

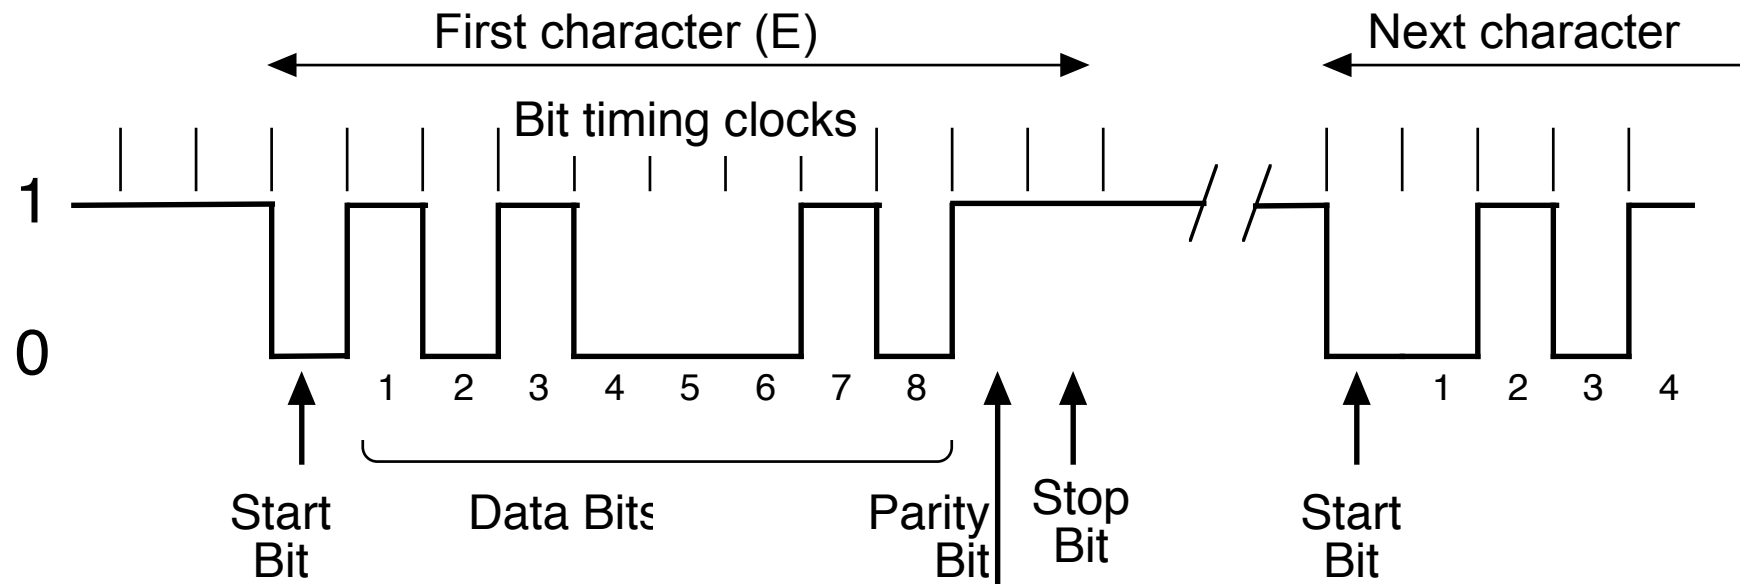
# Data Framing

Most communication sends bit-serial data. How do we isolate the characters?

## Asynchronous Communication.

- This is a very old technique, developed about 100 years ago for automatic teletypewriter equipment.
  - It is meant for slower speed communication, where the send and receive bit timing clocks may differ by  $\pm 5\%$ .
1. The single communication line rests in a “1” – “sender is connected”
  2. A character starts with a “Start bit”, always 0, 1 bit time
  3. The data bits (1–8) are sent in order, LSB first
  4. There may be a parity bit, to check transmission accuracy
  5. The character ends with a “Stop bit”, always 1, and 1, 1.5 or 2 bits long.

# Character 'E' 0x45, 01000101



- Character is “synchronised” from leading edge of Start Bit
- Hardware device is UART  
“Universal Asynchronous Receiver Transmitter”
- UART can select —

Data Bits	5,6,7,8 data bits
parity	Odd, Even, 0, 1, or none
Stop Bits	1, (rarely 1.5 or 2)

# UART Operation


1. A UART is usually given a clock at  $16\times$  the desired bit rate.
2. The clock is “counted down” by 16 times for transmission
3. For reception, look for a change  $1\rightarrow 0$ ; this is the start of the Start Bit
4. Then count 8 clocks; this should be the middle of the Start bit
5. Then count off successive 16 clocks for each data bit.
6. Read the parity bit; is it correct? (if not then Parity Error)
7. Read the Stop Bit; is it 1? (if not then Framing Error)
8. “Idle” until the next Start Bit ( $8 \leq \text{delay} < \infty$  clocks)

# Asynchronous characteristics and usage

1. Originally used for electromechanical equipment, hence
  - 1.5 stop bits (with 5-bit data, 75 bps),
  - 2 stop bits (ASCII 7-bit+parity, 110 bps)Now use only 1 stop bit with modern electronic communication
2. Asynchronous communication is relatively inefficient, at least 25% overhead from start/stop bits
3. Used mainly where there is no shared clock from sender to receiver, OR
4. with very short messages, perhaps just 1 character.
5. Data is now either 7 or 8 bits, and parity may be present or absent.

## A Simple Asynchronous Block-mode protocol

This protocol was used on Burroughs terminals, to about 38,400 bps, during the 1970's.



- The connection was like Ethernet (but much slower)
- Each station had an address, a single ASCII character
- The CPU cycled round all stations — POL (do you have a message?) or SEL (can you receive a message?)
- Characters were 7-bit ASCII, with even parity,
- Messages had an even parity longitudinal parity (BCC – Block Check Character), including all characters from address to ETX.

