$4

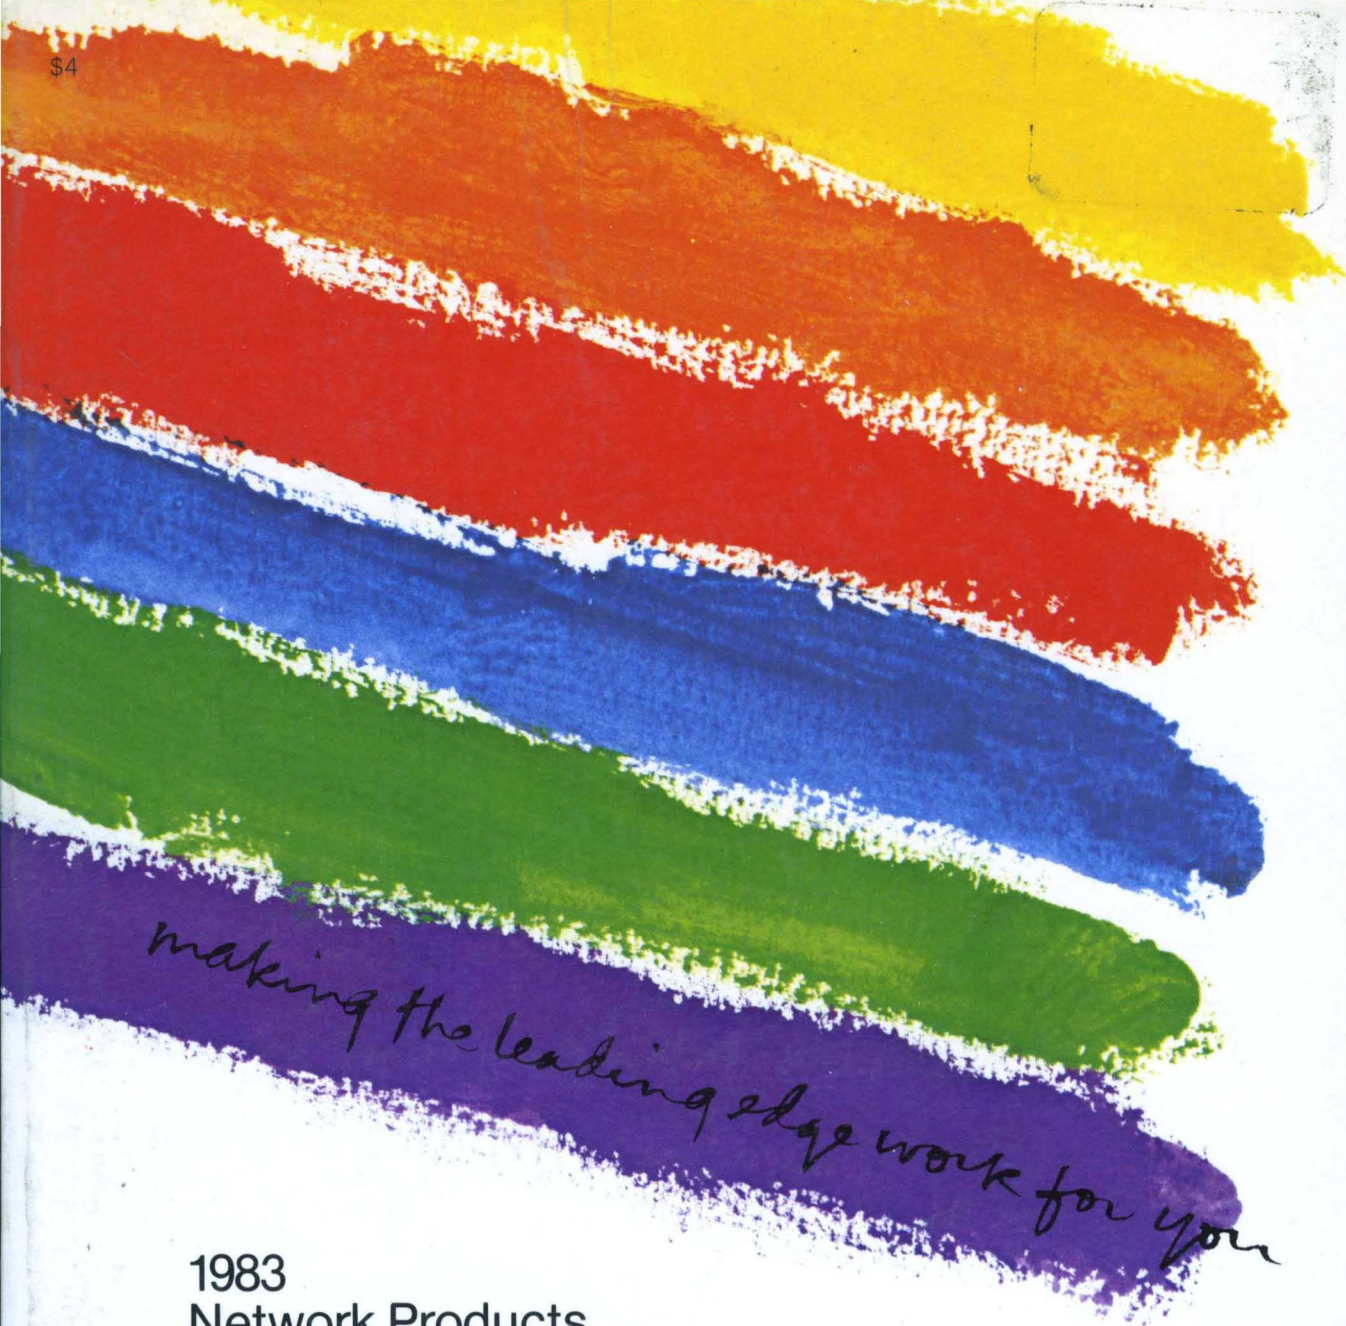

1983 Network Products Handbook

WESTERN DIGITAL
CORPORATION

Western Digital
1983 Network Products
Handbook

1983 Network Products Handbook

Western Digital is on the move, growing and expanding to meet your needs. This handbook, our first dedicated to network products, underscores our commitment to provide you with advanced, innovative solutions to your communication challenges. From the early UARTS we pioneered over a decade ago, to the products described herein, our mission has remained unchanged: to extend the leading edge in communications technology.

