CONTROL DATA'S CDCCP A USER'S PERSPECTIVE OF BIT-ORIENTED DATA LINK CONTROL PROTOCOLS

By

J.W. CONARD

CONTROL DATA CORPORATION

CDCCP
A USER'S PERSPECTIVE OF BIT-ORIENTED DATA LINK CONTROL PROTOCOLS

James W. Conard

Control Data Corporation

INTRODUCTION

The advent of bit-oriented protocols continues to generate a flurry of activity in the world of data communications. Engineers and programmers are struggling with flags, frame structures, commands, and responses. Users are trying to assess the impact of this activity on their data communications planning. To do this, they must sort through a maze of acronyms: ADCCP, SDLC, HDLC, CDCCP, and so on. They must know the difference between an "element of procedure" and a "class of procedure." They must decide if their equipment should be a "primary" or a "secondary", or both. They must give consideration to a variety of options, extensions, and parameters. It is our purpose here to bring this seemingly frenetic activity into perspective; to provide you, the user, with a road map through the new area of bit-oriented protocols.

We'll begin with an overview of bit-oriented protocols, and continue with a discussion of the evolution of the new technique. This is followed by a review of the major technical characteristics of Control Data's CDCCP. New material is included on classes of procedure, comparison of standards, and definition of various implementation parameters.

AN OVERVIEW OF BIT-ORIENTED PROTOCOLS

A data link control protocol is a set of very specific rules under which data is exchanged between business machines via a communications circuit. The business machines may be terminals, concentrators, message switches, or computers, in any mix. A link protocol typically defines initialization of an established link, control of normal data interchange, termination of the link, and perhaps most important to the user, abnormal condition recovery techniques which serve to assure message integrity.

Strictly speaking, the term link control excludes other levels within the communications procedure hierarchy (Figure 1). One of the objectives of the new protocol was, in fact, to clearly delineate the interface between link control and higher levels such as device and message control. The characteristics of these levels do, however, impact on link level control. The prudent system designer keeps a wary eye on their requirements.

Link control protocols have traditionally been character-oriented. They utilized, either singularly or in sequence, defined character structures from a given code set to convey supervisory information. Even though character-oriented protocols represent the vast majority of protocols in use today, it has long been recognized that they suffer from many deficiencies. Among these are:

1. The necessity to distinguish between data and control characters within a code set places a burden on hardware and software implementation.

2. The assignment of characters for link control subtracts from the combinations otherwise available for information transfer.

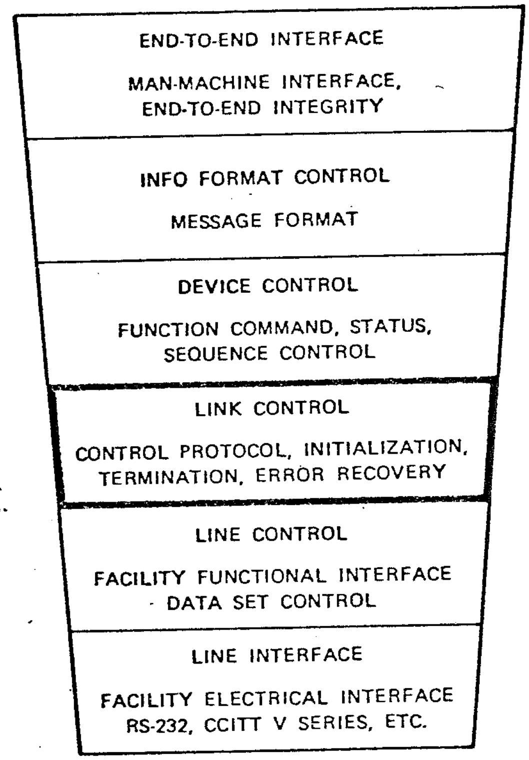

Figure 1. Communications Control Hierarchy

3. The character orientation meant that they were not naturally transparent to the structure or encoding of the text.

4. Transparency could only be achieved by invoking complicated escape techniques and at the expense of incompatibility with non-transparent protocols.

5. The mixture of message control, device control, and link control forced a significant amount of processing at a low functional level, and blurred the interface between these logically independent functions.

6. Error checking is usually done only on the text, thus exposing supervisory sequences to undetected errors which complicate error recovery.

7. The inherent two-way alternate nature of these protocols does not economically utilize full-duplex facilities.

8. The rigid structure of character-oriented protocols lack flexibility and expandability.

Bit-oriented protocols are an outgrowth of attempts to overcome these deficiencies. The inherent characteristics of the new protocol which, when properly applied, overcome the disadvantages of character protocols include:

1. Bit orientation. They utilize positionally located control fields rather than code-set combinations for link control.

2. Code independence. The use of framing flags and control fields divorces link control totally from the pattern or code structure of the information content. Thus bit-oriented protocols are inherently transparent.

3. Reliability. The use of one standard format for all information and supervisory transmission permits error checking of control as well as text information.

4. Flexibility. Bit-oriented protocols permit implementation in a variety of applications using a variety of communications facilities without modification of basic link control procedures.

5. Efficiency. The techniques applied are designed to take full advantage of full-duplex facilities, while retaining the ability to operate efficiently on half-duplex facilities where desired.

6. Hierarchical Independence. Bit-oriented protocols separate the functions of link control from those of device and message control.

Bit-oriented protocols combine these characteristics to provide greater utilization of facilities than is possible with the older, character-oriented methods. Their application permits the user to more fully achieve the benefits of his data communications system.

While reviewing what the bit-oriented protocols are, it is equally important to consider what they are not. The new protocols are not the total solution to the communications problem. They are only a link-level control mechanism, and thus are concerned solely with the transfer of data at that level. They are not a network protocol. They do not control the flow of information between users in a multinodal network. They can, however, be applied between nodes or between a node and a user. Any necessary end-to-end controls must be imbedded in the bit-oriented frame as information.

HISTORICAL PERSPECTIVE

It is natural for the user to ask: Do the new bit-oriented protocols represent an evolutionary or revolutionary departure from the older approaches? The answer may be found in a brief review of the evolution of data communications protocols.

Data link control protocols are as old as data communications. Over the years these protocols have been evolving typically to fulfill the requirements of a particular application. Early systems, using Baudot code, had no inherent link control capability. They relied totally on sequences of data characters to implement supervisory functions. The advent of other character sets led to protocols using controls derived from these sets. Each manufacturer developed protocols reflecting the needs of his product line and usually optimized for a specific implementation.

Control Data and IBM, to cite two examples, have each developed several protocols which achieved fairly widespread application. Control Data developed Mode II and Mode IV, among others, each with different characteristics and areas of application. IBM did the same with GPD, STR and BSC. Other manufacturers and users' groups also constructed protocols to meet their unique requirements. All of these various protocols were character-oriented in approach and generally incompatible with each other.

Standards organizations here and abroad, especially ANSI, ISO and ECMA, recognized the problem and struggled to resolve the incompatibilities. For lack of standardization, the protocols developed by the larger dominant manufacturers tended to fill the vacuum by becoming, in effect, de facto standards. This has certainly been the case with IBM's BSC developed in the late 1960s.

The standards organizations finally reached agreement with the publication in 1971 of ANSI's X3.28 on the user of ASCII control characters for information interchange and ISO's R1745 Basic Mode Control Procedures. It is worthy of note, however, that even before publication of these standards, both bodies were already at work on bit-oriented protocols. This was the result of recognition of deficiencies in character-oriented protocols that surfaced during the standardization process, and as a result of on-line experience with these protocols.

In late 1969, ANSI and ISO began formal work leading toward the development of bit-oriented standards. Other groups such as IATA, ICAO and ECMA also initiated study efforts, as did the various manufacturers. These efforts reached fruition and caught the attention of users in mid-1973 with the announcement by IBM of their bit-oriented protocol known as Synchronous Data Link Control (SDLC). ANSI followed in early 1974 with the first draft of Advanced Data Communications Control Procedures (ADCCP). ISO also began to formulate their High Level Data Link Control (HDLC) Procedures.

This brief review demonstrates that the "new" approach to link control, the bit-oriented protocol, represents no more than a natural and evolutionary milestone in the continued effort to improve data communications. It is, perhaps, revolutionary in the sense that a large degree of standardization is being achieved before widespread implementation.

TODAY AND TOMORROW

After having traced the evolution of bit-oriented protocols, it is appropriate to review the present "state of the protocol," and to attempt to assess the probable future impact of this approach.