# ASCII Control characters

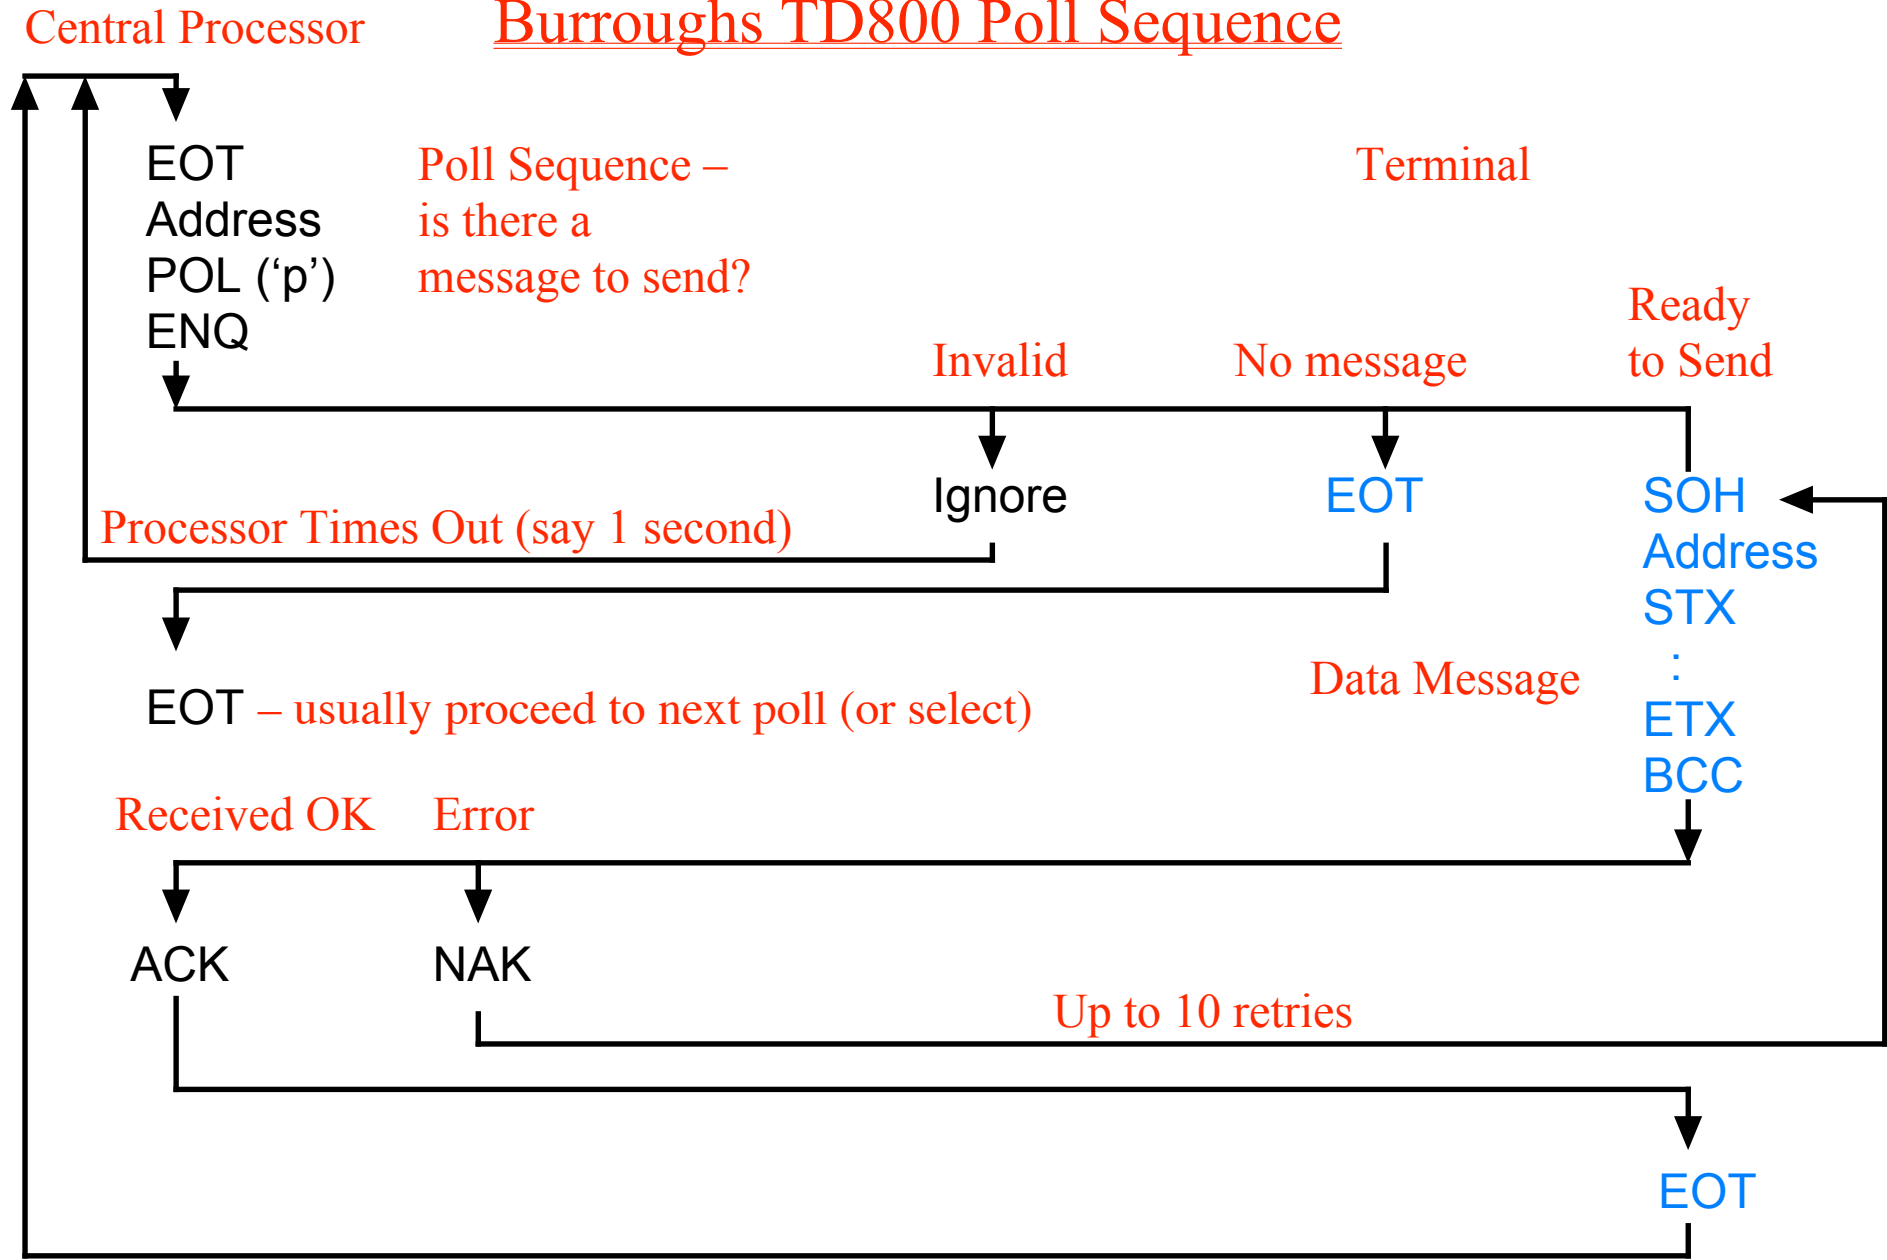
The protocol used the following ASCII control (non-data) characters

