bits).

The U frame cannot acknowledge acceptance of I frames because it cannot convey N(R), the sequence number of the expected I frame.

6.2 Extended Control Field Formats

The control field defined in 6.1 may be extended by the addition of a second contiguous octet immediately following the normal control field.

This capability provides the use of N(S) and N(R) sequence numbers based on m = 128.

Control field extension for the three formats is defined below. The least significant bits of the 7 bit sequence numbers are bit 2 and 10. Bit 1 is transmitted first.

	EXTENDED CONTROL FIELD BITS															
	1st Octet									2nd Octet						
BIT ORDER	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
I frame format	0	N(S)							P	N(R)						
S frame format	1	0	S	S	0	0	0	0	/	N(R)						
U frame format	1	1	M	M	*	M	M	M	F	0	0	0	0	0	0	0

* The value of this bit is undefined

LEGEND

N(S) = Send Sequence Number

N(R) = Receive Sequence Number

P/F = Poll/Final Bit

S = Supervisory Bits

M = Modifier Bits

6.3 State Variables and Sequence Numbers

Each Secondary maintains two state variables V(S) and V(R). Each transmitted and received I frame contains two sequence numbers N(S) and N(R). Each transmitted and received S frame contains one sequence number N(R). The function of and the relation between these variables and sequence numbers are specified below.

6.3.1 Send State Variable V(S)

The send state variable V(S) denotes the sequence number of the next in sequence I frame to be transmitted. This send state variable V(S) can assume the value zero through m-1. The value of the send state variable V(S) is incremented by one with each successive I frame transmission, but cannot exceed N(R) of the last received frame by more than m-1, i.e. there cannot be more than m-1 unacknowledged I frames (see Note 2).

6.3.2 Send Sequence Number N(S)

Only I frames contain N(S), the sequence number of transmitted frames. Prior to the first transmission of an I frame, the value of N(S) is set equal to the value of the send state variable V(S).

6.3.3 Receive State Variable V(R)

The receive state variable V(R) denotes the sequence number of the next in sequence I frame to be received. This receive state variable V(R) can assume the value zero through m-1. The value of the receive state variable V(R) is incremented by one on the reception of an in-sequence

I frame, i.e. N(S) = V(R).

6.3.4 Receive Sequence Number N(R)

All I frames and S frames contain N(R), the expected sequence number of the next received frame. Prior to transmission of a frame of the above types, the value of N(R) is updated to equal the current value of the receive state variable V(R). N(R) indicates that the station transmitting the N(R) has correctly received all I frames numbered up to N(R)-1. See 7.3.18 for definition of invalid N(R).

6.4 Poll/Final Bit (P/F bit)

The Poll bit (P bit) is used in a command to solicit a response or sequence of responses.

The Final bit (F bit) is used in a response as follows:

- in NRM to indicate the final frame transmitted as the result of a command with the P bit set to ONE,
- in ARM and ABM to indicate the response transmitted in reply to the reception of a command with the P bit set to ONE.

6.4.1 Poll Bit Function

The P bit is used in a command to solicit a response.

On a link only one frame with the P bit set to ONE may be outstanding at a given time. Before another command with P bit set to ONE can be issued, a response with the F bit set to ONE should be received.

If no valid response is obtained within a system-defined time-out the transmission of a command with the P bit again set to ONE for error recovery purposes is permitted.

In NRM the Secondary cannot transmit until a command with the P bit set to ONE is received. The Primary can solicit I frames by either sending an I frame with the P bit set to ONE or by sending certain S frames (RR, REJ or SREJ) with the P bit set to ONE.

In ARM and ABM I frames can be transmitted on an asynchronous basis, i.e. without being solicited by a command with the P bit set to ONE. The P bit set to ONE is used to solicit a response at the earliest opportunity with the F bit set to ONE. For example, if the Primary needs a positive acknowledgement that a particular command has been received, it may set the P bit in the command to ONE. This will force a response from the Secondary, as outlined above.

6.4.2 Final Bit Functions

In NRM the Secondary shall set the F bit to ONE in the last frame of its response. Following transmission of

the frame with the F bit set to ONE the Secondary shall stop transmitting until a subsequent frame with P bit set to ONE is received.

In ARM and ABM, following the receipt of a command with the P bit set to ONE, a frame with the F bit set to ONE shall be transmitted.