The pin-compatible family of products described in this Handbook represents the foundation for a growing selection of intelligent, high level protocol products. These products enable you to quickly and easily add popular protocols, such as X.25, to your systems. And pin compatability simplifies system expansion and protocols, such as local networking. Variations of these off-the-shelf protocol solutions are also available for custom applications.

In 1983 we'll be expanding this family of devices, and announcing additional LSI products and board level solutions, as well.

We encourage you to review the information in this handbook and, when you begin your design, to call our Network Products Hotline for design assistance at 714/966-7828.

Making The Leading Edge Work For You.

This handbook is designed for you, the design engineer. It's intended to be a useful tool, to enable you to make a preliminary evaluation of our products and, later, with samples in hand, to design our products into your own systems.

The data in these pages have been reviewed by our Marketing, Engineering, Manufacturing and Quality groups. Now we would like you to review the information we've provided and tell us how we can improve it. Please feel free to suggest any changes, additions, or clarifications that occur to you. And don't hesitate to call to our attention any sins of omission or commission we may have made.

We're eager to help upgrade the quality of information our industry provides to its customers. So, please, help us. Direct your comments to:

Director of Corporation Communications
WESTERN DIGITAL CORPORATION
2445 McCabe Way
Irvine, CA 92714
(714) 557-3550

Table of Contents

WD2501 X.25 Packet Network Interface (LAP)

FEATURES

- Packet switching controller, compatible with CCITT recommendation X.25, level 2, LAP (WD2501)
- Programmable primary timer (T1) and retransmission counter (N2)
- Programmable A-field which provides a wider range of applications than defined by X.25. These include: DTE-to-DTE connection, multipoint and loop-back testing.
- Direct memory access (DMA) transfer: two channels; one for transmit and one for receive. Send/receive data accessed by indirect addressing method. Sixteen output address lines.
- Zero bit insert and delete
- Automatic appending and testing of FCS field
- Computer bus interface structure: 8 bit bi-directional data bus. CS, WE, RE and four input address lines
- DC to 1.1 MBPS data rate
- TTL compatible
- 48 pin dual in-line packages
- Higher bit rates available by special order

APPLICATIONS

X.25 PACKET SWITCHING CONTROLLER
PART OF DTE OR DCE
PRIVATE PACKET NETWORKS
LINK LEVEL CONTROLLER

GENERAL DESCRIPTION

The WD2501 is a MOS/LSI device which handles bit-oriented, full-duplex serial data communications with DMA, which conforms to CCITT X.25 with programmable enhancements.

The device is fabricated in N-Channel silicon gate MOS technology and is TTL compatible on all inputs and outputs.

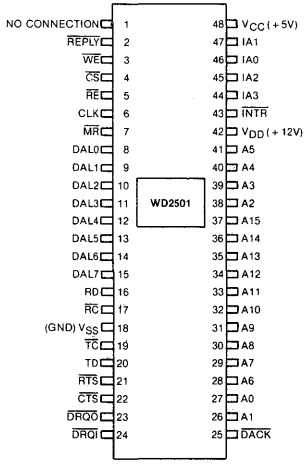

PIN ASSIGNMENTS

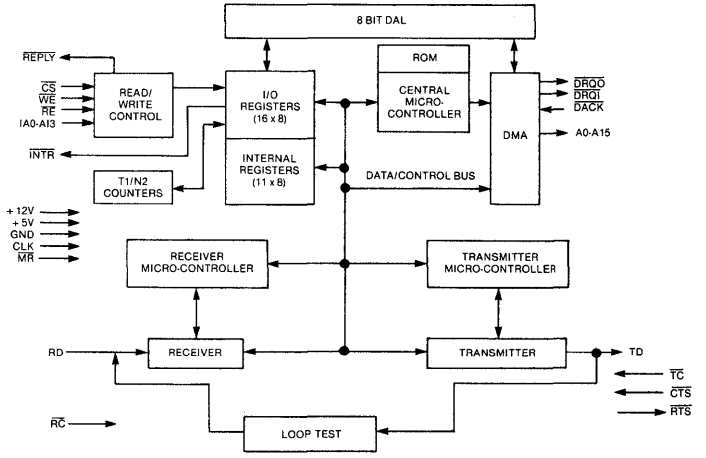

WD2501 BLOCK DIAGRAM

INTERFACE SIGNALS DESCRIPTION (All signals are TTL compatible.)

PIN NUMBER	SYMBOL	PIN NAME	FUNCTION
1		No Connection	Leave pin open
2	$\overline{REPLY}$	Reply	An active low output to indicate the WD2501 has either a $\overline{CS}$•$\overline{RE}$ or a $\overline{CS}$•$\overline{WE}$ input.
3	$\overline{WE}$	Write Enable	The data on the DAL are written into the selected register when $\overline{CS}$ and $\overline{WE}$ are low.
4	$\overline{CS}$	Chip Select	Active low chip select for CPU control of I/O registers.
5	$\overline{RE}$	Read Enable	The contents of the selected register are placed on DAL when $\overline{CS}$ and $\overline{RE}$ are low.
6	CLK	Clock	Clock input used for internal timing. Must be square wave, and greater than 250 KHz.
7	$\overline{MR}$	Master Reset	Initialize on active low. All registers reset to zero, except control bits MDISC and $\overline{LINK}$ are set to 1. $\overline{DACK}$ must be stable high before $\overline{MR}$ goes high.
8-15	DAL0-DAL7	Data Access Lines	An 8-bit bi-directional three-state bus for CPU and DMA controlled transfers.
16	RD	Receive Data	Receive serial data input.
17	$\overline{RC}$	Receive Clock	This is a 1x clock input, and RD is sampled on the rising edge of $\overline{RC}$. RD changes occur on the falling edge of $\overline{RC}$.
18	V_{SS}	Ground	Ground.
19	$\overline{TC}$	Transmit Clock	A 1x clock input. TD changes on the falling edge of $\overline{TC}$.
20	TD	Transmit Data	Transmitted serial data output.
21	$\overline{RTS}$	Request-To-Send	An open collector (drain) output which goes low when the WD2501 is ready to transmit either flags or data.
22	$\overline{CTS}$	Clear-To-Send	An active low input which signals the WD2501 that transmission may begin. If high, the TD output is forced high. May be hard-wired to ground.
23	$\overline{DRQO}$	DMA Request Out	An active low output signal to initiate CPU bus request so the WD2501 can output onto the bus.
24	$\overline{DRQI}$	DMA Request In	An active low output signal to initiate CPU bus request so that data may be input to the WD2501. $\overline{DRQO}$ and $\overline{DRQI}$ will not be low at the same time.
25	$\overline{DACK}$	DMA Acknowledge	An active low input from the CPU in response to $\overline{DRQI}$ or $\overline{DRQO}$. $\overline{DACK}$ must not be low if $\overline{CS}$ and $\overline{RE}$ are low or if $\overline{CS}$ and $\overline{WE}$ are low.
26-41	A0-A15	Address Lines Out (See front page for Pin Assignments)	Sixteen address outputs from the WD2501 for DMA operation. If the control bit ADRV is 1, the outputs are TTL drives at all times. If ADRV is 0, the outputs are three-state, and are HI-Z whenever $\overline{DACK}$ is high. (ADRV is in Control Register #1.)
42	V_{DD}	Power Supply	+ 12VDC power supply input.