0x01	SOH	Start of Header	Introduces message header
0x02	STX	Start of Text	
0x03	ETX	End of Text	
0x04	EOT	End of Transmission	no more data to send
0x05	ENQ	Enquire	terminal must respond
0x06	ACK	Acknowledge	message received correctly
0x15	NAK	Negative ACK	error in message
	pol	p	“is there a message to send?”
	sel	q	“can you receive a message?”

SOH | header | STX | .....data..... | ETX | BCC

The final Block Check Character (BCC) is data dependent and can be anything.

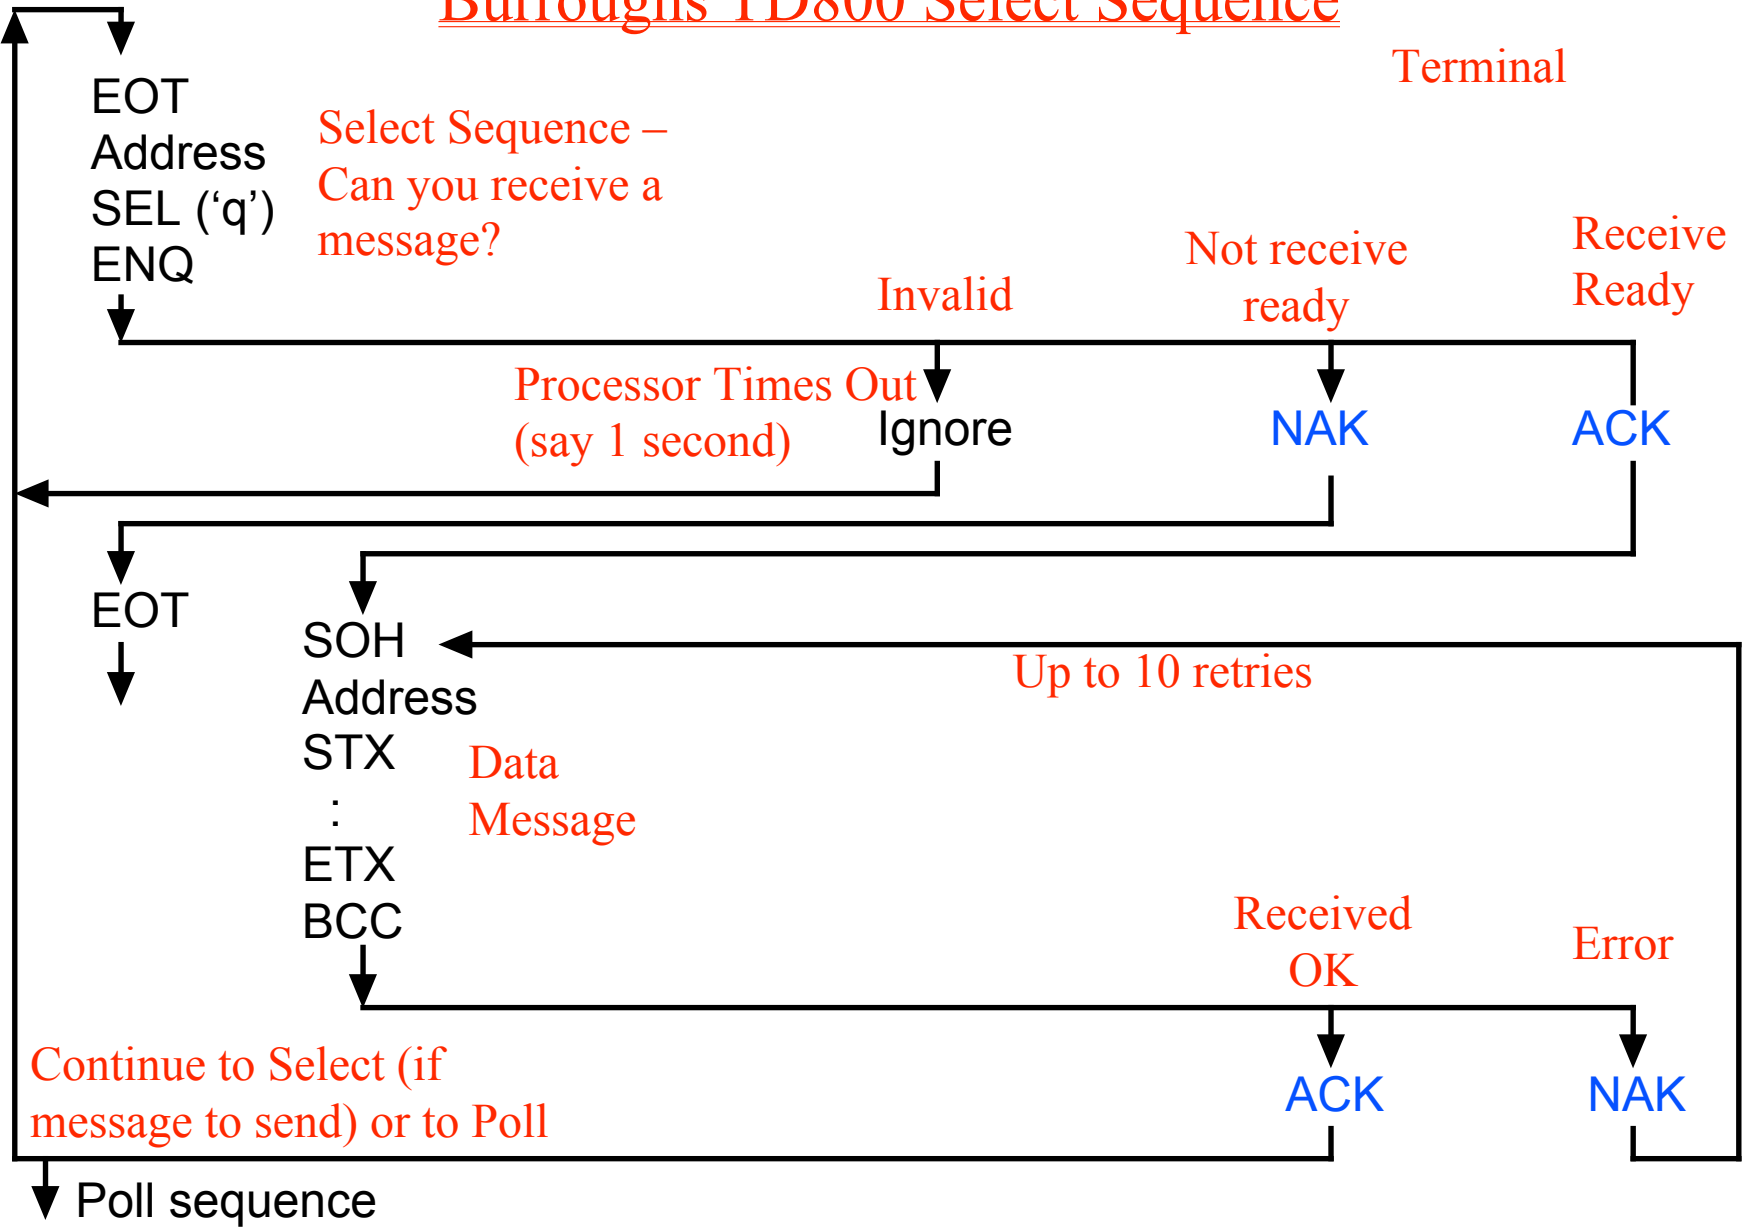
# Burroughs TD800 Poll Sequence



Central Processor

# Burroughs TD800 Select Sequence

Terminal





# Example of Parity Calculation

- The character parity (most-sig bit) is set to give an even number of bits in the character ('a' =  $61_{16} = 110\ 0001$  has 3 1-bits and becomes 1110 0001).
- The message parity is the Exclusive-OR ( $\oplus$ ) of the **7-bit** message character codes (the BCC of the first 4 octets is  $01 \oplus 31 \oplus 02 \oplus 44 = 76$ ).
- Data message is **Data Example**
- Both character and longitudinal parity are even

Character	SOH	1	STX	D	a	t	a		E	x	a	m	p	l	e	ETX	BCC
Hex value	01	31	02	44	61	74	61	20	45	78	61	6D	70	6C	65	03	69
with parity	81	B1	82	44	E1	74	E1	A0	C5	78	E1	ED	F0	6C	65	03	69

- The receiver checks both character parity and overall message parity and requests retransmission if either fails.
- The receiver also checks for the correct structure, eg **SOH adr STX** and ignores bad messages
- **Character parity, including BCC, is added on transmission.**

- This example shows many features of data communications protocols —
- Well-defined message formats, with extensive checking
- Well-defined question-response sequence, such as —

Do you have a message? (poll)	
	Yes, and here it is (the data)
I have it correctly (ACK)	
	I have no more messages (EOT)

- Retry in case of error, up to some agreed limit (say 10 times).  
If still fails, report an error to higher software layers.
- Control station can “timeout”, so it can recover if nothing is received.  
*A timeout is an essential part of any useful protocol*
- This is a “stop and wait” protocol – it sends a message and waits for a reply.

# Synchronous Communication

1. Introduced for computers and semi-intelligent terminals which send blocks of data.