The present status of bit-oriented protocols may be characterized as rapidly approaching maturity. Looking at the progress of the standards activity, first we see that ANSI X3S34, which bears the responsibility for data communications protocol procedures, has completed the fifth draft of ADCCP. Balloting on this draft has been completed at the parent X3S3 (Data Communications) level. Comments from the various members are now being resolved; and it is expected that ADCCP will go to ballot at the X3 (Computers and Information Processing) level by the end of 1976. This should result in publication of an American National Standard within a year. CDC is a very active member of task group S34, as well as its parent bodies.

ISO, the International Standards Organization, and more specifically ISO/TC97/SC6, has chosen to divide the HDLC standard into three or more standards. The frame structure standard, IS 3309, has been approved and published. The elements of the procedure standard, DIS 4335, have been approved at the SC6 level, and are now out on a 6-month ballot at the TC97 level. ISO has elected to standardize classes of procedure as separate documents. Two attempts at defining several classes have been unsuccessful. Further meetings to resolve differences are planned for November 1976 and early 1977. CDC also participates in this activity through its ANSI representatives.

While the standards are not binding, they will exert strong influence on the industry. They will provide guidelines that will point future software and hardware development in the same direction.

Manufacturers, meanwhile, are busy developing and announcing their own bit-oriented protocols. IBM has their SDLC, Burroughs their BDLC, NCR has BOLD, and CDC has CDCCP. The best information available indicates that all of these protocols are close to complete subsets of ADCCP. Some manufacturers have also announced products including bit-oriented protocol packages. Some of these are operational in limited applications.

The federal government is also in the process of preparing standards for publication as FIPS. These are also ADCCP-compatible. ECMA and CCITT are expected to publish standards compatible with HDLC.

Although not yet fully mature, bit-oriented protocols can be expected to have a major impact on data communications over the next five years. A primary reason for this is the impetus provided by IBM. SDLC is expected to be the only bit-oriented protocol that IBM will support. This will require manufacturers of terminals and processors, as well as software suppliers, who hope to interface IBM equipment to adopt SDLC which is a subset of ADCCP.

Another impetus toward adoption is that, for perhaps the first time, a broad base of standardization exists before widespread implementation begins. This fact has been recognized by IC manufacturers who are now developing chips to handle portions of the bit-oriented protocol functions, especially those related to frame structure.

CONTROL DATA'S CDCCP

Control Data Corporation formally initiated an effort to define a corporate-standard, bit-oriented link control protocol in mid-1974. The objective of this effort was to generate a standard protocol which would facilitate the exchange of information in a variety of applications, and be capable of accommodating simple to complex, low to high-speed synchronous sources and sinks. The minimum requirements were that the new protocol provide for two-way alternate to two-way simultaneous operation, permit multidrop configuration, be suitable for satellite transmission, provide for nonsymmetric and symmetric operation, and include effective levels of error detection. It was also required that the protocol be modular in definition and implementation to permit wide application, and to permit revision with minimum impact on implementation.

To meet these objectives, a task force was established with representatives from various divisions of the corporation with an interest in communications. The protocol to be standardized was called Control Data Communications Control Procedure (CDCCP).

The efforts of the task force resulted in a draft of a proposed standard for CDCCP. This draft is now in the process of revision to incorporate the results of the past year's activity, especially in the area of definition of classes of procedure. The revised draft will be reviewed by the various concerned divisions and should become a corporate standard in early 1977.

CDCCP spans the entire set of bit-oriented protocols now in the process of standardization and implementation. These include IBM's SDLC, ANSI's ADCCP, and ISO's HDLC. CDCCP is, therefore, geared to satisfy any of these requirements by use of a subset of the CDCCP protocol. The CDCCP draft is already serving to provide design guidelines to various developing divisions.

CDCCP, as typical of the bit-oriented protocols now being considered for supplementation, may be characterized as being comprised of three major constituent parts. These are: the frame structure, the elements of procedure, and the classes of procedure.

Before delving into the characteristics of these, however, it will be useful to set the stage by defining the interconnect arrangements and station configurations which will be typical of CDCCP applications.

CDCCP has a wide variety of potential applications. It is suitable for two-way alternate and two-way simultaneous operation using a variety of data link configurations, including full and half-duplex, point-to-point, multipoint, switched, and non-switched. The three facility configurations expected to be most common in CDCCP applications are illustrated in Figure 2.

A point-to-point facility is one which interconnects two and only two stations. Point-to-point facilities may be either non-switched, sometimes referred to as private line or dedicated, or they may be switched. The difference between switched and non-switched is one of facility acquisition. In the switched case the facility must be

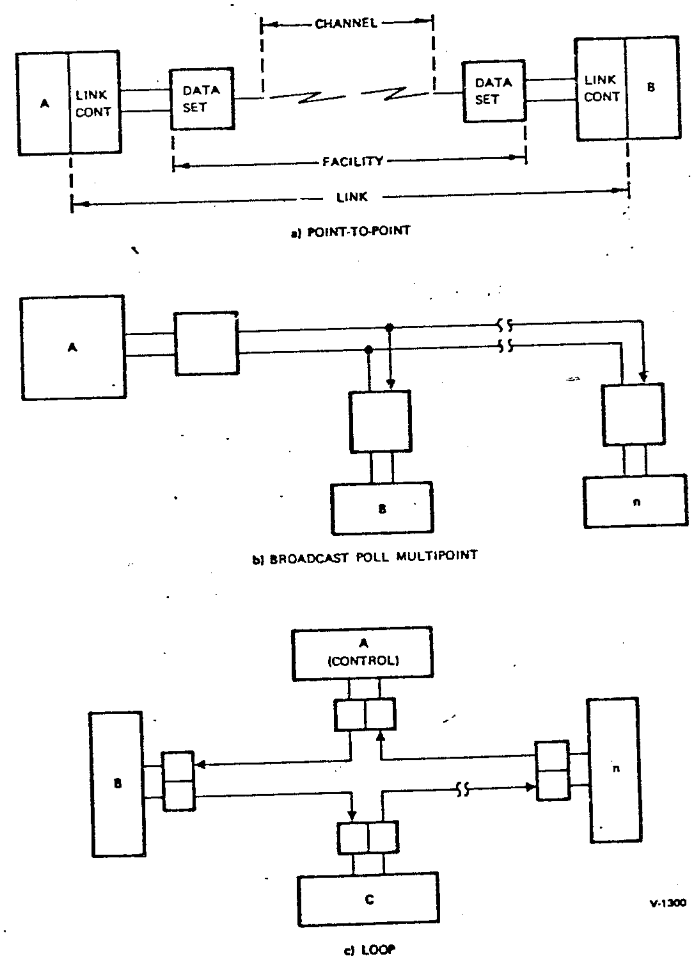

Figure 2. Interconnect Arrangements

acquired prior to the transfer of data and released at the end of the transfer. Non-switched facilities are dedicated nd usable on demand.

A multipoint arrangement, expected to be very common for these applications, is the broadcast polling arrangement which consists of a single master and two or more remote stations. Transmissions from the master are received by all remotes. Transmissions from the remotes are received only by the master. This multipoint arrangement requires four-wire channels.

Many special and hybrid combinations of interconnect arrangements are possible. The most likely to be encountered in these applications is the loop arrangement. The loop configuration consists of two or more point-to-point facilities arranged such that the loop starts and ends at the same location. The point-to-point facilities are normally two-wire channels, and operate in simplex mode: A transmits to B, B transmits to C, and so on around the loop. Transmission in the reverse direction is not possible. Each station on the loop operates as a repeater. Loop facilities may be encountered which are completely user-owned, especially when located within the confines of a building. Others may use common carrier facilities when geographically dispersed.

Operating with these interconnect arrangements, CDCCP recognizes the existence of four types of logical stations operating in any one of three station configurations. Logical station types are: Primary, Secondary, Balanced, and Configurable. Station configurations are: Unbalanced, Balanced, and Symmetrical. These are illustrated in Figure 3 and described below.