INTERFACE SIGNALS DESCRIPTION CONTINUED

PIN NUMBER	SYMBOL	PIN NAME	FUNCTION
43	$\overline{INTR}$	Interrupt Request	An active low interrupt service request output, and returns high when Status Register #1 is read.
44-47	IA0-IA3	Address Lines In (See front page for Pin Assignments)	Four address inputs to the WD2501 for CPU controlled read/write operation with registers in the WD2501. If ADRV = 0, these may be tied to A0-A3. (ADRV is in Control Register #1.)
48	V_{CC}	Power Supply	+5VDC power supply input.

ORGANIZATION

A detailed block diagram of the WD2501 is shown in Figure 1.

Mode control and monitor of status by the user's CPU is performed through the Read/Write Control circuit, which reads from or writes into registers addressed by IA0-IA3.

Transmit and receive data are accessed through the DMA control. Serial data is generated and received by the bit-oriented controllers.

Internal Control of the WD2501 is by means of three internal microcontrollers; one for transmit, one for receive, and one for overall control.

Parallel transmit data are entered into the Transmitter Holding Register (THR), and then presented to the Transmitter Register (TR) which converts the data to a serial bit stream. The Cyclic Redundancy Check (CRC) is computed in the 16-bit CRC register, and the results become the transmitted Frame Check Sequence (FCS).

Parallel receive data enter the Receiver Holding Register (RHR) from the 24-bit serial Receiver Register (RR). The 24-bit length of RR prevents received FCS data from entering the RHR (unless residual bits are received). The receiver CRC register is used to test the validity of the received FCS. A 3-stack FIFO is included in the receiver.

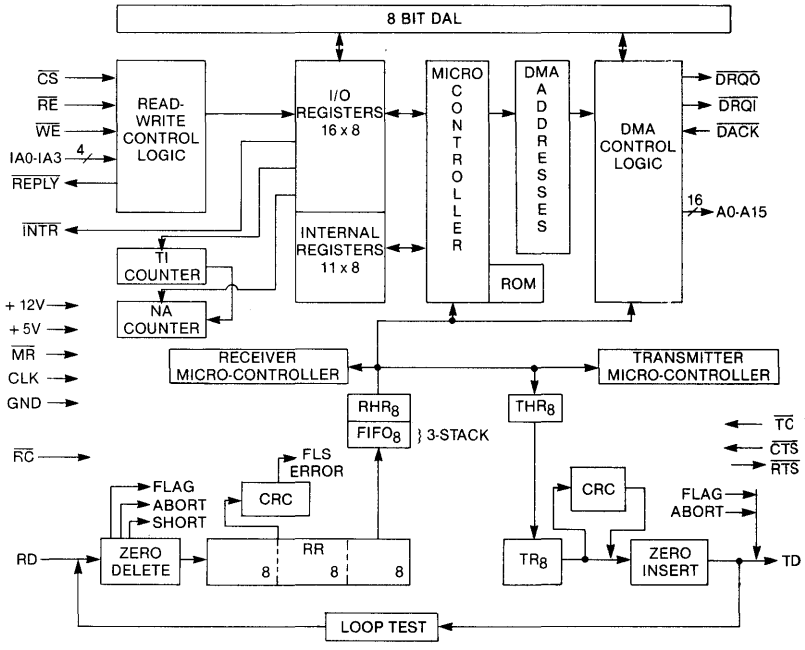

Figure 1. WD2501 BLOCK DIAGRAM — DETAILED

FRAME FORMAT

The WD2501 performs "bit-oriented" data communications control. According to the bit-oriented procedures (HDLC, SDLC, ADCCP), each serial block of data is called a frame.

Each frame starts and ends with a Flag (01111110). A single flag may be used both as the closing flag of one frame and the opening flag of the next frame. In between flags, data transparency is provided by the insertion of an 0 bit after all sequences of 5 contiguous 1 bits. The receiver will strip the inserted 0 bits. The last 16-bits before the closing flag is in the Frame Check Sequence (FCS). Each frame also includes address and control fields (A and C fields).

The FCS calculation includes all data between the opening flag and the first bit of the FCS, except for 0's inserted for transparency. The 16-bit FCS has the following characteristics:

Polynomial = $X^{16} + X^{12} + X^5 + 1$
Transmitted Polarity —Inverted
Transmitted Order —High Order Bit First
Preset Value —All I's

After the frame is received, if there were no errors, then the remainder in the CRC register (internal in the WD2501) will be:

1111000010111000 F0B8

The WD2501 generates and tests the Flag, FCS, A-Field, C-Field, and performs zero bit insertion and deletion. According to the X.25 protocol, there are three types of frames: supervisory (S-frame), unnumbered (U-frame), and information (I-frame).

The WD2501 performs frame level (level 2) link access control. All S- and U-frames are automatically generated and tested by the WD2501. The user need only be concerned with the I-frames, which are packets.

The WD2501 will transmit contiguous flags for interframe time fill.