2. Used with modems which convey bit-timing and tell the receiver when to accept a bit — the receiver does not have to infer bits.
3. All data is sent as blocks or frames, each preceded by 2 or 3 special SYN (Synchronisation) codes.

ASCII	0x16	01101000	(lsb first)
EBCDIC	0x32	01001100	(lsb first)
4. Receiver searches bit by bit for the pattern SYN SYN (in binary ASCII 0001011000010110).
5. It then delivers a character each 8 bits afterwards (seldom use character parity).

# Binary Synchronous Control (BSC) protocol

1. Used by IBM at about the same time as Burroughs TD800 protocol, but for similar purposes.
2. Intended for 8-bit EBCDIC, not 7-bit ASCII, so no character parity and need a better message parity. (Use 16-bit cyclic redundancy check.)
3. Can be used for true binary data, with “byte stuffing”.

The simplest message format is for normal text



This is just like the TD800 protocol, except that it uses EBCDIC and has two leading SYN characters (but TD800 could also use synchronous ASCII).

- A very long message may be broken into blocks, each ending with ETB (End of Transmission Block), and last block ending with ETX.

# Bisync Transparent mode

- What happens if the data might contain an ETX character? Text from terminal certainly will not, but we might want to send pure binary data.
- The answer is to use “byte stuffing”, adding DLE (Data Link Escape) codes in appropriate places.



- Transparent mode is signalled by a **DLE STX** in place of the beginning STX and a message is ended by the pair **DLE ETX**.
- But in case the message contains a genuine DLE ETX sequence, etc, replace any data DLE by DLE DLE and treat the following character as data, (unless it is another DLE).

# The BSC protocol

- The BSC protocol is similar in principle to the TD800 protocol, but differs of course in detail.
- Both are stop and wait protocols.
- Both have a controlling, or Primary, station with one or more Secondary stations. The Primary station always initiates an operation and a Secondary always responds.
- Both protocols are now obsolete and were replaced by the HDLC (or SDLC) protocols to be described (which is itself largely obsolete as well!).