- In the case of two way simultaneous frame transfer, where the Secondary is transmitting when the command with the P bit set to ONE is received, the F bit will be set to ONE in the earliest possible subsequent response to be transmitted.
- The transmission of a frame with the F bit set to ONE does not require the Secondary to stop transmitting. Additional frames may be transmitted following the frame which had the F bit set to ONE. Thus in ARM and ABM the F bit is not to be interpreted as the end of transmission by the Secondary. It is only to be interpreted as indicating the response of the previous frame with the P bit set to ONE.

6.4.3 Use of the P/F Bit for Error Recovery

As stated in 6.4, P bits are generated by the Primary and F bits by the Secondary. The P/F bit may be used by both functions as a basic error recovery technique. This technique is called checkpointing.

The Secondary executes checkpointing as follows:

- When a frame with the P bit set to ONE is received the Secondary expects it to acknowledge all I frames up to and including the frame sent by the Secondary with the F bit set to ONE. This is called Secondary checkpointing.

  A Combined does not execute Secondary checkpointing.

The Secondary assists the Primary to execute checkpointing as follows:

- When the Secondary transmits a frame with the F bit set to ONE it shall acknowledge all in-sequence I frames up to and including the frame sent with the P bit set to ONE.

7. COMMANDS AND RESPONSES

This section describes the formats for the Information, Supervisory and Unnumbered frames (I frames, S frames and U frames).

COMMANDS	RESPONSES
I frame format I — Information	**I frame format** I — Information
S frame format RR — Receive Ready RNR — Receive Not Ready REJ — Reject SREJ — Selective Reject	**S frame format** RR — Receive Ready RNR — Receive Not Ready REJ — Reject SREJ — Selective Reject
U frame format SNRM — Set Normal Response Mode SNRME — Set Normal Response Mode Extended SARM — Set Asynchronous Response Mode SARME — Set Asynchronous Response Mode Extended SABM — Set Asynchronous Balanced Mode SABME — Set Asynchronous Balanced Mode Extended SIM — Set Initialization Mode TEST — Test UI — Unnumbered Information XID — Exchange Identification DISC — Disconnect	**U frame format** UA — Unnumbered Acknowledgement FRMR — Frame Reject DM — Disconnect Mode RIM — Request Initialization Mode RD — Request Disconnect UI — Unnumbered Information XID — Exchange Identification TEST — Test

7.1 Information Frames (I frames), Commands and Responses

The function of the I frames is to transfer across a data link sequentially numbered frames containing an information field.

The control field of the I frame is defined in 6.1 and 6.2. Bit 1 is set to ZERO and the control field contains two sequence numbers, N(S) and N(R), which are defined in 6.3.2 and 6.3.4.

7.2 Supervisory Frames (S frames), Commands and Responses

S frames are used to perform basic supervisory functions such as acknowledgement, polling, flow control and error recovery.

S frames shall not contain an information field.

The control field of the S frame is defined in 6.1 and 6.2.

Bit 1 is set to ONE, bit 2 is set to ZERO and bits 3 and 4 define the function (commands or responses) to be performed as follows:

S FRAMES	BITS	
COMMANDS / RESPONSES	3	4
RR — Receive Ready	0	0
REJ — Reject	0	1
RNR — Receive Non Ready	1	0
SREJ — Selective Reject	1	1

An S frame contains a receive sequence number N(R), as defined in 6.3.4.

7.2.1 Receive Ready Command and Response (RR)

RR is used to indicate that the originating station is ready to receive an I frame and to acknowledge previously received I frames numbered up to N(R)-1.

RR may be used to clear busy, previously reported by the transmission of RNR. See 9.1 for further details.

The Primary may use the RR command with the P bit set to ONE to solicit a response (Poll).

7.2.2 Reject Command and Response (REJ)

REJ is used to request retransmission of I frames starting with the frame numbered N(R). I frames numbered N(R)-1 and below are acknowledged.

When a receiving station detects a sequence error it may return a REJ and then enter a REJ exception. After a station transmits a REJ it shall not transmit SREJ or REJ for an additional sequence error until the first REJ exception has been cleared. To do so, would acknowledge as correctly received all I frames with N(S) lower than N(R) in the second REJ or SREJ.
The REJ exception is cleared upon the receipt of an I frame with an N(S) equal to V(R). See 9.3.1 and 9.4 for sequence error recovery procedures.

7.2.3 Receive Not Ready Command and Response (RNR)

RNR is used to indicate busy: i.e. temporary inability to accept additional incoming I frames. I frames numbered up to N(R)-1 are acknowledged. I frame numbered N(R) and subsequent I frames received, if any, are not acknowledged, the acceptance status of these frames will be indicated in subsequent exchanges.