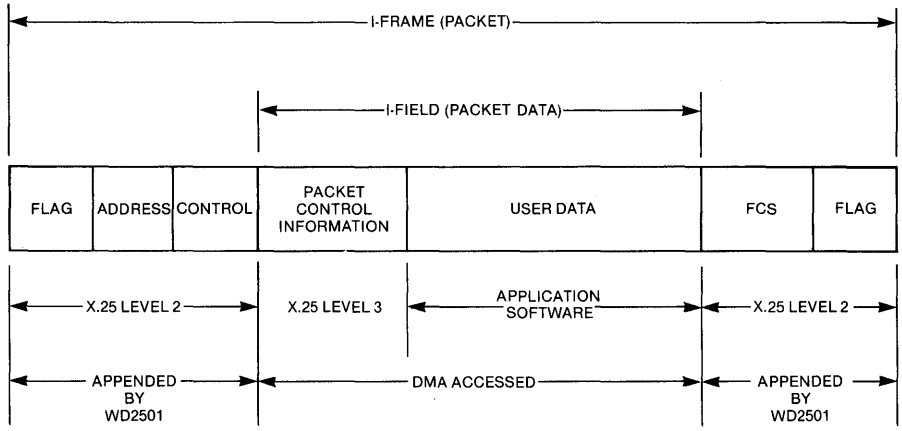

X.25 MODE

NOTE: X.25 Level 1, is the Physical Interface

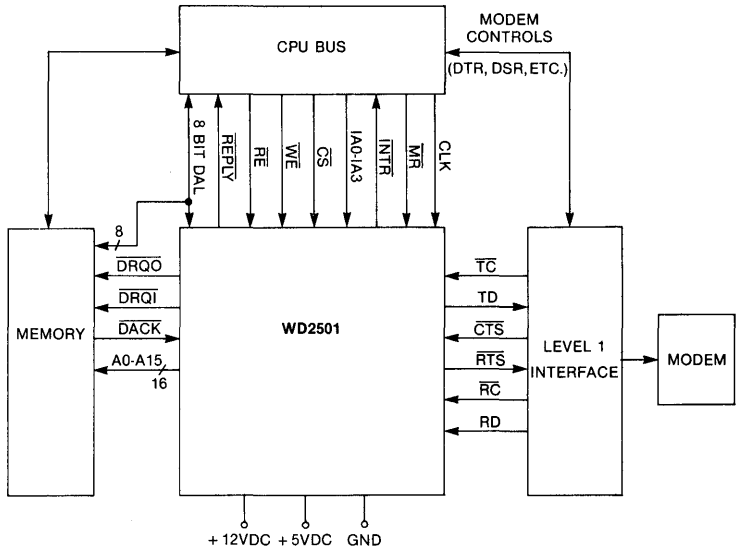

Figure 2. SYSTEM CONNECTION

II. PROGRAMMING THROUGH REGISTERS

The WD2501 is controlled and monitored by sixteen I/O registers.

Control, status, and error bits will be referred to as CR, SR, or ER, respectively, along with two digits. For example, SR16 refers to status register #1 and bit 6, which is "XBA."

REGISTER DEFINITION

REG #	IA3	IA2	IA1	IA0	REGISTER	REGISTER GROUPING
0	0	0	0	0	CR0	OVERALL CONTROL AND MONITOR
1	0	0	0	1	CR1	
2	0	0	1	0	*SR0	
3	0	0	1	1	*SR1	
4	0	1	0	0	*SR2	
5	0	1	0	1	*ER0	
6	0	1	1	0	*CHAIN MONITOR	RECEIVER MONITOR
7	0	1	1	1	*RECEIVED C-FIELD	
8	1	0	0	0	T1	TIMER
9	1	0	0	1	N2/T1	
A	1	0	1	0	TLOOK HI	DMA SET-UP
B	1	0	1	1	TLOOK LO	
C	1	1	0	0	CHAIN/LIMIT	
D	1	1	0	1	NOT USED	
E	1	1	1	0	XMT COMMAND "E"	"A" FIELD
F	1	1	1	1	XMT RESPONSE "F"	

*CPU READ ONLY. (Write Not Possible)

NOTE:
1. Registers E and F should be set-up only with MDISC = 1.

CONTROL, STATUS, ERROR REGISTERS

REGISTER	BIT # 7	6	5	4	3	2	1	0
CR0	0	0	READY BUT IDLE	ACTIVE/ $\overline{\text{PASSIVE}}$	LOOP TEST	2RAMT	RECR	MDISC
CR1	0	0	0	ADRV	0	0	0	SEND
SR0	NA2	NA1	NA0	RNRR	NB2	NB1	NB0	RNRX
SR1	1PKR	1XBA	1ERROR	0	NE2	NE1	NE0	0
SR2	T1OUT	$\overline{\text{IRTS}}$	REC IDLE	0	0	0	0	$\overline{\text{LINK}}$
ER0	ER07	ER06	ER05	ER04	ER03	ER02	ER01	ER00

1 Causes Interrupt ($\overline{\text{INTR}}$ Goes Low).

2 RAMT is used to initiate a self-test feature (see section on Self-Test). For normal operation, program to 0.

CONTROL REGISTER 0

REGISTER	CR07	CR06	CR05	CR04	CR03	CR02	CR01	CR00
CR0	0	0	READY BUT IDLE	ACTIVE/ $\overline{\text{PASSIVE}}$	LOOP TEST	RAMT	RECR	MDISC

BIT	DESCRIPTION
CR00	MDISC is a mandatory disconnect command. MDISC will cause a logical disconnect in the link. No DMA accessed data will be transferred as long as MDISC = 1. After Master Reset (MR pin transition from low to high), MDISC will be set. The WD2501 will neither transmit nor accept received data until MDISC = 0.
CR01	This bit is RECR which defines the CPU's receiver buffer as initially Ready (CR01 = 1). If RECR = 1, this bit indicates that the WD2501 may begin receiving I-frames. (See SR00)
CR02	RAMT — This bit activates an internal register test. See Section on Self tests for description.
CR03	The LOOP TEST bit will connect the transmitted data output to the receiver input. The receiver input pins $\overline{\text{RD}}$ and $\overline{\text{RC}}$ are logically disconnected. The "E" and "F" registers of the A-field should be equal.
CR04	This bit will cause the WD2501 to initiate link set-up if CR04 = 1, or to wait for a link set-up from the remote device if CR04 = 0.
CR05	RBI bit. For compatibility with the 2511, let RBI = 0.
CR7, CR6	Unused control bits, like CR07, should be 0.