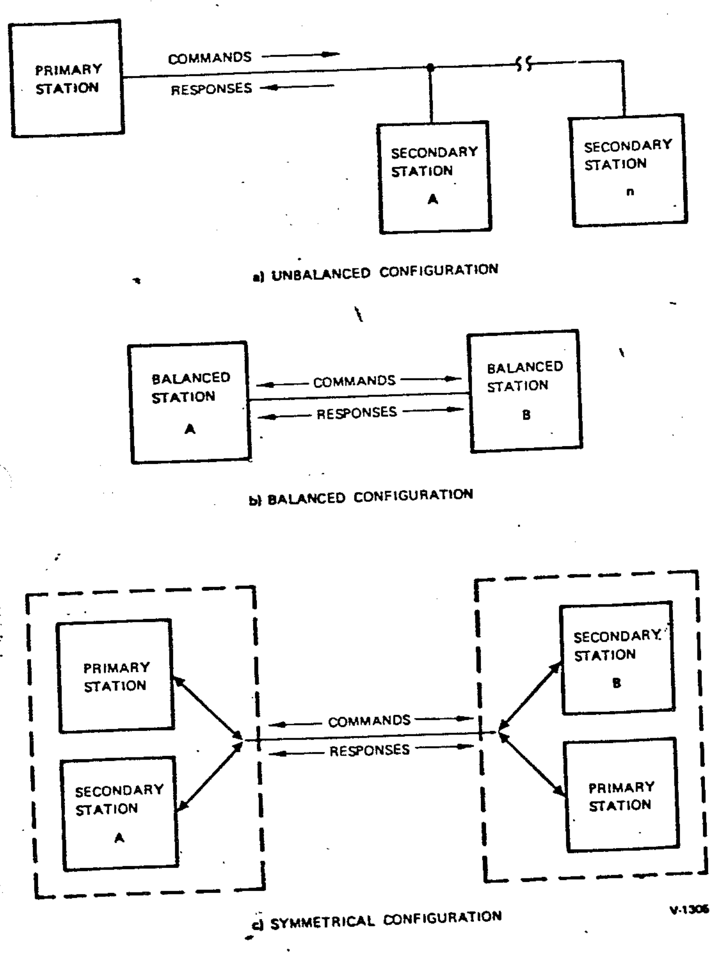

Figure 3. Logical Station Configurations

A logical primary station is an entity which has only primary link control capability. A primary station assumes responsibility for organization of data flow and for link-level error recovery. The frames transmitted by a primary are always referred to as command frames. A primary transmits command frames to, and receives and interprets response frames from, a secondary or secondaries on the link. A primary maintains a single and separate information exchange capability with each secondary on the link.

A logical secondary station is an entity which has only secondary link control capability. A secondary station has no direct responsibility for control of the link but instead responds to primary station control. Frames transmitted by a secondary are always referred to as response frames. A secondary transmits response frames to, and receives and interprets command frames from, a primary station on the link. A secondary station maintains a single information exchange capability with a primary station.

A logical balanced station is an entity which shares equal and complementary functional link control capability with another balanced station on a link. A balanced station has no identifiable primary or secondary entity. A balanced station transmits commands and response to, and receives and interprets commands and response from, another balanced station. A balanced station maintains a single information exchange capability with another balanced station.

A configurable station is one which has the capability of being, at different times and as the result of mode setting commands, more than one type of logical station, i.e., a primary station, a secondary station, or a balanced station.

Turning to configurations we have the following:

An unbalanced configuration is one which consists of a primary station and one or more secondary stations. Such a configuration is unbalanced in the sense that prime responsibility for initialization of the link, for information transfer, and for link-level error recovery rests with the primary station. The primary station commands a secondary station to assume modes and states as appropriate.

A balanced configuration is one which consists of two balanced stations connected via a dedicated or switched point-to-point facility. In the balanced configuration each station has equal and complementary responsibility for control of the data link.

It is possible to combine a logical primary and a logical secondary in one physical station and to operate, with another such station, two totally independent, unbalanced configurations. It is further possible to operate this configuration on a single communications facility by multiplexing transmissions on a frame-by-frame basis. The result is not a new configuration but simply the symmetric combination of two independent, unbalanced, point-to-point configurations.

With this brief summary of the various arrangements of station types and configurations in mind, we can turn to the realm of frame structures, elements, and classes.

FRAME STRUCTURE

The frame structure provides a common structure for all supervisory and information transfers in the bit-oriented protocols. The frame structure governs the structure, formatting, and significance of the various fields in the frame, as well as the frame-delimiting flags and frame-check sequences. The following paragraphs provide a broad overview of the technical aspects of the frame structure.

A frame is a sequence of contiguous bits bounded by, and including, opening and closing flag sequences. A valid frame is a minimum of 48 bits in length and must conform to the structure illustrated in Figure 4. Frames containing only link-control sequences form a special case where no I field is present.

FLAG SEQUENCE (F). All frames open and close with the flag sequence. This sequence has the binary configuration 01111110; that is, a zero bit followed by six one bits, followed by a zero bit.

The opening flag serves as a position reference for the address and control fields, and initiates transmission error checking. The closing flag serves as a position reference for the flag-check sequence.

Transmitters must send only complete 8-bit flags. All receivers attached to the data link must search continuously, on a bit-by-bit basis, for the flag sequence. Thus, the flag sequence provides frame synchronization.

An F may be followed by a frame, another F, or an idle line. An F which closes a frame may also be used as the opening F on a following frame. Any number of F's may be transmitted between frames.

Since the F sequence brackets and synchronizes the frame, it must be prevented from occurring in any field of the frame. This is accomplished by the zero insertion technique described below.

Each transmitter must insert a zero bit following five contiguous one bits anywhere between the opening and closing flag sequences. The insertion of the zero bit thus applies to the address, control, information, and FCS fields, and effectively prevents the fortuitous transmission of the F sequence 01111110.

After detecting the opening flag (start of frame), each receiver continuously monitors the received bit stream, and removes any zero bit which follows a succession of five contiguous one bits. Note that zero insertion at the transmitter follows the computation of FCS, and that zero deletion at the receiver precedes the FCS check process.

ADDRESS FIELD (A). The address field (A) immediately follows the opening flag of a frame and precedes the control field. This field always contains the address of the secondary station. The primary station is never identified. The address field is N octets in length where $N \geq 1$. The contents of the field may be a single, group, or global address.

Two addressing modes are defined for the secondary station link-address field. These are the basic and extended modes described below. For a specific link the maximum number of octets must be explicitly defined.

In the basic mode, the address field contains one address. In this mode, address extension is not permitted. All 256 combinations are available for addresses. This basic mode field consists of one 8-bit octet with the format illustrated in Figure 5.

In the extended mode, the secondary link address field is a sequence of octets which comprise a single secondary or balanced station address. The least significant bit is used as an extension indicator. When this bit is zero, the following octet is an extension of the address field. The address field is terminated by an octet having a one in bit position one (least significant bit). Thus the address field is recursively extendable. The format of the extended address field is also illustrated in Figure 5.

Two or more secondaries may be required to recognize the same group or global address. Each secondary, however, responds with its individual address.

CONTROL FIELD (C). The control field (C) is located immediately following the address field and preceding the information field in the frame structure. The control field is used to convey commands, responses, and sequence numbers necessary to control the data link.

There are two modes defined for the control field. These are the basic and extended modes described in the following paragraphs. For a given link the mode must be specifically identified.