A Secondary receiving RNR when in the process of transmitting (in two way simultaneous frame transfer) shall stop transmitting I frames at the earliest possible time. A Secondary may return a frame with the F bit set to ONE before suspending transmission. See 9.1 for further details.

7.2.4 Selective Reject Command and Response (SREJ)

SREJ is used to request retransmission of the single I frame numbered N(S) = N(R).

I frames numbered up to N(R)-1 are acknowledged.

When a receiving station detects a sequence error it may return an SREJ and then enter an SREJ exception. The SREJ exception is cleared upon receipt of an I frame with an N(S) equal to V(R).

After a station transmits an SREJ it shall not transmit SREJ or REJ for an additional sequence error until the first SREJ exception has been cleared. To do so would acknowledge as correctly received all I frames with N(S) lower than N(R) in the second SREJ or REJ.

NOTE 3: Once SREJ has been transmitted, I frames, following the I frame requested, are accepted but not acknowledged. These shall be held for processing until the requested I frame is received. The sequence following the I frame requested may contain a sequence gap which is recoverable by another SREJ (see 9.4.4).

A sequence error detected by a station may consist of more than one missing frame. In this case each missing frame shall be treated as a separate exception. Each shall be cleared in turn as described above if SREJ is to be used for this purpose.

See 9.3.2 and 9.4 for sequence error recovery procedures.

7.3 Unnumbered Frames (U frames), Commands and Responses

U frames are used to extend the number of control functions. Some U frames reset state variables. When such a U frame is used, it may cause previously transmitted I frames not to be acknowledged.

The control field of the U frame is defined in 6.1 and 6.2.

Bit 1 and bit 2 are set to ONE, and bits 3, 4, 6, 7 and 8 are modifier bits defining the additional functions.

U FRAMES		CONTROL FIELD BITS				
COMMANDS	RESPONSES	3	4	6	7	8
SNRM		0	0	0	0	1
SNRME		1	1	0	1	1
SARM	DM	1	1	0	0	0
SARME		1	1	0	1	0
SABM		1	1	1	0	0
SABME		1	1	1	1	0
DISC	RD	0	0	0	1	0
	UA	0	0	1	1	0
SIM	RIM	1	0	0	0	0
TEST	TEST	0	0	1	1	1
XID	XID	1	1	1	0	1
UI	UI	0	0	0	0	0
	FRMR	1	0	0	0	1

7.3.1 Set Normal Response Mode Command (SNRM)

SNRM is used to place the addressed Secondary in the normal response mode (NRM). No information field is permitted with SNRM. The Secondary confirms acceptance of SNRM by transmission of UA. Upon acceptance of this command the send and receive state variables V(S) and V(R) of the Secondary are set to zero.

7.3.2 Set Normal Response Mode Extended Command (SNRME)

SNRME is used to place the addressed Secondary in the normal response mode extended (NRME), where all control fields will be two octets in length as defined in 6.2.

Except for requesting the control field extension, the operations of SNRME and SNRM are identical.

7.3.3 Set Asynchronous Response Mode Command (SARM)

SARM is used to place the addressed Secondary in the asynchronous response mode (ARM). No information field is permitted with SARM. The Secondary confirms acceptance of SARM by transmission of UA. Upon acceptance of this command the send and receive state variables V(S) and V(R) of the Secondary are set to zero. The Secondary may continue transmission following the return of UA.

7.3.4 Set Asynchronous Response Mode Extended Command (SARME)

SARME is used to place the addressed Secondary in the asynchronous response mode extended (ARME), where all control fields will be two octets in length as defined in 6.2.

Except for requesting the control field extension, the operations of SARME and SARM are identical.

7.3.5 Set Asynchronous Balanced Mode Command (SABM)

SABM is used to place the addressed Combined in the asynchronous balanced mode (ABM). No information field is permitted with SABM. The Combined confirms acceptance of SABM by the transmission of UA at the first opportunity. Upon acceptance of this command the send and receive state variables V(S) and V(R) of the Combined are set to zero.

7.3.6 Set Asynchronous Balanced Mode Extended Command (SABME)

SABME is used to place the addressed Combined in the asynchronous balanced mode extended (ABME), where all control fields will be two octets in length as defined in 6.2.

Except for requesting the control field extension, the operations of SABME and SABM are identical.

7.3.7 Disconnect Command (DISC)

DISC is used to place the addressed station into a disconnected mode.

No information field is permitted with DISC. Prior to actioning the command the Secondary confirms the acceptance of DISC by the transmission of UA. See 8.1 for further details.