# Synchronous Data Link Control (SDLC), or High-level Data Link Control (HDLC)

- These were introduced in conjunction with the ISO Open System Interconnection (OSI) model, which included X.25.
- Although HDLC itself is little used now, many of its details have carried over into present protocols.

The HDLC frame format is

width (bits)	8	8 or 16	8 or 16	8 x N bits	16 or 32	8
Description	Flag	Address	Control	.....data.....	FCS	Flag

- The address and control sizes are agreed at set-up time
- Apart from the flags there is little framing or other overhead

# Bit Stuffing

- The Flags which surround an HDLC frame are always the 8-bit pattern 01111110, six consecutive 1s, surrounded by 0's.
- Binary data is handled by “bit stuffing” with the rules –
  1. On transmission, always force a 0 after 5 consecutive 1's (ie stuff in a extra 0 bit after the pattern 11111.)
  2. On reception, ignore any 0 bit which is preceded by 5 1's. In other words 111110 → 11111.
  3. Flags are inserted after bit-stuffing on transmission, and detected before “un-stuffing” on reception.
- On random data, we can expect the pattern ...011111... to occur in 1 out of 64 bits, so on average we expand by 1 extra bit for 64 data bits, or about 1.5%.
- Note that a complete HDLC frame (including flags!) can be tunnelled within another HDLC frame



# HDLC Control fields

- The control field is usually 8 bits, in one of three formats
- An Information Frame starts with a 0 (in the LSB position!) and conveys normal data. The fields will be discussed later

1	3	1	3
0	N(S)	P/F	N(R)

- A Supervisory Frame is used for acknowledgements, including not-ready indications

1	1	2	1	3
1	0	S	P/F	N(R)

- An Unnumbered Frame is used for various control functions

1	1	2	1	3
1	1	M	P/F	M

# HDLC Control Fields

- P/F (Poll/Final).

If sent by a Primary station it is a “poll” bit and demands a response.

If sent by the Secondary (Final) it signals the last in a sequence of frames

- N(S) and N(R) are send and receive sequence numbers, used in flow control.

- S is a function field –

RR	Receive Ready	ready to receive frame N(R)
----	---------------	-----------------------------

REJ	Reject	reject frame N(R); send it and all following
-----	--------	--

RNR	Receive Not Ready	received frame N(R), but send no more
-----	-------------------	---------------------------------------


SREJ	Selective Reject	resend only frame N(R)
------	------------------	------------------------

- M is also a function field, mostly for link control, connection and disconnection

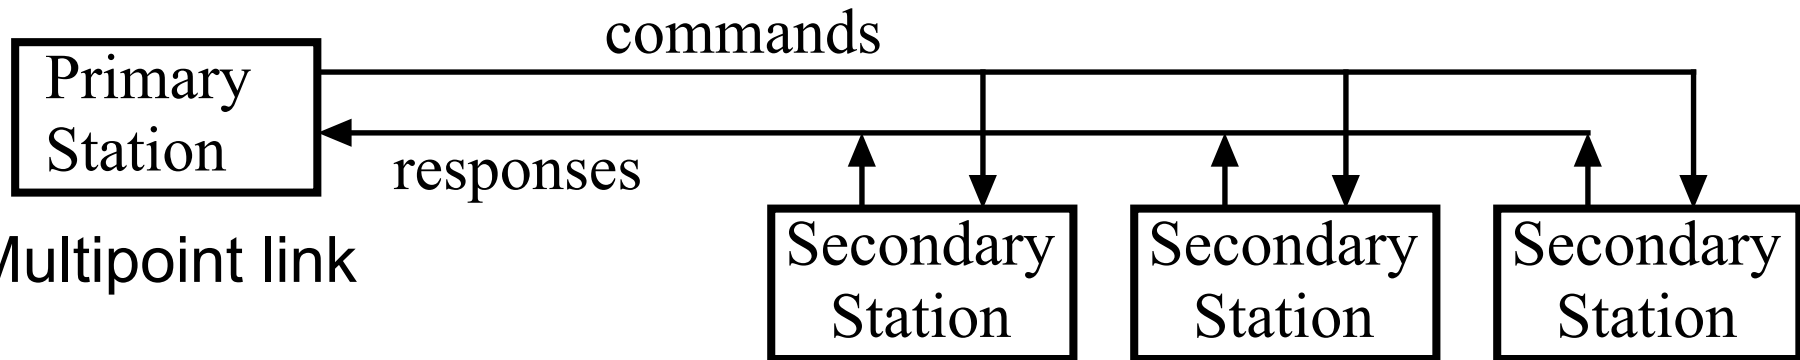
- The final CRC (Cyclic Redundancy Check) is a 16-bit checksum (possibly 32-bit) as discussed later.

# HDLC Configurations


- A Primary station can send commands; a Secondary station responses



Point-to-point link



Multipoint link



Point-to-point link between combined stations

# HDLC communication modes

1. In **Normal Response Mode (NRM)** the primary stations controls the operation, just as in the earlier protocols.
2. **Asynchronous Response Mode (ARM)** is like NRM, but allows a secondary to send data or control information without explicit permission, but it cannot send commands. Both NRM and ARM are intended to connect dumb terminals to a computer.
3. **Asynchronous balanced mode (ABM)** allows both stations to send both commands and responses. It is the usual method between computers.
4. All modes have an extended mode, with 7 bit sequence numbers
5. There are **unnumbered commands** to set these modes