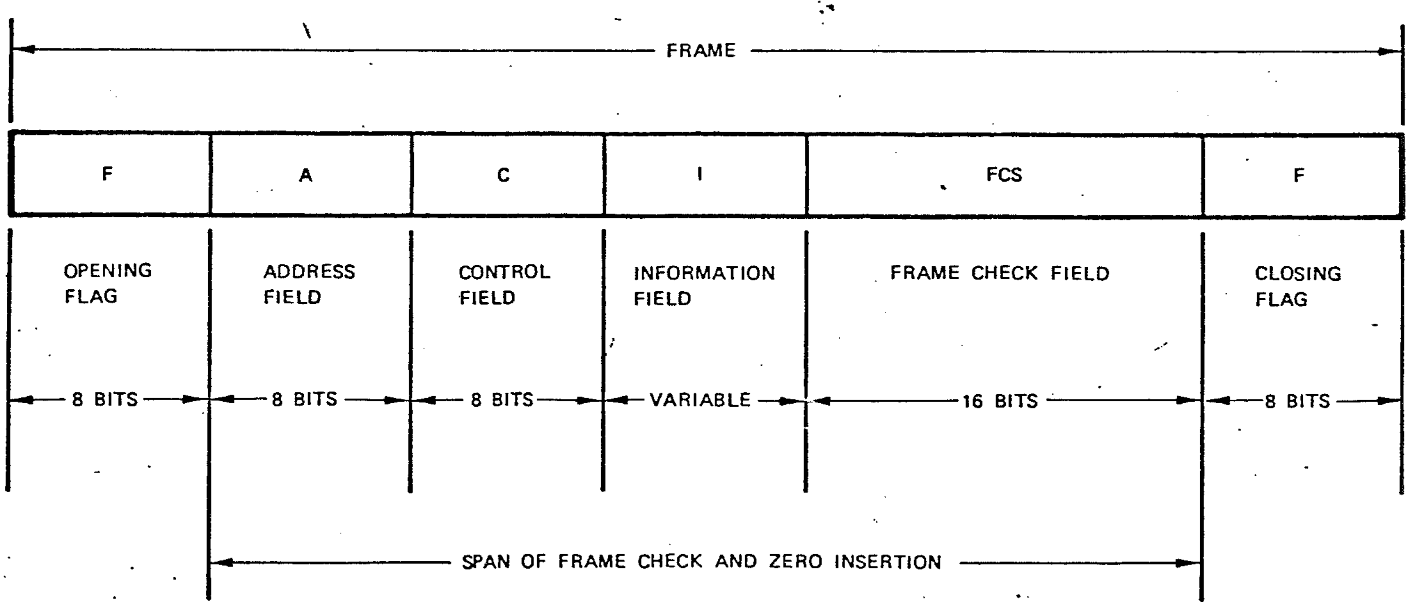

Figure 4. Frame Structure

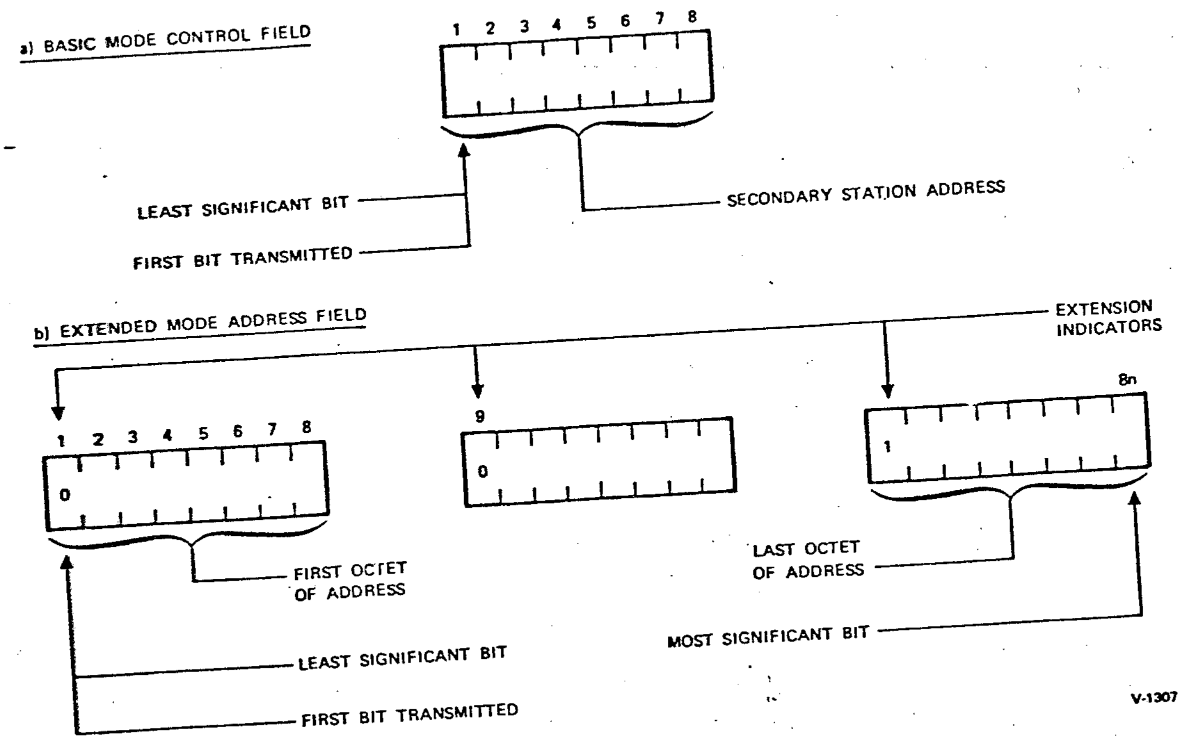

Figure 5. Address Field Format

The basic control field consists of a single, 8-bit octet. This field is structured into one of three formats. These are the information transfer format used by all stations to transfer information, the supervisory format used to convey link supervisory data, and the unnumbered format used to provide additional link control functions.

In addition, each format includes a format identifier and a poll/final bit. The poll/final bit serves as the send/receive control. A poll (P) bit is sent only by a primary or balanced station to authorize secondary or balanced transmission. The final (F) bit is used only in response to a P bit. Only one P bit is outstanding, i.e., unanswered by an F bit, on a data link.

Figure 6 illustrates the basic mode control field.

The basic mode control field provides for a modulo 8 sequence count. On long propagation delay links, e.g., satellite links, it may be necessary to extend the sequence number modulus. The extended mode control field provides this capability.

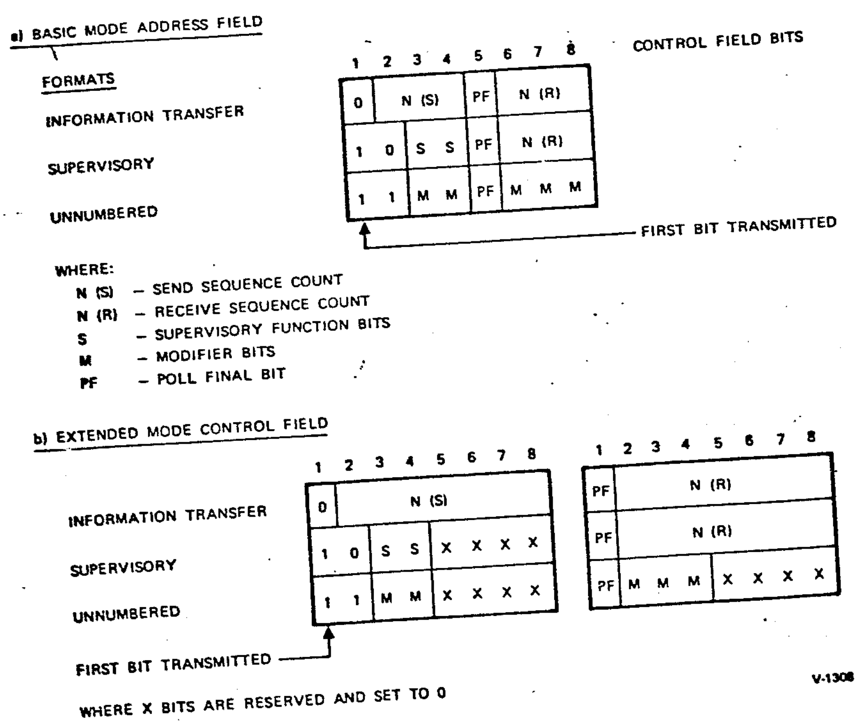

Figure 6. Control Field Formats

The control field is extended by the addition of a second contiguous octet immediately following the basic field. This extension increases the modulus count to 128. The three formats for an extended mode control field are also illustrated in Figure 6.

INFORMATION FIELD (I). A frame exists as a vehicle for transporting the data contained in the information field (I). The data link control is completely transparent to the contents of the I field. The I field may, therefore, consist of any number of bits, in any code, related to character structure or not. The I field is unrestricted as to length. It should be recognized, however, that typical length is contingent on system requirements and limitations beyond the link level. Factors limiting I-field length may include channel error characteristics, station buffer sizes, and the logical properties of the data.

The occurrence of a flag or abort sequence within the I field is prevented by the zero insertion technique described previously.

I fields are normally included in every frame having a C field with an information transfer format. These information transfer frames are the only ones which are sequence-numbered. An information field with a length of zero is specifically permitted.

Provisions are also made for an I field in an unnumbered C-field format. Such frames are not protected by sequence checking.

FRAME CHECK SEQUENCE (FCS). Each frame includes a 16-bit frame check sequence (FCS) immediately following the I field (or the C field if there is no I field) and preceding the closing flag. The FCS field serves to detect errors induced by the transmission link and to validate transmission accuracy. The 16 bits result from a mathematical computation on the digital value of all binary bits (excluding inserted zeros) in the frame including the address, control and information fields.