7.3.8 Request Disconnect Response (RD)

RD is used to indicate that the Secondary needs to be placed in a disconnected mode (NDM or ADM). No information field is permitted with RD.

If the Secondary is in NRM, RD may be sent as a response to a command with the P bit set to ONE. RD may be sent asynchronously if the Secondary is in ARM or the Combined in ABM.

A Secondary which has sent RD and receives a command other than DISC must accept the command if it is able to do so. If it is not able to do so, it may respond RD again. The acceptance cancels the request to be placed in the disconnected mode.

7.3.9 Unnumbered Acknowledge Response (UA)

UA is used by the Secondary to acknowledge the receipt and acceptance of certain unnumbered commands. Whether or

not this response is used has been stated in the particular command definition. No information field is permitted in UA. The V(S) and V(R) of a station which has issued a mode setting command are set to zero upon receipt of the UA response to such command.

7.3.10 Disconnect Mode Response (DM)

DM is used to report a status where the Secondary is logically disconnected from the link; and is, by system definition, in NDM or ADM.

DM may be sent by the Secondary in ADM to request the remote Primary to issue a mode setting command. If sent in response to a mode setting command, DM reports that the Secondary is still in NDM/ADM and cannot action the mode setting command. No information field is permitted with DM.

Further details are to be found in 8.1 and 8.3.

7.3.11 Set Initialization Mode Command (SIM)

SIM is used to cause the addressed Secondary to start a procedure to initialize its link level control functions. No information field is permitted with SIM. The Secondary confirms the acceptance of SIM by the transmission of UA at the first opportunity and enters the initialization mode (IM). Upon acceptance of this command, the send and receive state variables V(S) and V(R) of the Secondary are set to zero.

7.3.12 Request Initialization Mode Response (RIM)

RIM is used to report the need for initialization of a Secondary. No information field is permitted with RIM. Once a Secondary has transmitted RIM, additional commands subsequently received (other than SIM, DISC, TEST or XID) are monitored only to detect an opportunity to retransmit RIM, i.e. no additional transmissions are accepted or actioned until the receipt of SIM, DISC, TEST or XID.

Further details are to be found in 8.1 and 8.3.

7.3.13 Test Command (TEST)

The TEST command is used to cause the addressed Secondary to respond with TEST at the first opportunity, thus performing a basic data link control test.

An information field is optional with the TEST command. If present, however, the received information field will be returned by the addressed Secondary with the TEST response.

The Primary considers the link level test terminated upon receipt of the TEST response or when a timer has ex-

pired. The results of the TEST Command/Response may be made available for interrogation by a higher level.

The TEST command has no effect on the mode or state variables maintained by the Secondary.

A TEST command with an unallocated address may facilitate the detection of an inadvertent loop.

7.3.14 Test Response (TEST)

The TEST response is used to reply to the TEST command. An information field, if present in the TEST command, will be returned with the corresponding TEST response.

If a received TEST command has an information field which exceeds the available storage of the Secondary, an FRMR exception may be entered. If the FRMR exception is not entered, the TEST response will be returned without the information field.

A Secondary receiving a TEST command shall transmit a TEST response when able to do so, unless:

- a response to a mode setting command is pending transmission (e.g. UA, DM),
- an FRMR exception exists.

If the Secondary is not able to respond to a TEST command, it shall:

- return an FRMR when in operational mode,
- return DM when in disconnected mode and the received command had the P bit set to ONE.

7.3.15 Exchange Identification Command (XID)

The XID command is used to cause the addressed Secondary to identify itself. An information field is optional with the XID command; if present, the information field will contain the Primary identification.

7.3.16 Exchange Identification Response (XID)

The XID response is sent as a reply to an XID command. An information field containing the Secondary identification is optional. A Secondary receiving an XID command will, if capable, send the XID response in any mode unless UA is pending or an FRMR exception exists.

On switched networks when the Secondary is constrained to send first because of the characteristics of the transmission network, it may use the XID response to request an XID exchange.

7.3.17 Unnumbered Information Command and Response (UI)

UI is used to send information (e.g. status, operation interruption, link level programs or parameters) without

impacting the V(S) or V(R) state variables at any station.

Reception of UI is not verified by sequence numbers, therefore the frame may be lost or duplicated.

There is no specific reply required upon receipt of UI.

7.3.18 Frame Reject Response (FRMR)

FRMR is used by the Secondary in an operational mode to indicate that one of the rejection conditions specified below exists:

- the receipt of a frame with a control field which is not supported,
- the receipt of a frame with an information field exceeding the agreed size.
- the receipt of a frame with an illegal information field (e.g. an S frame with an information field).
- the receipt of an invalid N(R). N(R) is defined as invalid if it acknowledges an I frame which has already been acknowledged or which has not yet been transmitted.