CONTROL REGISTER 1

REGISTER	CR17	CR16	CR15	CR14	CR13	CR12	CR11	CR10
CR1	0	0	0	ADRV	0	0	0	SEND

BIT	DESCRIPTION
CR10	The SEND bit (CR10) is used to command the WD2501 to send the next packet or packets. If SEND = 1, the WD2501 will read from TLOOK the BRDY bit of the next segment for transmission. If BRDY = 0, the WD2501 will clear SEND and no action occurs. If BRDY = 1, the WD2501 will then read TSADR and TCNT, followed by the transmission of that buffer. After transmission, the WD2501 clears BRDY of the segment just transmitted, and reads BRDY of the next segment. If 1, the next segment is transmitted. If 0, the SEND bit is cleared, and transmission of packets is stopped. As a matter of good practice, the CPU should set SEND each time one of the BRDY bits is set.
CR11-13	Unused bits, write in 0's.
CR14	The ADRV bit (CR14) is the control for the 16 bit output addresses (A0-A15). If ADRV = 0, the outputs are tri-state and are in HI-Z, except when DACK goes low. If ADRV = 1, the outputs are always low impedance (TTL), and are high-level when DACK is high.
CR15-17	Unused bits, write in 0's.

STATUS REGISTER 0

REGISTER	SR07	SR06	SR05	SR04	SR03	SR02	SR01	SR00
SR0	NA2	NA1	NA0	RNRR	NB2	NB1	NB0	RNRX

BIT	DESCRIPTION
SR00	RNRX. An RNR has been transmitted, or will be at next opportunity. The CPU should set RECR, when receive buffers are available.
SR03-SR01	NB2-NB0. Next block to be transmitted.
SR04	RNRR. This bit is set when an RNR frame is received. Once set, it is cleared when an RR, REJ, or UA is received.
SR07-SR05	NA2-NA0. Next block of transmitted data to be Acknowledged.

STATUS REGISTER 1

REGISTER	SR17	SR16	SR15	SR14	SR13	SR12	SR11	SR10
SR1	$_1$PKR	$_1$XBA	$_1$ERROR	0	NE2	NE1	NE0	0

BIT	DESCRIPTION
SR10	0
SR13-SR11	NE2-NE0. Next Expected packet segment number of RLOOK.
SR14	0
$SR15_1$	The ERROR bit indicates: 1) An error has occurred which is not recoverable by the WD2501, or 2) A significant event has occurred. The "significant events" are: change in link status (link-up or down), the WD2501 is progressing to the next segment in a chained receive buffer, or one direction of the link has been reset. The exact nature of the reason for the ERROR bit is given in ER0.
$SR16_1$	The XBA bit means that a previously transmitted Block, or Blocks, have been acknowledged by the remote device. Upon acknowledgement, the ACK'ED bit is set to "1" for each segment in TLOOK which was acknowledged.
$SR17_1$	The PKR bit stands for Packet Received. This means that a packet has been received error-free and in correct sequence according to the received N (S) count. The data (I-field) has been placed in the CPU's RAM memory. NE is advanced.

NOTE 1: The three interrupt-causing bits are SR17, SR16, and SR15. Any of the three will cause an interrupt request ($\overline{INTR}$ goes low). After SR1 is read, all three bits are reset to 0, and $\overline{INTR}$ returns high.

STATUS REGISTER 2

REGISTER	SR27	SR26	SR25	SR24	SR23	SR22	SR21	SR20
SR2	T1OUT	$\overline{IRTS}$	REC IDLE	0	0	0	0	$\overline{LINK}$

BIT	DESCRIPTION
SR20	If the link is established, $\overline{LINK}$ = 0. If the link is logically disconnected, $\overline{LINK}$ = 1.
SR24-21	Unused Bits — 0
SR25	REC IDLE indicates that the WD2501 has received at least 15 contiguous 1's.
SR26	$\overline{IRTS}$ stands for the Internal Request-To-Send bit, and indicates that the transmitter is attempting (successful or not) to send either data or flags. If the $\overline{RTS}$ pin is not tied to ground or WIRE-ORED with another signal, then $\overline{IRTS}$ = $\overline{RTS}$.
SR27	T1OUT bit means that timer T1 has timed-out. This bit returns to 0 when T1 is re-started. When T1OUT = 1, T1 is not running. NOTE: This bit could be a 1 for a few micro-seconds in between intervals when T1 stops and is restarted.

ERROR REGISTER (ER0)

ER07	ER06	ER05						
0	0	0	ER01 = ROR ER02 = TUR ER04 = RLNR					
0	0	1	**ER04**	**ER03**	**ER02**	**ER01**	**ER00**	
			0	0	0	0	1	LINK is up. (Was down)
			1	0	0	0	0	Received DISC while LINK up, or partially up.
			0	0	1	0	0	DISC sent, since SARM sent N2 times without UA. (Note 3).
			0	0	0	1	0	DISC sent, REC IDLE for T1xN2.
0	1	0	**CHAIN STATUS** ER00 = GNCS ER01 = CNR					
1	0		LINK RESET RECEIVED if ER05-ER00 = 000000 LINK RESET TRANSMITTED if ER05-ER00 = non-zero ER00 similar to W ER01 similar to X ER02 similar to Y ER03 similar to Z ER05 means received F = 1, but did not send P = 1 ER04 — I-frame was sent N2 times without acknowledge					
1	1		COMMAND REJECT RECEIVED if ER05-ER00 = 000000 TRANSMITTED if ER05-ER00 = non-zero ER00 = W ER01 = X ER02 = Y ER03 = Z					

Unassigned bits are always 0.

TERMS USED IN ERROR REGISTER

GNCS	Going to Next Chain Segment.
RLNR	RLOOK Not Ready. REC RDY bit of next segment is 0, but RECR = 1. This interrupt will not occur if RECR = 0.
ROR*	Receiver Over-Run. The Receiver Register (RR) had a character to load into the FIFO, but the FIFO was full.
TUR*	Transmitter Under-Run. The Transmitter Register (TR) needed a character from the Transmitter Holding Register (THR), but the THR was not ready. The frame being transmitted is ABORTed.
CNR	Next Chain Segment of Receiver was Not Ready.