The process is known as cyclic redundancy checking, using a generator polynomial of $X^{16} + X^{12} + X^{5} + 1$. The transmitter's 16-bit remainder value is initialized to all ones before a frame is transmitted. The binary value of the transmission is premultiplied by X^{16}, then divided by the generator polynomial. Integer quotient values are ignored, and the transmitter sends the complement of the resulting remainder value, high-order bit first, as the FCS field.

The receiver will discard a frame in error and will not advance the receive sequence count, thus causing a retransmission of the errored block.

ELEMENTS OF PROCEDURE

The elements of procedure comprise the building blocks of CDCCP. All elements employ the common frame structure discussed previously. Elements of procedure include operational modes, commands and responses. Using these common elements, various classes of procedure which meet the requirements of various application situations can be constructed. The paragraphs which follow summarize the various elements and their characteristics.

CDCCP defines six logical modes for station operation. These are divided into an initialization mode, three mutually exclusive information transfer modes, and two mutually exclusive disconnect modes.

In the initialization mode, a station is operating under control of a system-defined initialization procedure, and may exchange information in any manner defined for that station. This may include unformatted bit streams as well as CDCCP frames.

In the information transfer modes, a station is fully operational and may exchange information, supervisory, and unnumbered frames under CDCCP protocol rules.

In the disconnected modes, a station is logically disconnected from the link and may not transmit or receive information frames, but may accept mode setting commands and issue appropriate responses. Disconnected modes are provided to prevent a station from appearing on the link in a fully operational mode which could cause ambiguity as to station mode and status.

A CDCCP station is minimally capable of operating in one information transfer mode and one disconnected mode. An initialization mode capability is a functional extension to basic capability. A station may be capable of operating in more than one information transfer or disconnected mode at different times. Such a station would be a configurable station capable of operating in more than one class of procedure.

Each of the six logical modes is described in the following paragraphs.

Normal Response Mode (NRM) is an information transfer mode for unbalanced configurations. In NRM, a secondary station may initiate transmission only as the result of receiving explicit permission to do so from the primary station. Explicit permission is defined as transmission by the primary of a command frame with the poll bit set to one. After receiving permission, the secondary can initiate a response transmission. The response transmission may consist of one or more frames while maintaining an active link state. The last frame of the transmission will be explicitly indicated by the secondary by means of a final bit set to one. Following transmission of the last frame, the secondary will stop transmitting until explicit permission is again received from the primary.

Asynchronous Response Mode (ARM) is an information transfer mode for unbalanced configurations. In ARM, a secondary, after it is commanded to this mode by the primary, may initiate transmission without receiving explicit permission from the primary. Such an asynchronous transmission may contain single or multiple frames, and is used for information-field transfer and/or status changes in the secondary. Examples of status changes are the number of the next expected frame, change from a ready to a busy condition or vice versa, or establishment of an exception condition. In ARM, a secondary will transmit a frame with a final bit set to one only in response to a received command frame with the poll bit set to one. Additional response frames may be transmitted following the frame which has the final bit set to one.

Asynchronous Balanced Mode (ABM) is the information transfer mode for balanced stations in a balanced configuration. A balanced station in ABM may initiate transmission at any respond opportunity without receiving permission from the other balanced station. Such a transmission may contain single or multiple frames used for information transfer or status. In ABM, a balanced station must transmit a frame with the final bit set to one only in response to a received command frame with the poll bit set to one. Additional response frames may be transmitted following the response frame with the final bit set to one.

Normal Disconnected Mode (NDM) is a disconnected mode in which the secondary or balanced station is logically disconnected from the link and may not initiate or receive information frames. A secondary or balanced station in NDM may initiate transmission of a response frame only as the result of receiving a command frame with the poll bit set to one. A station is system-predefined as to the conditions, in addition to a disconnect command, which cause it to assume NDM. Examples of such conditions are initial power on, power on after temporary loss of power, or manual reset of link-level logic.

Asynchronous Disconnected Mode (ADM) is a disconnected mode in which the secondary or balanced station is logically disconnected from the link and may not initiate or receive information frames. A station in ADM may initiate transmission without receiving a command frame with the poll bit set to one. Such transmission is, however, limited to a request for disconnect mode (DM) or a request for initialization mode (RIM). A station is system-predefined as to the conditions, in addition to a disconnect command, which cause it to assume ADM. Examples of such conditions are initial power on, power on after temporary loss of power, or manual reset of link-level logic.

Initialization Mode (IM) is provided to allow a secondary or balanced station to be initialized by the remote primary or balanced station. A secondary or balanced station enters initialization mode upon responding to a set initialization mode (SIM) command. A station may request initialization mode by sending a request initialization mode (RIM) response. While in IM, a station may exchange information in any manner specified for that station, including unformatted bit streams as well as UI or I frames. IM ends when a station accepts, actions, and acknowledges a mode-setting command.

TRANSMISSION FORMATS. Three control field formats are used to perform information transfer, basic supervisory control functions, and special or infrequent control functions.

The information (I) format is used to perform an information transfer. It is the only format which may contain an information field. The functions of sequence counts and poll/final bit are independent; that is, each frame has a transmit/send sequence count. The receive sequence count may or may not acknowledge additional frames at the receiving station, and the P/F bit may or may not be set to one.

The supervisory (S) format is used to perform link supervisory control functions such as to acknowledge information frames, to request retransmission of information frames, or to indicate temporary interruption of receive capability.

The unnumbered (U) format is used to provide additional link control functions. This format contains no sequence numbers. As a result, five modifier bit positions are available which allow definition of up to 32 additional supervisory functions.

TRANSMISSION PARAMETERS. The parameters associated with the three transmission formats are described in the following paragraphs.

Each information frame is sequentially numbered and may have the value zero through modulus minus one (where modulus is the modulus of the sequence numbers). Modulus equals eight for the unextended control field, and the sequence numbers cycle through the entire range.

The maximum number of sequentially numbered information format frames that the primary or secondary may have outstanding (i.e., unacknowledged) at any given time may never exceed one less than the modulus of the sequence numbers. This restriction is to prevent any ambiguity in the association of transmission frames with sequence numbers during normal operation and/or error recovery action.

Each station maintains a separate (independent) send sequence number N (S) and a receive sequence number N (R) on the information frames it sends and receives. Each secondary station then maintains an N (S) count on the information format frames it transmits to the primary, and an N (R) count on the information format frames it has correctly received from the primary. In the same manner, the primary maintains separate N (S) and N (R) counts for information format frames sent to and received from each secondary on the link.

POLL/FINAL (P/F) BIT. In CDCCP, each frame contains a poll/final (P/F) bit in the control field. This bit serves a function in both command and response frames. In command frames, it is referred to as the poll (P) bit. In response frames, it is referred to as the final (F) bit. In both cases, the bit is set to one.

The P bit is used to solicit a response or sequence of responses from a secondary or balanced station. When the poll frame contains the receive sequence number N (R), it constitutes a specific solicitation for the station to send I frames numbered N (S) = N (R) and following sequential I frames. In addition, it serves as an acknowledgment to the station for all I frames through send sequence number N (S) = N (R) -1. If the station has no I frames to send, it responds with a supervisory frame with the F bit set to one.

In NRM, the P bit is set to one in command frames to solicit response frames from the secondary station. In this mode the secondary station cannot transmit until a command frame with the P bit set to one is received. The primary station can also restrict the secondary station from transmitting I frames by sending a RNR supervisory frame with the P bit set to one or zero.

In ARM and ABM, I response frames can be transmitted by the secondary or balanced station on an asynchronous basis. The P bit set to one is used to solicit a response with the F bit set to one at the earliest opportunity. For example, if the remote primary or balanced station wants to get positive acknowledgment 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 or balanced station.

An F bit is used to acknowledge a P bit. A station may not send a final frame without prior receipt of a poll frame. In NRM, the secondary is required to set the F bit to one in the last frame of its response, which may consist of one or more frames. Following the transmission of a frame with the F bit set to one, the secondary must halt transmission until a command frame with a P bit set to one is received.

In ARM or ABM, the secondary or balanced station is required to transmit a response frame with the F bit set to one in response to a P bit, but is not required to halt transmission. The F bit shall be sent at the earliest opportunity as a function of link configuration, i.e., TWA or TWS. Since additional frames may be transmitted by a station in ARM following an F bit response, the F bit is not to be interpreted by the primary or balanced station as the end of transmission. It simply serves to finalize the response to the command frame with the P bit set.