The Secondary shall transmit FRMR at the first opportunity.

An information field shall be used with FRMR. This information field shall provide the reasons for the rejection, and shall be coded as follows:

– without control field extension:

	INFORMATION FIELD BITS																			
BIT ORDER	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
BIT ASSIGNMENT	Rejectet Frame Control Field								0	V(S)			C / R	V(R)			w	x	y	z

– with control field extension:

	INFORMATION FIELD BITS																															
BIT ORDER	1 16																17	18 ... 24							25	26 ... 32				33	34	35	36
BIT ASSIGNMENT	Rejected Frame Extended Control Field																0	V(S)							C / R	V(R)				w	x	y	z

Bit 1 is transmitted first.

The meaning of the bits in the information field shall be, in both cases:

- bits 1 to 8 or 1 to 16 contain the control field of the received frame which caused the frame reject exception condition.
- V(S) is the current send state variable at the Secondary (bit 10 or 18 is low-order bit).
- C/R: . set to ONE indicates that the frame which caused the FRMR containd a response,
  . set to ZERO indicates that the frame which caused the FRMR contained a command.
- V(R) is the current receive state variable at the Secondary (bit 14 or 26 is low-order bit).
- w: set to ONE indicates that the control field received and returned in bits 1 to 8 or 1 to 16 is not supported.
- x: set to ONE indicates that the control field received and returned in bits 1 to 8 or 1 to 16 was accompanied by an information field which is not permitted with this command. Bit w must be set to ONE in conjunction with this bit.
- y: set to ONE indicates that the information field received exceeded the maximum information field length which can be accommodated by the Secondary or Combined.
- z: set to ONE indicates that the control field received and returned in bits 1 to 8 or 1 to 16 contained an invalid N(R).

The w, x, y and z bits in the information field of the FRMR may all be set to ZERO indicating an unspecified rejection of the frame for one or more of the conditions cited above. If required, the information field contained within the FRMR may be padded with ZERO bits so as to end on any convenient mutually agreed character, byte, word or system-dependent boundary.

8. DISCONNECTED MODES AND INITIALIZATION

8.1 Disconnected Modes

A Secondary in a disconnected mode shall be, as a minimum, capable to respond with DM to a command with the P bit set to ONE.

8.1.1 Normal Disconnected Mode (NDM)

In this mode the Secondary has a restricted normal response mode capability and shall, therefore, respond only as a result of receiving a command with the P bit set to ONE.

A Secondary in NDM actions only mode setting commands, the exchange identification (XID) command and the TEST command.

A Secondary which is unable to continue functioning may send RIM (see 7.3.12).

Non implemented commands received in NDM will be answered with DM (see 7.3.10). Other cases will be treated as follows:

- reception of a mode setting command: this will be executed normally (see 7.3.1 and 7.3.2),
- reception of an XID command: an XID response will be transmitted,
- reception of a TEST command: a TEST response will be transmitted.

8.1.2 Asynchronous Disconnected Mode (ADM)

In this mode the Secondary has restricted asynchronous response capability and may respond as defined in 5.2.2.

A Secondary in ADM will only act on mode setting commands, the exchange identification (XID) command and the TEST command. Other commands with the P bit set to ZERO will be ignored.

A Secondary which is unable to continue functioning may send RIM (see 7.3.12).

Non implemented commands received in ADM will be answered with DM (see 7.3.10). Other cases will be treated as follows:

- reception of a mode setting command: this will be executed normally (see 7.3.3 to 7.3.6),
- reception of an XID command: an XID response will be transmitted,
- reception of a TEST command: a TEST response will be transmitted.

8.2 Initialization Mode (IM)

IM is invoked when a Primary concludes that a Secondary either is behaving abnormally and needs its link control program changed or requires initialization.

In similar manner a Secondary may determine that it is unable to continue functioning due to program checks and request IM in order to obtain a correction.

A Secondary enters the IM as explained in 7.3.11.

IM is terminated when the Secondary receives and acknowledges (via UA) a mode setting command different from SIM.

8.3 Interrelation between Modes

The interrelationship between modes is graphically explained below.