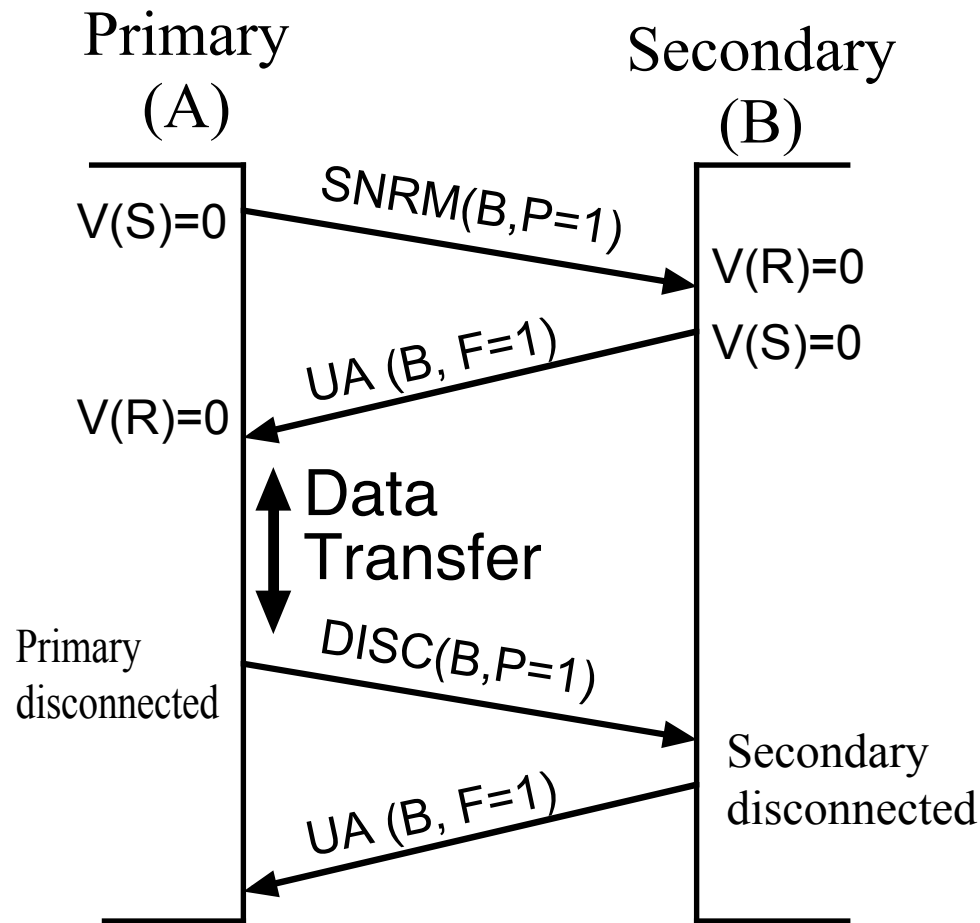
SARM	SARME	Set Asynchronous Response Mode (E)
SNRM	SNRME	Set Normal Response Mode (E)
SABM	SABME	Set Balanced Response Mode (E)

## Other important commands and responses

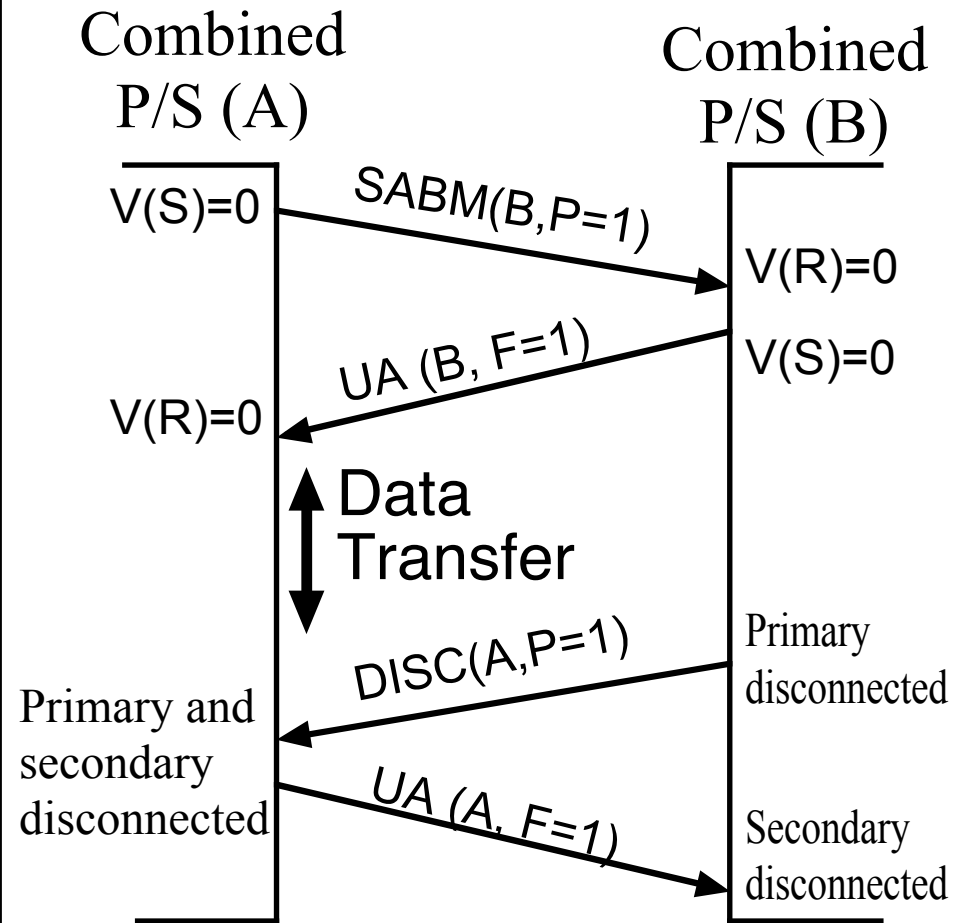
CDMR	Command Reject	Unacceptable command
DISC	Disconnect	Disconnect, or break, connection
DM	Disconnect Mode	Station is disconnected
FRMR	Frame Reject	
RSET	Reset	Reset to a known state
UA	Unnumbered acknowledge	General command acknowledge
UI	Unnumbered Information	Not part of a sequence of frames

# Link control

Two examples of link connection and disconnection



Normal Response Mode or Multidrop



Asynchronous Balanced Mode

# Flow Control

Flow control allows a receiving station to control transmissions so that it (the receiver) is not overwhelmed. There are several types of flow control —

1. Stop-Go, or ARQ, or XON/XOFF.

The receiver completely starts or completely stops transmission, often on a single message basis.

2. Buffer Control (Window control).

The receiver notifies the sender of how many buffers are available to receive data. Increasing the number of buffers often allows higher user data rates (if link latency is important). Decreasing the number of buffers reduces the traffic; zero buffers implies traffic stopped.