Since P and F bits are exchanged on a one-for-one basis, and only one P bit can be outstanding at a time, the N (R) count of a frame containing a P or F bit set to one can be used to detect I-frame sequence errors. This capability is referred to as checkpointing, and can be used not only to detect frame sequence errors, but to indicate the frame sequence number to begin retransmission.

COMMANDS AND RESPONSES. The following paragraphs briefly describe each of the sets of commands and responses used in each of the three transmission formats.

The function of the information-transfer command and response is to transfer sequentially numbered frames containing an information field across a data link. The I-command and response-control field is illustrated in Figure 7. Bit 1 of the I control field is always zero and identifies this frame as an I frame. Bit 5 is the poll/final bit described previously.

The information-format-control field contains two sequence numbers. Bits 2, 3 and 4 comprise N (S), the send sequence count, which indicates the sequence number associated with this information frame. Bits 6, 7, and 8 comprise N (R), the receive sequence count, which indicates the sequence number of the next expected information-format frame to be received. The N (R) implicitly acknowledges correct receipt of information frames numbered up to N (R) -1.

Supervisory format commands and responses are used to perform basic link supervisory control functions such as acknowledgment, polling, and error recovery. Frames with the supervisory format do not contain an information field, and therefore do not increment the sequence counts at either the transmitter or the receiver. The supervisory-command and response control fields are also illustrated in Figure 7.

Bits 1 and 2 of the S control field identify the frame as an S frame. Bit 5 is the poll/final bit. Bits 6, 7, and 8 comprise the N (R), receive sequence count, which indicates the sequence number of the next expected information format frame to be received. It also implicitly acknowledges correct receipt of information frames numbered up to and including N (R) -1. Bits 3 and 4 of the S control field define the supervisory function and are encoded identically for both command and response frames.

The unnumbered (U) format commands and responses are used by the primary and secondary to extend the number of link supervisory functions. Frames transmitted with the unnumbered format do not increment the send sequence counts N (S) at either the transmitting or receiving station. Five modifier bits are defined which allow up to 32 additional supervisory functions. Of these, ten are defined. The remaining combinations are reserved for future assignment. The unnumbered command and response control field is also illustrated in Figure 7. Bits 1 and 2 of the U format control field identify the frame as a U frame. Bit 5 is the poll/final bit. Bits 3, 4, 6, 7 and 8 are the modifier bits and are encoded as illustrated in Figure 7.

The following paragraphs delineate each of the commands and responses beginning with the supervisory frames. The complete repertoire is illustrated in Figure 8.

The receive-ready (RR) supervisory frame is used by a station to indicate that it is ready to receive an information frame, and to acknowledge previously received information frames numbered up to and including N (R) -1. A primary or balanced station may use the RR command with the poll bit set to one to solicit responses from (i.e., "poll") secondary or balanced stations.

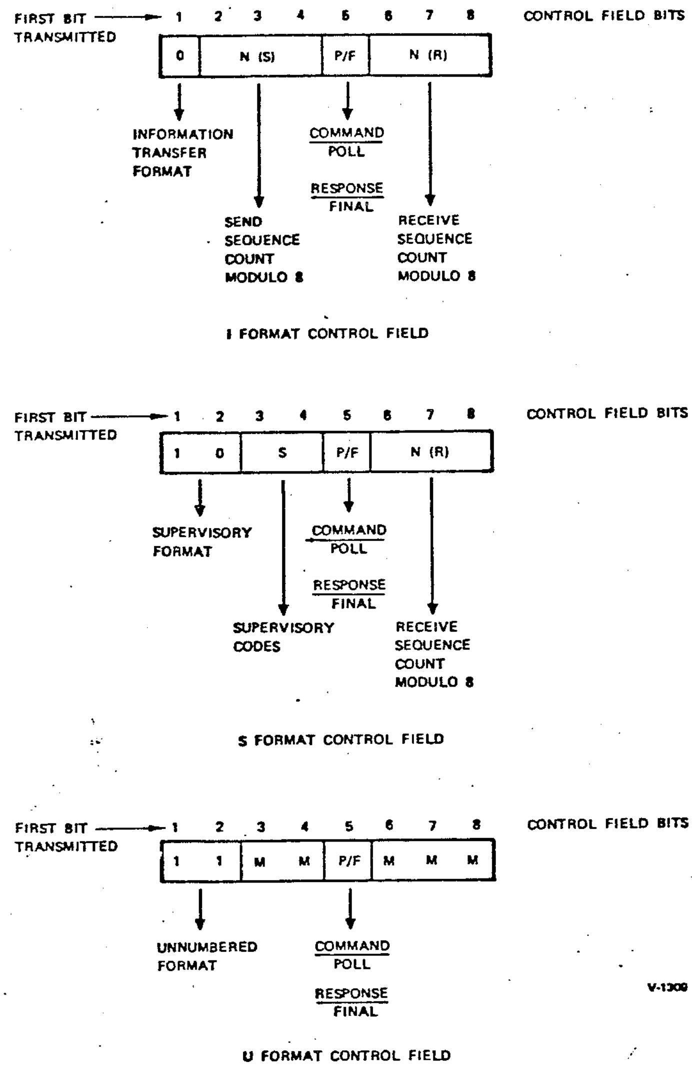

Figure 7. Control Field Formats

The receive-not-ready (RNR) supervisory frame is used by a station to indicate temporary inability to accept additional, incoming information-format frames. Information-format frames numbered up to and including N (R) -1 are acknowledged; information frame N (R) and any subsequent information format frames received, if any, are not acknowledged. A station receiving an RNR frame when in the process of transmitting (i.e., a FDX station) is to stop transmitting at the earliest possible time by completing or aborting the frame in process.

The reject (REJ) supervisory frame is used by a station to request retransmission of information-format frames starting with the frame numbered N (R). Information-format frames numbered N (R) -1 and below are acknowledged. Additional I frames pending initial transmission may be transmitted following the retransmitted I frame(s).

The selective-reject (SREJ) supervisory frame is used by a station to request retransmission of the single information frame numbered N (R). Information format frames numbered through N (R) -1 and below are acknowledged. Once a SREJ has been transmitted, the only I frames accepted are those which are numbered contiguously and in sequence following the I frame requested, and the specific retransmitted I frame indicated by the N (R) in the SREJ command/response.

FORMAT	CONTROL FIELD BIT ENCODING 1 2 3 4 5 6 7 8	COMMANDS	RESPONSES
INFORMATION	0 –N(S) – • –N(R) –	I – INFORMATION	I – INFORMATION
SUPERVISORY	1 0 0 0 • N(R)	RR – RECEIVE READY	RR – RECEIVE READY
	1 0 0 1 • N(R)	REJ – REJECT	REJ – REJECT
	1 0 1 0 • N(R)	RNR – RECEIVE NOT READY	RNR – RECEIVE NOT READY
	1 0 1 1 • N(R)	SREJ – SELECTIVE REJECT	SREJ – SELECTIVE REJECT
UNNUMBERED	1 1 0 0 • 0 0 0	UI – UNNUMBERED INFORMATION	UI – UNNUMBERED INFORMATION
	1 1 0 0 • 0 0 1	SNRM – SET NORMAL RESPONSE MODE	
	1 1 0 0 • 0 1 0	DISC – DISCONNECT	RD – REQUEST DISCONNECT
	1 1 0 0 • 1 0 0	UP – UNNUMBERED POLL	
	1 1 0 0 • 1 1 0		UA – UNNUMBERED ACKNOWLEDGE
	1 1 0 1 • 0 0 0	USER 0 –	USER 0 –
	1 1 0 1 • 0 0 1	USER 1	USER 1
	1 1 0 1 • 0 1 0	USER 2	USER 2
	1 1 0 1 • 0 1 1	USER 3	USER 3
	1 1 1 0 • 0 0 0	SIM – SET INITIALIZATION MODE	RIM – REQUEST INITIALIZATION MODE
	1 1 1 0 • 0 0 1	RSPR – RESPONSE REJECT	CMDR – COMMAND REJECT
	1 1 1 1 • 0 0 0	SARM – SET ASYNC RESPONSE MODE	DM – DISCONNECT MODE
	1 1 1 1 • 0 1 0	SARME – SET ARM EXTENDED MODE	
	1 1 1 1 • 0 1 1	SNRME – SET NRM EXTENDED MODE	
	1 1 1 1 • 1 0 0	SABM – SET ASYNC BALANCED MODE	
	1 1 1 1 • 1 0 1	XID – EXCHANGE IDENTIFICATION	XID – EXCHANGE IDENTIFICATION
	1 1 1 1 • 1 1 0	SABME – SET ABM EXTENDED MODE	

• = P/F

Figure 8. CDCCP Command/Response Repertoire

Continuing into the unnumbered format commands and responses, we can logically divide these into several groups: mode setting, other-mode-related, recovery, and miscellaneous.