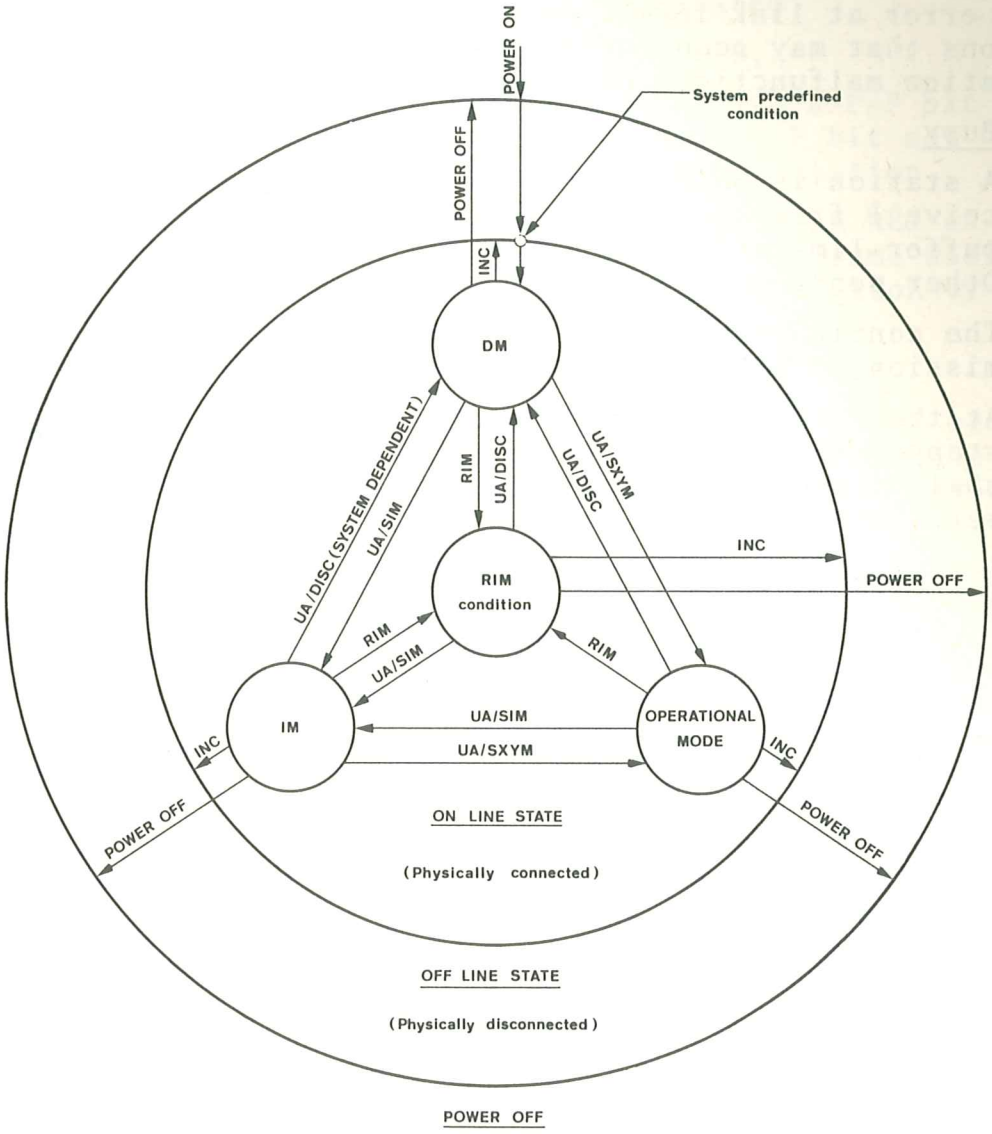

DM : Disconnected Mode
IM : Initialization Mode
UA/SIM : UA answering SIM
UA/DISC : UA answering DISCONNECT
UA/SXYM : UA answering an Operational Mode Setting Command
INC : Abnormal Incident (e.g. an exception resulting from an equipment malfunction)
RIM Condition : A RIM condition can be established in any mode by a Secondary using a RIM response. Once such a condition is established, additional commands subsequently received (other than SIM, DISC or XID) are monitored only to detect a respond opportunity to retransmit RIM. No additional transmissions are accepted or actioned until this condition is reset by the receipt of SIM or DISC.

9. EXCEPTION REPORTING AND RECOVERY

This section describes the error recovery procedures which are available to take recovery action following the detection of an error at link level. Exceptions described are those situations that may occur as the result of transmission errors, station malfunctions or operational situations.

9.1 Busy

A station is busy when it cannot receive or continue to receive I frames due to internal constraints, e.g. receive buffer limitation. In this case RNR shall be transmitted. Other pending frames may be transmitted before or after RNR.