3. Rate Control.

Allow packets to be sent at some maximum rate — good for for some long-delay networks. Not treated here.

# Sliding Window flow control

- **Stop-and-Wait** protocols (ARQ protocols) may give slow data transfers because every message involves several **send-wait-receive-wait** sequences. Over long distances the waits can be large and seriously reduce the useful network speed.  
But they may be necessary in a *half-duplex* connection.
- With a **full-duplex** connection both stations can send and receive simultaneously and we may send several messages while waiting for the first to be acknowledged, leading to a *sliding window* protocol.
- With a sliding window protocol every message has a sequence number, incremented for successive messages, so that the transmitter and receiver can agree on which messages have been sent and which received.
- Acknowledgements are, if possible, “piggy-backed” onto reverse traffic (using the N(R) field of the message). Otherwise use RR response.




Start with messages 0–7 sent, and messages 0–2 acknowledged. If message 3 is acknowledged, the window can advance and allow message 8 to be sent.

These (earlier) messages have been sent and received

These messages are outstanding  
“**window**”

These (later) messages are still to be sent



These (earlier) messages have been sent and received

These messages are outstanding, but the **window** has moved

These (later) messages are still to be sent

(Earlier) message 7 sent;  
message 2 received

**Window size = 5  
messages**

(Later) message 9 sent;  
message 4 received

After message 4 has been acknowledged, the window can advance and allow another two messages to be sent

- HDLC usually works with a 3-bit sequence number, or frame numbers 0–7, which increments modulo-8.
- The sender has a variable  $V(S)$ , the number for the next frame to be sent.
- The receiver has a variable  $V(R)$ , for the next frame which it expects.

The HDLC sequence numbers can be increased to 7 bits (window size = 127) to handle connections with long delays, such as satellite links.

**NOTE :** The end-nodes maintain variables  $V(S)$  and  $V(R)$ ; the values of these variables appear as the numbers  $N(S)$  and  $N(R)$  in the messages.

Data Message (seq = 1)  
Data Message (seq = 2)  
Data Message (seq = 3)