In the mode setting category we have the Set Normal Mode (SNRM) command which is used to place the addressed secondary station in the normal response mode (NRM) where all control fields are one octet in length. No information field is permitted with the SNRM command.

The Set Asynchronous Response Mode (SARM) command is used to place the addressed secondary station in an asynchronous response mode (ARM) where all control fields are one octet in length. No information field is permitted with the SARM command.

The Set Asynchronous Balanced Mode (SABM) command is used to place the addressed balanced station in ABM where all control fields are one octet in length. No information field is permitted with the SABM command.

An equivalent set of three extended mode commands exist to place the addressed station in the extended mode, where all control fields are two octets in length.

Other-mode-related commands and responses include the disconnect (DISC) command, which is used to perform a logical disconnect (i.e., inform the receiving station that the transmitting station is suspending operation with the receiving secondary or balanced station). In a switched network operation, the DISC command may also be used to initiate a physical disconnect at a lower level.

The Set Initialization Mode (SIM) command is used to initiate system-specified, link-level initialization procedures at the secondary or balanced station. No information field is permitted with this command.

The request disconnect (RD) response is used to indicate to the remote primary or balanced station that the transmitting secondary or balanced station wishes to be placed in disconnected mode (NDM or ADM) or, in switched networks, to request a physical disconnect operation.

The request for initialization mode (RIM) is transmitted by a station to notify the primary or balanced station of the need for a SIM command. The receipt of any command except a SIM or DISC will cause the station to repeat the RIM, i.e., no command transmissions are accepted (except to detect a respond opportunity) until the RIM condition is reset by receipt of a RIM or DISC. No information field is permitted with the RIM response.

The unnumbered acknowledge (UA) response is used to acknowledge receipt and acceptance of the SNRM, SARM, SABM, SNRME, SARME, SABME, SIM, DISC, and RSPR unnumbered commands. The UA response is transmitted in the normal or extended control field format as directed by the received unnumbered command. No information (I) field is permitted with the UA response.

The disconnected mode (DM) response is used to report a non-operational status where the secondary or balanced station is logically disconnected from the link, i.e., the station is, per system definition, in NDM or ADM.

The DM response is sent by a station in NDM or ADM (or equivalent) to request the remote primary or balanced station to issue a set-mode command. If sent in response to the reception of a set-mode command, it is to inform the addressed primary or balanced station that the transmitting secondary or balanced station is still in NDM/ADM and cannot action the set-mode command.

There are two command/responses used primarily for recovery purposes. The first is the command reject (CMDR) response which is used to report that an exception condition, not recoverable by retransmission of the identical frame, resulted from the receipt of an error-free frame from the primary or balanced station. A status field is returned with a CMDR to provide the reason for issuance of the CMDR. This status field immediately follows the basic control field in the frame.

The other is the response reject (RSPR) command which is used in an identical manner to report exception conditions not recoverable by link level at the secondary or balanced station.

The group of miscellaneous commands and responses includes the following.

The Unnumbered Information (UI) command and response is used to transfer non-sequence numbered information fields across a link. This can be useful for higher level status, link initialization data, etc. Neither send or receive state variables S and R are impacted by the transmission or reception of a UI frame. Since reception of UI frames is not sequence-number verified, the I field may be lost or duplicated if a link exception condition occurs during transmission of the UI frame.

The Unnumbered Poll (UP) command is used to solicit response frames from a single secondary or balanced station (individual poll), or from a group of secondary stations (group poll), by establishing a logical operational condition that exists at each addressed station for one respond opportunity.

The Exchange Identification (XID) command is used to cause the addressed station to report its station identification and, optionally, to provide the station identification of the transmitting station to the remote station. An information field is optional with the XID command; but if used, will contain the station ID of the transmitting primary or balanced station. The XID command may use the global address if the unique address of the secondary or balanced station is unknown. The XID response is used to reply to an XID command. Again, an information field containing the station ID of the transmitting secondary or balanced station is optional.

Four commands and four responses (User 0 Through User 3) are specified to permit the definition of special, system-dependent functions not having general applicability. Any such functions are beyond the scope of CDCCP standards.

CLASSES OF PROCEDURES

Procedural differences among applications, based on overall system considerations such as network configuration, recovery procedures, terminal sophistication, etc., are accommodated in CDCCP by defining various classes of procedure. These classes combine modes of operation (ARM, NRM and ABM), commands and responses and exception recovery procedures. These combinations are supplemented by a set of optional functional extensions providing additional capability for certain applications. Each class forms an implementation subset of CDCCP procedures.

The six defined classes of procedures (see Table 1) are composed of:

- three types of stations: primary stations, secondary stations and balanced stations;
- two types of configurations: unbalanced (for primary and secondary stations) and balanced (for balanced stations);
- two types of transmission response: normal and asynchronous; and
- two sizes of modulus: 8 and 128.

TABLE 1. CLASSES OF PROCEDURE

Designation	Description
UAB	Unbalanced operation, asynchronous response mode, Modulo 8
UAE	Unbalanced operation, asynchronous response mode, Modulo 128
UNB	Unbalanced operation, normal response mode, Modulo 8
UNE	Unbalanced operation, normal response mode, Modulo 128
BAB	Balanced operation, asynchronous transmission, Modulo 8
BAE	Balanced operation, asynchronous transmission, Modulo 128

Figure 9 summarizes the command and response usage of the four unbalanced and two balanced classes currently defined. It also illustrates the eight currently defined functional extensions.

There are eight optional functional extensions to the six classes of procedure. These functional extensions are achieved by the addition of commands and responses to the basic repertoire of any given class of procedure.

Option 1: Switched Network Operation - This option provides the capability in switched network operation to exchange identification of stations and to request logical disconnection from the link.

Option 2: Two-Way Simultaneous Operation - This operation provides the ability for more timely reporting of I-frame sequence errors in two-way simultaneous operation.