The continued existence of busy shall be reported by transmission of RNR at each P/F exchange.

At the reception of RNR, transmission of I frames shall be stopped at the earliest possible time. A Secondary in NRM shall finish its transmission with a frame with the F bit set to ONE.

The clearing of busy is indicated by the transmission of S frames, except RNR, and by those U frames which cause the reset of the state variables V(S) and V(R). Clearing of busy at the Primary is also indicated by the transmission of an I frame with the P bit set to ONE. Clearing of busy at a Secondary is also indicated by the transmission of an I frame with the F bit set to ONE.

9.2 FCS Errors

All frames that are received with an FCS error shall be discarded without any further action being taken.

9.3 Sequence Errors During Reception of I Frames

At the reception of an I frame the correctness of the sequence number is checked. If it is incorrect a sequence exception is entered.

During a sequence exception the receiving station does not increment V(R) and does not acknowledge the I frame with a sequence error, nor any other I frame which may follow, until the sequence exception condition is cleared.

A station which retransmits I frames shall use the N(R) and P/F bit values applying at the time of retransmission.

A station which receives one or more I frames having sequence errors but no FCS error shall use the received values of N(R) and the P/F bit for data link control purposes.

9.3.1 REJ Condition

The transmission of REJ, in a two way simultaneous frame transfer, may provide an earlier sequence exception re-

covery than is possible by P/F checkpointing recovery.

The REJ condition is entered when a sequence error has been detected and an REJ frame is transmitted to report this fact. The REJ condition is terminated:

- when the requested I frame is received,
- when the receiving station responds with a P/F bit set to ONE acknowledging receipt of a P/F bit set to ONE in a frame from the transmitting station.

In the latter case, the station remains in an REJ exception and perceives that the requested I frame will not be received. It may reenter the REJ condition by sending a further REJ (see also 7.2.2).

9.3.2 SREJ Condition

SREJ is primarily used to initiate more efficient recovery by requesting the retransmission of a single I frame following the detection of a sequence error. This prevents the unnecessary retransmission of all additional I frames which may have been subsequently transmitted.

SREJ can only be used to recover from sequence errors up to size d.

d is the maximum difference between the receive state variable V(R) and the send sequence number N(S) of the I frame which was used to detect the sequence error. d = m-1-k, where k is the maximum number of outstanding I frames that the stations have defined by bilateral agreements.

The station selects the SREJ exception and enters the SREJ condition by the transmission of SREJ. Further I frames which are received without FCS error but with $N(S) \neq V(R)$ are buffered in sequence of reception for processing after the sequence exception has been cleared.

The SREJ condition is terminated when the requested I frame is received.

9.4 Error Recovery

9.4.1 Time Out Recovery

When a station, due to a transmission error, does not receive (or receives and discards) the later I frames in a sequence it will not detect a sequence error. It cannot therefore issue an REJ or SREJ. The station which transmitted the unacknowledged I frames may, following the completion of a system specified time out period, take appropriate recovery action to determine the point at which retransmission shall begin.

NOTE 4: The station initiating such time out recovery shall ensure that any action it takes solicits a response from the receiving station. Simple retransmission of all the outstanding I frames may not be sufficient. The acknowledgement of these frames may have been lost due to transmission errors and the receiving station may be awaiting further frames in the sequence.

Thus retransmission shall be accompanied by one of the following techniques:

- *the initiation of P/F checkpointing by the station effecting time out recovery,*
- *the reporting of sequence errors by the receiving station using REJ or SREJ frames.*

9.4.2 Checkpointing

The operation of checkpointing is defined in 6.4.3. Checkpoint recovery is not initiated under the following conditions:

- when an REJ with the P/F bit equal to ZERO has been received and actioned, checkpoint retransmission is inhibited on the next frame received with P/F bit equal to ONE. REJ recovery is already in operation,
- when an SREJ with the P/F bit equal to ZERO has been received and actioned, checkpoint recovery is inhibited on the next frame with the P/F bit equal to ONE if this frame is SREJ and contains the same N(R) as the first SREJ,
- when a U frame with the P/F bit equal to ONE is received, P/F bit recovery is inhibited, because of the lack of N(R) in the U frame,
- when an SREJ with the P/F bit equal to ONE is received, SREJ retransmission takes place instead of checkpoint retransmission.

9.4.3 REJ Recovery

A station receiving an REJ effects error recovery by retransmission starting with the lowest numbered unacknowledged I frame followed by new I frames waiting for a first transmission.