Window size  $w = 3$

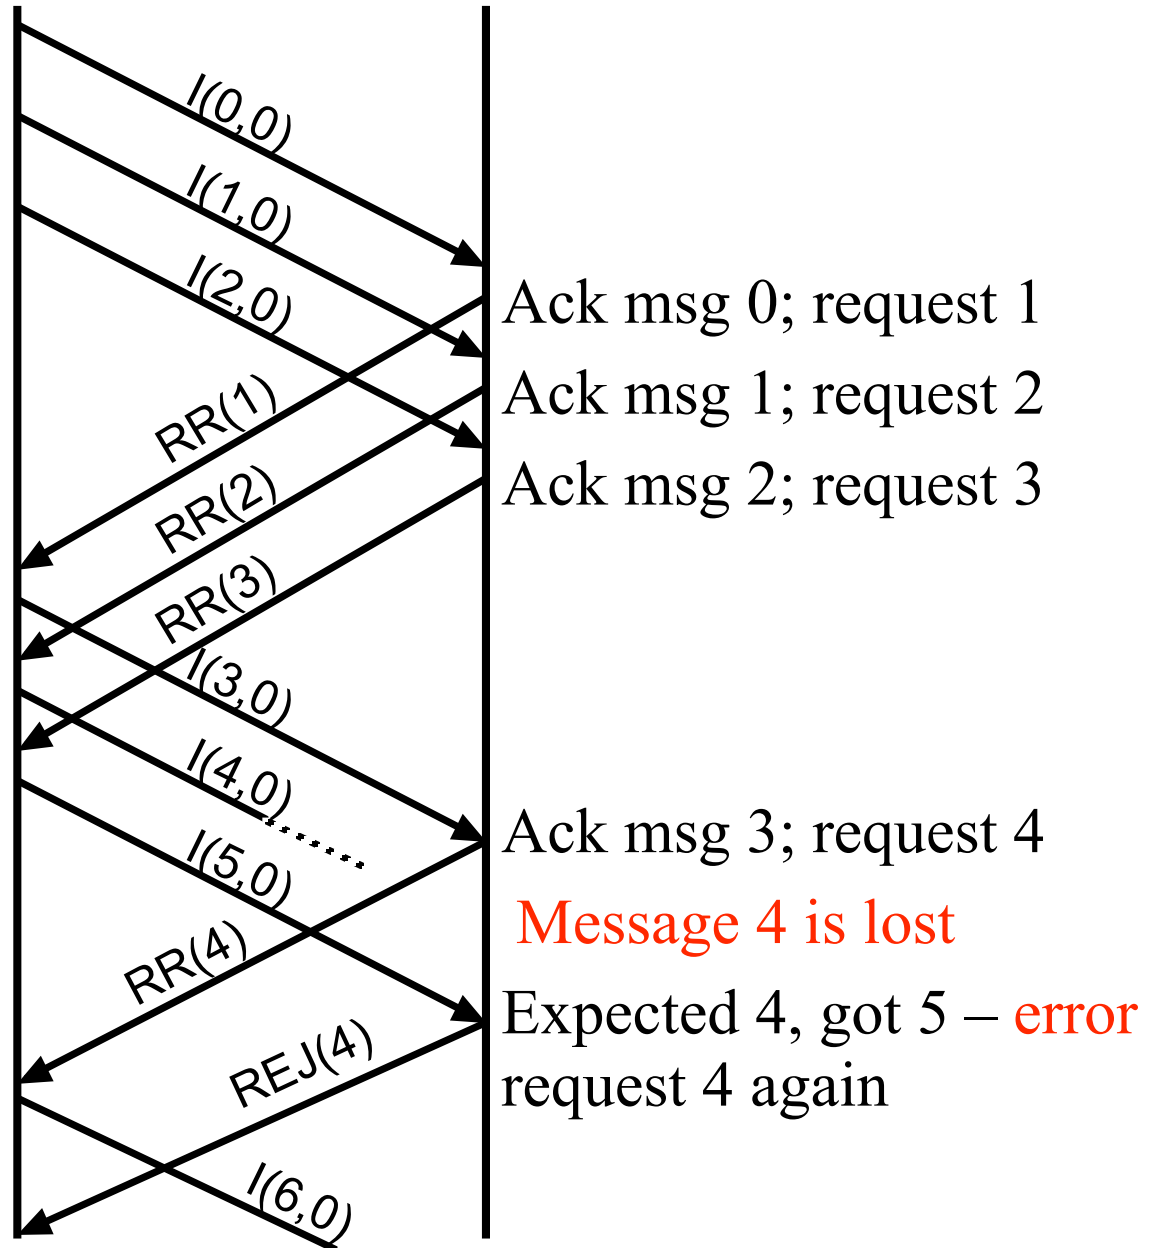
Can send  $1+w-1=3$

Can send  $2+w-1=4$

Can send  $3+w-1=5$

Can send  $4+w-1=6$

Message 4 lost,  
resend 4, 5 & 6



## Two retransmission disciplines

1. **Go-back-N**. If a message is lost, resend *that message and all later ones*. This is simplest if errors are rare and with small window sizes.
  - The receiver discards all messages until the sequence is reestablished.
  - It may be necessary to repeat a whole “window-size” of messages.
  - In HDLC/X.25 it is signalled by a REJ response frame
2. **Selective repeat** repeats only missing or faulty messages. It is best if there is a large link latency and therefore a large window size.
  - The receiver needs more complex memory management to reassemble messages in the right order (and they may have differing sizes).
  - In HDLC/X.25 it is signalled by a SREJ response frame. In the above example frames 5 and 6 will have been sent and the repeated frame 4 could be acknowledged by RR(7), acknowledging all frames up to 6.

# Transmission times and link performance

The link “performance” almost always means

*“How many correct bytes/octetets can be transferred per second”?*

(Sometimes data loss rate, delay, or delay variation may be important)

- This means working out how many bytes of *user data* are transferred in a typical exchange, and how long that exchange takes.
- Recommend working out from first principles in each case, rather than learning difficult formulæ.
- Remember that anything that happens takes time; calculate these times and add them up.
- If you have a 1 Mb/s link between Auckland and Wellington, but a 14.4 kbps modem to connect to the Auckland end, the speed is 14.4 kbps.
- If you have a heavily loaded network, which can transfer at 10 kbps, changing a 14.4 kbps modem to 56 kbps will have little effect.

• Things to consider are –

1. **Frame overheads**. If a frame has an overhead (header, trailer, etc) of 20 octets, 100 bytes of user data needs a 120 byte frame and this 120 bytes must be used in working out times, but still 100 for data.
2. **Acknowledgement overheads**. The total length of the acknowledgement frame (RR or piggy-backed) must be added to the length of the data frame in calculating time.
3. **Link latency**. Information travels at 200,000 km/s (cable) or 300,000 km/s (radio) whether we like it or not. This transmission delay or latency must be included, especially for stop-and-wait protocols.
4. **Turnaround delays**. Many systems take time to finish transmitting, start receiving and stabilise to deliver a reliable signal. This delay can be significant in the older stop-and-wait protocols with frequent line “reversals”.

## Example

Calculate the user data rate for TD800 protocol sending 250 character blocks to a terminal from the CPU at 9600 bps using asynchronous transmission (8 data, 1 start, 1 stop bits). (Address = 1 character.) The line turnaround time is 5 ms and transmission is over 1,000 km. (This uses the “select” sub-protocol.)

	times (ms)
Time per character (10 bits at 9,600 bps) ms	1.04
Select sequence (4 chars + turnaround)	9.17
Latency to terminal (1,000 km) @ 200,000 km/s	5.00
ACK reply, plus turnaround	6.04
Latency back to CPU	5.00
Message + 5 byte overhead + turnaround	270.62
Latency to terminal	5.00
ACK reply, + turnaround	6.04
Latency back to CPU	5.00
Total time for 250 bytes of data	311.88
312 ms to send $250 * 8 = 2000$ bits (bits/second)	<b>6,412.83</b>

It is instructive to repeat the exercise, but with a message of 2,000 characters, instead of the previous 250 characters.

	times (ms)
Time per character (10 bits at 9,600 bps) ms	1.04
Select sequence (4 chars + turnaround)	9.17
Latency to terminal (1,000 km) @ 200,000 km/s	5.00
ACK reply, plus turnaround	6.04
Latency back to CPU	5.00
Message + 5 byte overhead + turnaround	2,093.75
Latency to terminal	5.00
ACK reply, + turnaround	6.04
Latency to terminal	5.00
Total time for 2,000 bytes of data	2,135.00
2.135 s to send $2000 * 8 = 16000$ bits (bits/second)	<b>7,494.15</b>

The effective data rate is now just under 7,500 bps.

(The maximum with asynchronous characters is  $9600 * 8 / 10 = 7680$  bps; the larger message blocks have almost eliminated signal and protocol overheads.)



## Sliding windows performance

- With full-duplex modems we eliminate the turnaround delays, and with overlapped protocols, much of the effect of latency.
- But now must “fill” the link with data so the transmitter never waits.

Assume a satellite link (altitude = 36,000 km;

1-way latency = 72,000 km = 0.24 s, round-trip latency say 0.5 s)

As it takes 0.5 s to get an acknowledgement to a packet we need a window with 0.5 s of data. (At 9600 bps, 4800 bits =  $4800/(128*8) = 4.69$  pkts)

packet size bytes	data rate bps	packets to buffer	bytes to buffer
128	9,600	5	640
128	50,000	25	3,200
1,000	50,000	4	4,000
1,000	1,000,000	63	63,000

The first two lines correspond to X.25 (b1910) communication via satellite. A window of 7 can handle 9,600 bps, but no more; hence the 7 bit sequence number option for a window of 127 packets.

**For a more modern example**, consider a transpacific link (10,000 km each way) sending 53-octet “cells” (for ATM transmission) at 150 Mbps or greater.

The round-trip latency (20,000 km at 200,000 km/s) is 100 ms.

As before, how many packets “fill” the round-trip link before we get a reply?

packet size bytes	data rate bps	packets to buffer	bytes to buffer	
53	155,520,000	36,680	1,944,040	OC-3
53	622,080,000	146,717	7,776,001	OC-12
53	2,488,320,000	586,868	31,104,004	OC-48
53	9,953,280,000	2,347,472	124,416,016	OC-192