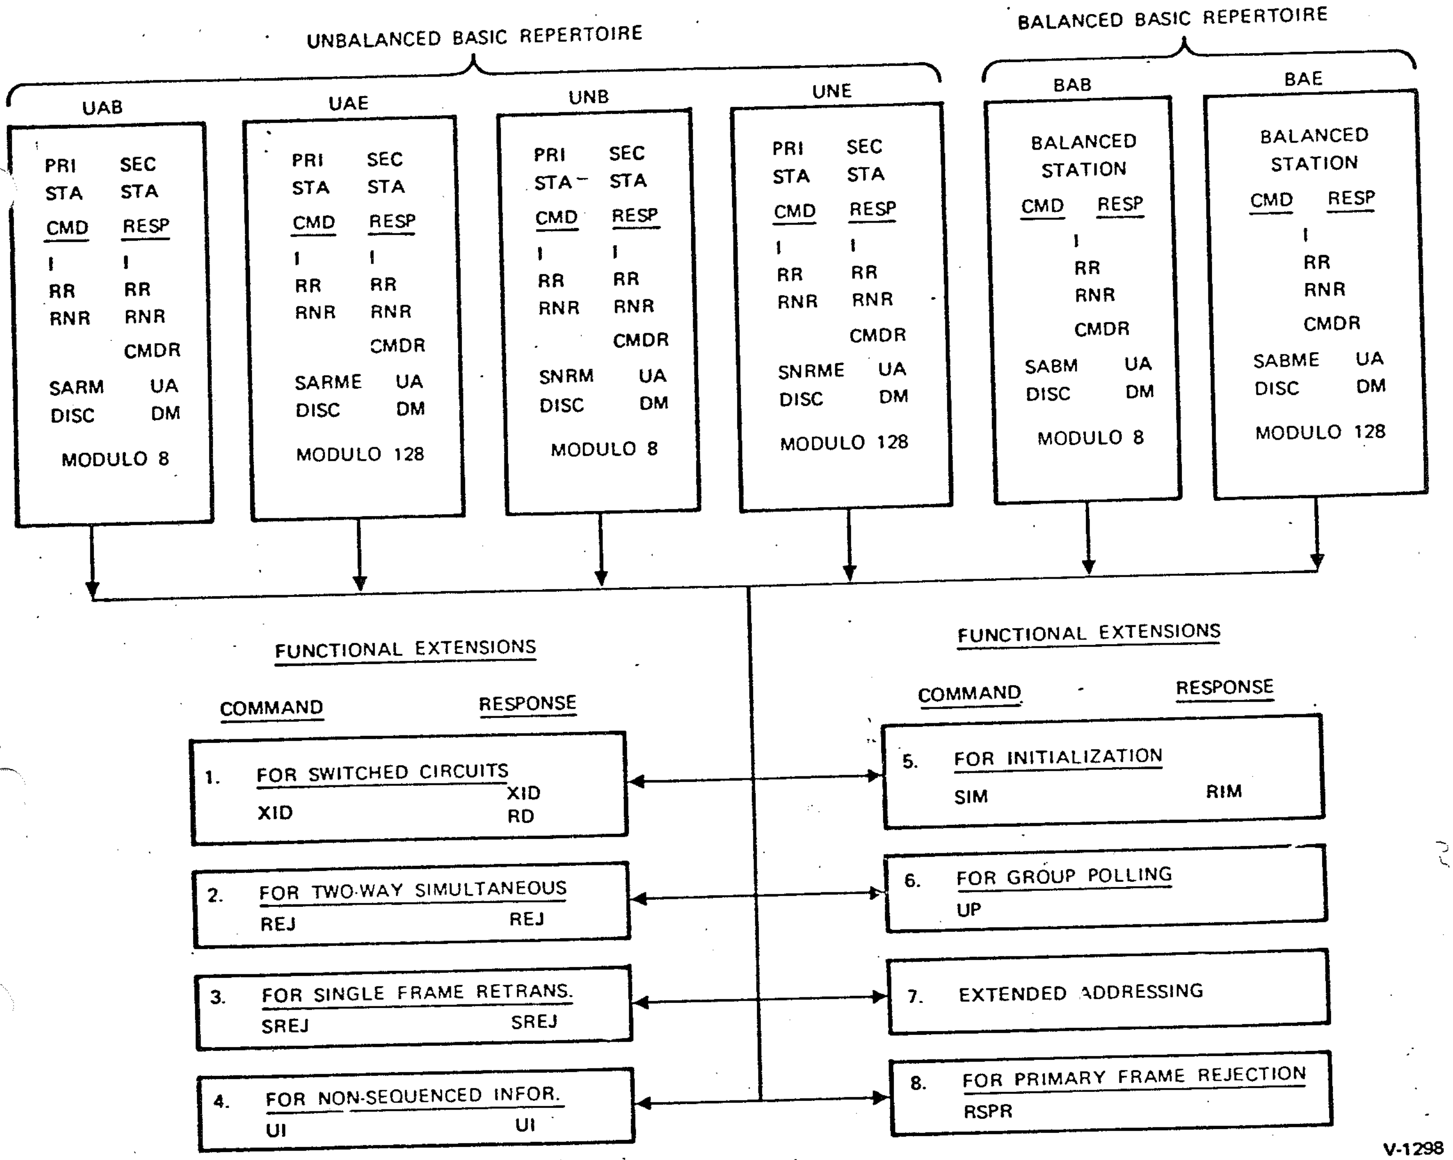

Figure 9. Command/Response Implementation by Class

Option 3: Single-Frame Retransmission - This option provides the ability for more efficient recovery from I-frame sequence errors by requesting retransmission of a single frame.

Option 4: Non-Sequenced Information Frame - This option provides the ability to exchange information fields without impacting the I-frame sequence counts.

Option 5: Initialization - This option provides primary/balanced ability to initialize remote stations and secondary/balanced ability to request initialization.

Option 6: Group Polling - This option provides a primary with the ability to perform group polling.

Option 7: Extended Addressing - This option provides a second octet for addressing and thus provides addressing for more than 254 secondaries on a single link.

Option 8: Primary Frame Rejection - This option provides the primary with the ability to report a frame rejection condition to a secondary.

Since classes of procedure are still in their formative stage, the standardization picture remains somewhat cloudy. ISO, in particular, is having difficulty in reaching agreement on the precise definition of classes and on methods of codifying them. It is expected that the final ISO classes will be very close to the present CDCCP and ANSI standards.

IMPLEMENTATION

The subject of compatibility between the various bit-oriented protocols was mentioned earlier. Since this subject is especially important to the user, the chart in Figure 10 has been prepared. This illustrates the complete set of commands and responses now defined and lists the ones being implemented for each protocol. The information presented is, of course, subject to change but represents the best data available. This chart indicates a high degree of basic compatibility between the standards of the various standards bodies, manufacturers, and user's organizations.

Given this basic compatibility, it remains for the user to carefully determine his requirements in terms of a "class

ADCCP		HDLC		SDLC		CDCP		BOLD		BDLC		SNAP/X25	
CMD	RES	CMD	RES	CMD	RES	CMD	RES	CMD	RES	CMD	RES	CMD	RES
I	I	I	I	I	I	I	I	I	I	I	I	I	
RR	RR	RR	RR	RR	RR	RR	RR	RR	RR	RR	RR		RR
REJ	REJ	REJ	REJ	REJ	REJ	REJ	REJ	REJ	REJ	REJ	REJ		REJ
RNR	RNR	RNR	RNR	RNR	RNR	RNR	RNR	RNR	RNR	RNR	RNR		RNR
SREJ	SREJ	SREJ	SREJ			SREJ	SREJ	SREJ	SREJ	SREJ	SREJ		
UI	UI	*	*	NSI	NSI	UI	UI	NSI	NSI				
SNRM		SNRM		SNRM		SNRM		SNRM		SNRM			
DISC	RD	DISC	*	DISC		DISC	RD	DISC		DISC		DISC	
UP		*		ORP		UP		ORP					
	UA		UA		NSA		UA		NSA		UA		UA
USR (4)	USR (4)					USR (4)	USR (4)						
SIM	RIM	*	*	SIM	RIM	SIM	RIM	SIM	RIM				
FRMR	FRMR	*	CMDR		CMDR	RSPR	CMDR	RSPR	CMDR	RSPR	CMDR		CMDR
SARM	DM	SARM	*		ROL	SARM	DM	SARM	ROL	SARM		SARM	
SARME		SARME				SARME		SARME		SARME			
SNRME		SNRME				SNRME		SNRME		SNRME			
XID	XID	*	*			XID	XID						
SABM		*				SABM							
SABME		*				SABME							

* = PENDING APPROVAL

V-1312

Figure 10. Command/Response Repertoire of Selected Standards

of procedure" to be used. This will define the operational modes and elements of procedure to be used. Following this it will be necessary to generate a system specification for the specific application. This document will identify and quantify many variables necessary to achieve successful on-line operation. It is here that the impact of lower and higher levels of the communications hierarchy will be reviewed and specified.

The list in Figure 11 typifies the type of operational parameters which must be bilaterally reviewed and defined in order to successfully specify, and ultimately implement, any bit-oriented link protocol.

While it is evident that the new bit-oriented protocols are rapidly reaching maturity, much work remains to be done, especially in the standardization of successively higher levels of the data communications hierarchy.

Through its representation on ANSI and through liaison with other groups, Control Data is closely following developments leading to standardization of device control and message formats. These functions which would be contained within the I field of CDCCP are, of course, of major interest to CDC and its users. Other areas being pursued include the emerging application possibilities of packet switching, and public and private data networks.

Control Data Corporation, as part of its total services concept, is dedicated to provide cost-effective and efficient solutions for the user's data communications problems. The development of bit-oriented protocols is but one example of this commitment.

- CLASS OF PROCEDURE TO BE USED
- FUNCTIONAL EXTENSIONS TO BE IMPLEMENTED
- TYPE OF STATION: PRIMARY, SECONDARY, OR BALANCED
- BASIC OR EXTENDED ADDRESSING
- BASIC OR EXTENDED CONTROL FIELD
- LOWER LEVEL LINK CHARACTERISTICS
- NUMBER OF FLAGS BETWEEN FRAMES
- TIMER VALUES
- RETRY COUNTERS
- STATION ADDRESSES
- CONTENTS OF XID FRAME
- USE OF AND CONTENTS OF UI FRAMES
- INITIALIZATION PROCEDURES
- DISCONNECT PROCEDURES
- MAXIMUM LENGTH OF INFORMATION FIELD
- ET AL

V-1310

Figure 11. Typical Link Control Parameters Subject to Bilateral Agreement