*TUR and ROR means that the bit rate clocks (TC and RC) are either too fast for the WD2501, or the DACK response to DRQI and DRQO is too slow, or both.

W, X, Y, Z OF CMDR

A command reject (CMDR) contains a three byte I-field. The first byte is the rejected frame control field. The second byte contains the current N (S) and N (R) counts of the terminal reporting the rejection condition. The third byte contains W-X-Y-Z-Z1-0-0-0 where W is the LSB.

W set to 1 indicates that the control field received and returned in the first I-frame byte was invalid.

X set to 1 indicates the rejected frame contained an I-field which is not permitted with this command. Bit W must be set to 1 when X = 1.

Y set to 1 indicates the I-field exceeded the total limit established. Y is mutually exclusive with W.

Z set to 1 indicates the control field contained an invalid N(R). Z is mutually exclusive with W. Since the WD2501 is LAP, only, a CMDR will not be sent to an invalid N (R). A link reset (SARM) is sent. Thus, Z is not used.

Upon receiving a CMDR, the WD2501 will place the 3 byte I-field in memory by DMA, just as if the CMDR were a packet. However, NE is not advanced, and the next received packet will over-write the CMDR I-field. Therefore, the CPU should "quickly" use the 3 byte I-field in the received CMDR.

When the WD2501 transmits a CMDR, the command reject condition is entered. Only a received SARM or DISC will clear this condition. If any other command is received, the WD2501 will re-transmit the CMDR. However, valid responses received will be acted upon accordingly. Also, the WD2501 may transmit packets while in the command reject condition.

MEMORY ACCESS METHOD

The WD2501 memory access is accomplished by the use of DMA and two look-up tables. These tables are set-up to allow up to 7I-frames to be outstanding in each direction of the communications link. The look-up tables are divided into two areas (T-LOOK and R-LOOK) and are in memory external to the WD2501.

TLOOK
RLOOK

These tables contain address and control information for individual send/receive packets.

To access TLOOK and RLOOK only the starting address of TLOOK is loaded into the WD2501 in registers A' and B'.

REG A	A15	A14	A13	A12	A11	A10	A9	A8
REG B	A7	A6	A5	A4	A3	A2	A1	A0

A0-A15 16 bit TLOOK starting address

The TLOOK and RLOOK tables are each divided into 8 segments and each segment contains 8 bytes. Figure 3 illustrates the segmentation of TLOOK and RLOOK. Figure 5 and 6 illustrate the contents of a single TLOOK and RLOOK segment.

In transmit, the WD2501 will have read from TLOOK the starting address and length of the first packet to be transmitted. The WD2501 will automatically transmit the flag, address, and control fields. Next, the information field data will be transmitted using DMA from the "SEND #0 PACKET" memory. At the end of the information field, the WD2501 will automatically send the FCS and closing Flag. The WD2501 will then move on to the next packet.

If retransmission of one or more (up to seven) packets becomes necessary, the WD2501 will automatically retrace the previous transmissions through the TLOOK table. The user's CPU software does not become involved in the retransmission. However, an ERROR COUNTER is incremented. (See Error Counter Section.)

To receive, each frame is checked for correct address and FCS fields and for type of control field. If the frame is a packet, the information field is placed in the assigned memory location in a method similar to that used in transmit. After the packet is received error-free and in proper N(S) sequence count, an

interrupt is generated, and the WD2501 is ready for the next packet which will be placed in the next location.

Figure 4 shows a "store-and-forward" example that is useful in a network node.

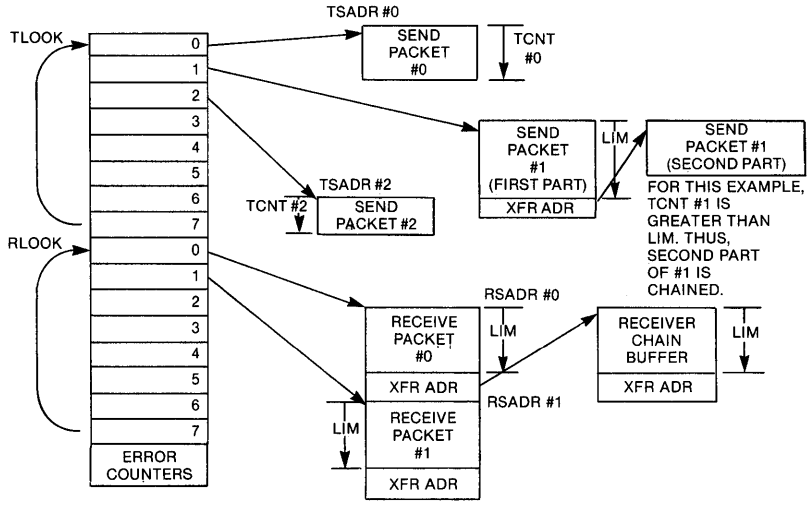

Figure 3. MEMORY ACCESS SCHEME

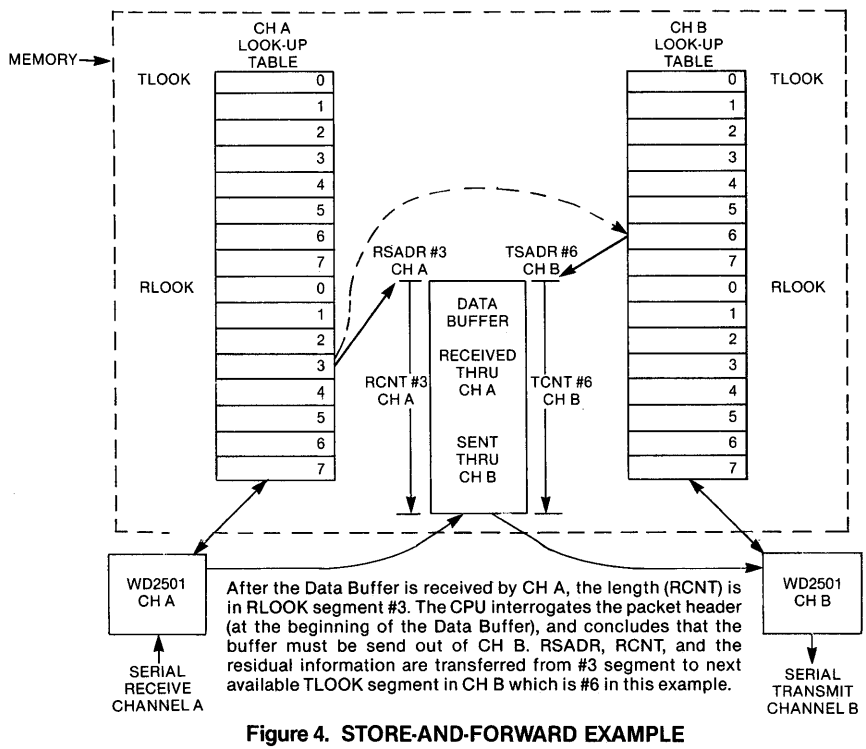

Figure 4. STORE-AND-FORWARD EXAMPLE

TLOOK AND RLOOK

Figures 5 and 6 detail the individual segments for TLOOK and RLOOK.