When a receiving station uses REJ it is the responsibility of the transmitting station to ensure that REJ recovery is prevented if P/F checkpoint recovery has already been initiated. This situation may arise when I frames are lost and P/F checkpointing is started using an S or U frame. A send sequence number N(S) is not available for the receiving station to detect missing frames. However, by acknowledging the P/F checkpoint

and transmitting back the N(R) of the last correctly received frame, checkpoint recovery may be initiated.

If subsequent I frames are received with N(S) ≠ V(R) REJ may be sent to report this detected sequence error. The transmitting station receiving this REJ shall ignore it.

9.4.4 SREJ Recovery

A station receiving SREJ effects error recovery by retransmitting the I frame specified by the SREJ.

I frames that may have been transmitted following the I frame indicated by the SREJ command/response are not retransmitted as the result of receiving an SREJ. Additional I frames awaiting initial transmission may be transmitted following the retransmission of the specific I frame N(S) = N(R) as requested by the SREJ.

When a receiving station uses SREJ, it is itself responsible for ensuring that recovery by SREJ is not initiated in addition to P/F checkpoint recovery. This situation may arise when I frames are lost and checkpointing is started by setting the P/F bit equal to ONE in an S or U frame. A send sequence number N(S) is not available for the receiving station to detect missing frames.

However, by acknowledging the P/F checkpoint and transmitting back the N(R) of the last correctly received frame, checkpoint recovery may be initiated.

If subsequent I frames are received with N(S) ≠ V(R) the receiving station will detect the sequence error. However, SREJ may not be sent. The station does not start the SREJ interval and an SREJ exception is not entered.

Thus a station will not send SREJ with the same N(R) as that in the previously sent frame with the P/F bit equal to ONE.

A station which has entered an SREJ exception, when ready to issue the next frame with the P/F bit equal to ONE shall send an SREJ with P/F bit equal to ONE. When a transmitting station receives an SREJ frame with P/F bit equal to ONE, it will inspect the N(R). If this N(R) is equal to an N(R) in a previous SREJ which has been actioned since the last P/F bit equal to ONE was transmitted, then no further action will be taken.

9.5 Time-Out Considerations

When a station issues any command or response that should cause the remote station to return any frame, it may check whether this expected frame is received within a defined time period. This is called a "time-out" function.

A station may implement different time-out functions, dependent on the events that are expected. Two of these time-out functions are defined below.

- Timer T 1

  A Primary that issues a P bit, should check whether the response or responses involved by this P bit are received within a defined time period. This time-out function is controlled by a timer T 1 and is called "wait for F" time-out.

- Timer T 2

  A Primary or a Secondary in ARM that issues I frames, may check whether acknowledgements are received within a defined time period. This time-out function is controlled by a timer T 2 and is called "wait for N(R)" time-out.

  In order that a station may have only one timing state at any one time, when a frame with the P bit set to ONE is transmitted, timer T 2 may be suppressed in favour of timer T 1.

  NOTE 5: It is recommended that, when timer T 2 expires, a station should

  - *either issue a P bit and start timer T 1 (only applicable to Primary),*
  - *or retransmit all outstanding I frames and restart timer T 2.*

  In the latter case, a station should not only retransmit the last transmitted I frame, as the remote station might possibly have detected the N(S) sequence error and consequently will only react on the receipt of the I frame with the expected N(S). On the other hand, if the remote station did not detect the N(S) sequence error, it may regard a retransmitted I frame as a sequence error and return an REJ frame with an N(R) greater than the V(S) of the station which retransmits. This should not be regarded by the retransmitting station as an "invalid N(R)" but the retransmitting station should increment V(S) appropriately.

9.6 Frame Rejection

A frame rejection is established upon the receipt of an error free frame (no FCS error) with a rejection condition as defined in 7.3.18. At the Primary this exception is subject to recovery at a higher functional level. As a result of higher level action, the Primary may issue the appropriate mode setting command.

At the Secondary this exception is reported by FRMR for appropriate Primary action. Once a Secondary has established an FRMR exception no additional transmissions are accepted or actioned except for examination of the state of the P bit

and the N(R) count. The FRMR response is repeated at every opportunity until reset by the Primary.

In balanced operation, if recovery is not effected by the other Combined within a specified time-out interval, the reporting Combined may repeat FRMR, or it may choose to assume control of the recovery functions as cited above. If the other Combined that receives FRMR is unable to effect an appropriate recovery action, it will reply with an FRMR of its own, rejecting the received FRMR. The Combined that sent the original FRMR will then initiate an appropriate recovery function as defined above.