BRDY means that the transmit buffer is ready. The WD2501 will send the block only after the CPU makes BRDY = 1. (BRDY is used in conjunction with the SEND bit.) At the completion of the transmission, the WD2501 will make BRDY = 0, and then read the BRDY of the next segment.

After transmitting a packet, an acknowledgement must be received from the remote device. The acknowledgement is contained in the received N(R) count of an I-frame or S-frame. Upon acknowledgement, the WD2501 will make ACK'ED = 1, and generate a block-acknowledged interrupt. Before assigning a new block to a segment in TLOOK, the CPU must make sure that the previous block which used that segment number has been acknowledged.

REC RDY informs the WD2501 that the receive buffer is ready. The WD2501 will not receive a packet into a buffer referenced by a particular segment until REC RDY = 1. If the WD2501 progresses to a segment which has REC RDY = 0, an error interrupt will be generated.

After receiving an error-free packet with correct N(S), the WD2501 will, in order: 1) Set FRCML (Frame Complete), clear REC RDY, and store received residual count. 2) Store the received length, in characters, of the I-field in RCNT HI and RCNT LO. 3) Advance the NE count, and generate a packet received interrupt. 4) Acknowledge the received packet at the first opportunity.

The addresses (TSADR and RSADR) are 16-bit binary addresses. HI represents the upper 8-bits and LO represents the lower 8-bits. The counts (TCNT and RCNT) are 12-bit binary numbers for the number of characters in the I-field.

TSADR is the starting address of the buffer to transmit, and TCNT is the binary count of the number of bytes to transmit.

RSADR is the starting address of the receive buffer. After successfully receiving the packets, the WD2501 will write the value of RCNT which is the binary length of the received packet.

Whether the WD2501 accesses a look-up table or a memory block, a DMA Cycle is required for each access.

Figure 5. TLOOK SEGMENT

BYTE # IN SEGMENT	BIT # 7	6	5	4	3	2	1	0
1	ACK'ED	NS	NS	NS	NS	NS	NS	BRDY
2	TSADR HI							
3	TSADR LO							
4	SPARE				TCNT HI			
5	TCNT LO							
6	SPARE FOR USER DEFINITION							
7	SPARE							
8	SPARE							

NS = Not Spare

Figure 6. RLOOK SEGMENT

BYTE # IN SEGMENT	BIT # 7	6	5	4	3	2	1	0
1	FRCML*	NS	NS	NS	RES2	RES1	RES0	REC RDY
2	RSADR HI							
3	RSADR LO							
4	NOT SPARE				RCNT HI			
5	RCNT LO							
6	SPARE FOR USER DEFINITION							
7	SPARE							
8	SPARE							

NS = Not Spare (NOTE: The "not spare" bits may be either 1 or 0).

*FRCML = Frame Complete

The control bits in TLOOK (BRDY and ACK'ED) and in RLOOK (FRCML and REC RDY) define various states for each segment. These states are shown below:

TLOOK STATES

ACK'ED	BRDY	STATE
0	1	Ready To Transmit (CPU set BRDY, cleared ACK'ED)
0	0	*Transmitted and Awaiting Acknowledge (WD2501 cleared BRDY)
1	0	Received Acknowledge (WD2501 set ACK'ED)
1	1	This state not allowed

*State 0-0 could also occur whenever there is no data ready to send.

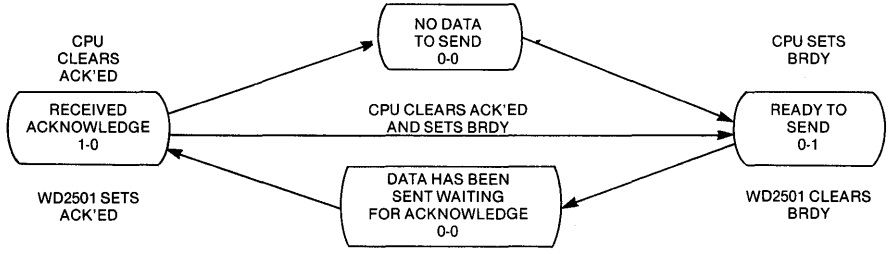

TLOOK SEGMENT STATE FLOW

Notice that in a TLOOK segment, the 0-0 state could have two meanings. Due to control internal to the WD2501, this will not pose an ambiguity to the WD2501. However, if it is a difficulty to the CPU, the CPU could do two things:

1. Avoid ever entering the "NO DATA TO SEND"
2. At start-up, the CPU should set all ACK'ED bits. (Since this would only be a start-up procedure, this would not violate the "deadly embrace" rule.)

As another alternative, the CPU could use one or more of the SPARE bits for additional state control.

In the "WAITING FOR ACKNOWLEDGE" state, one or more re-transmissions could occur.

REG = C	CHAIN				LIMIT			
CHAIN / LIMIT	Bit 7	Bit 6	Bit 5	Bit 4	Bit 3	Bit 2	Bit 1	Bit 0

CHAINING

The WD2501 includes a chained-block feature which allows the user more efficient use of memory particularly in situations where the maximum packet size is much larger than the average packet size.

Register C is used to program the chaining feature. The upper 4 bits define CHAIN which is the number of chain segments allowed in addition to the first segment. (If this feature is not used, make CHAIN all 0's).

The lower 4 bits of Register C defines LIMIT which is the size of the buffer in multiples of 64 bytes including the transfer address (XFR ADR). If LIMIT is 0000, the size is 64. For 0001, the size is 128, and so on.

For example, suppose that the LIMIT defines a segment size of 128 and that CHAIN defines 8 additional segments in addition to the first. (Note: Register C would be hex 81 in this example.) When 126 bytes of I-field have been received, the WD2501 will take the next two bytes as a transfer address (XFR ADR) pointing to another segment. At the end of that segment is another XFR ADR, and so on, up to a maximum of 9 total segments.

On the receiver, a XFR ADR of all 0's will mean that the next segment is not ready. If the WD2501 reaches a XFR ADR on the receiver, there will be an Error Interrupt code 42 if XRF ADR = all 0's. Otherwise, there will be an Interrupt code 41 which is a status indication that the WD2501 is going to the next segment.

On the receiver, Register 6 upper 4 bits gives a status of which chain segment is currently being used.

The transmitter chaining works just like the receiver with the following exceptions:

1. XFR ADR = all 0's will not indicate next segment not ready.
2. There is no interrupt when going from one segment to another.
3. There is no status of the current segment being used.

The total receiver limit is given by the expression:

TOTAL RECEIVER LIMIT = (64 X (1 + LIMIT) − 2) X (1 + CHAIN)

Also, note that the transmitter and receiver counts are modified by 2 for each time a chain boundary is crossed. For example, if LIMIT = 0001 (segment size = 128 including XFR ADR), and if an I-field of 270 bytes is to be transmitted, then there will be two times that a chain boundary is crossed. The TCNT must be made 274 to send 270 bytes. The same is true for RCNT.

RLOOK STATES

FRCML	REC RDY	STATE
0	1	Ready To Receive (CPU set REC RDY, cleared FRCML)
1	0	Received Packet (WD2501 set FRCML, cleared REC RDY)
0	0	Not Ready (CPU cleared FRCML)
1	1	This state not allowed

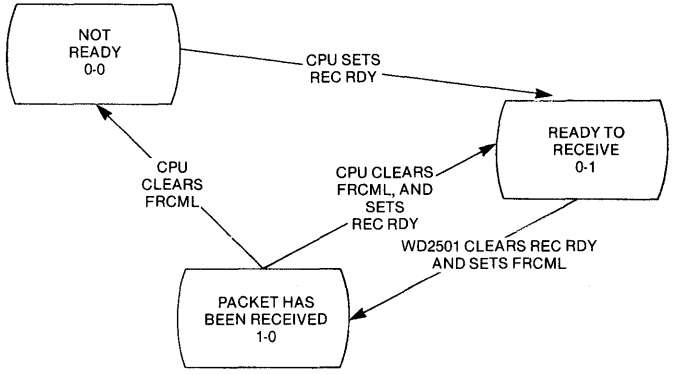

RLOOK SEGMENT STATE FLOW

"DEADLY EMBRACE" PREVENTION

A "deadly embrace" can occur when two processors reach a state where each is waiting for the other. In this case, the two processors are the user's CPU and the micro-controller inside the WD2501. Therefore, to prevent the "deadly embrace," the following rule is obeyed by the WD2501 and should also be obeyed by the user's CPU. This rule applies to TLOOK, RLOOK and to the I/O registers. The Error Counters do not apply to this rule.

> RULE: If a bit is set by the CPU, it will not be set by the WD2501, and vice versa. If a bit is cleared by the WD2501, it will not be cleared by the CPU, and vice versa.

As an example, the BRDY bit in the TLOOK segments is set by the CPU, only, but cleared by the WD2501, only.

TLOOK AND RLOOK POINTERS

There are three 3-bit counters for the status of the segments in TLOOK and RLOOK. Status Register #0 (SR0) contains counters NA and NB which are used in conjunction with TLOOK. NB is the segment number of the next block to be transmitted, and is advanced at the end of each block transmission. NA is the value of the segment of the next block which will be acknowledged. If all transmitted blocks have been acknowledged, then NA = NB.

In SR1 is a 3-bit counter, NE, used with RLOOK. NE is the value of the segment number where the next received packet will be placed.

NA = Next to be Acknowledged

NB = Next Block to be Transmitted

NE = Next Expected to be Received

VARIABLE BIT LENGTH AND RESIDUAL BITS

The WD2501 will send 8 bits per character, only. Also all transmitted frames will be integer multiples of bytes.

The WD2501 may receive a packet with, or without, an integer multiple of bytes. The "RES" bits in the RLOOK tables shows the number of residual bits. The residual bits occupy the lower portion of the last received character.

RES 2	RES 1	RES 0	Received Residual Bits
0	0	0	0
0	0	1	7
0	1	0	6
0	1	1	5
1	0	0	4
1	0	1	3
1	1	0	2
1	1	1	1

ERROR COUNTERS

Following contiguously after RLOOK are six 8-bit error counters. The WD2501 will increment each counter at the occurrence of the defined event. However, the WD2501 will not increment past 255 (all 1's). The CPU has the responsibility of clearing each counter. The first counter past RLOOK is #1, etc.

ERROR COUNTER	COUNT
1	*Received Frames with FCS Error (includes frames ABORTed in the I-field).
2	Received Short Frames (less than 32-bits)
3	**Number of times T1 ran-out (completed)
4	Number of I-Frame Retransmissions due to T1 completing
5	*REJ Frames Received
6	REJ Frames Transmitted

*These counters are incremented only if the received A-field is equal to either Register E or F.

**Incremented only when attempting to transmit a command.

The Error Counters are accessed by the WD2501 transmitter DMA channel. Therefore, if multiple errors are received while the WD2501 is transmitting a long frame, only the last error will be counted. The only Counters which could miss counts because of this are Counters #1, #2, and #5. The error Counters are incremented only when the link is up (LINK = 0).

OTHER I/O REGISTERS

RECEIVED C-FIELD

Register 7 is the C-field of the last received frame, provided the A-field of the frame was equal to either register E or F, the FCS was good, and the frame contained 32 or more bits, and the WD2501 is not waiting for a SARM or DISC in response to a transmitted CMDR.

TIMER

Registers 8 and 9 define a 10-bit timer (T1), and a 6-bit Maximum Number of Transmissions and Retransmissions counter (N2).

REGISTER	7	6	5	4	3	2	1	0
8	T1							LSB
9	N2							
	MSB					LSB	MSB	

MSB = Most Significant Bit
LSB = Least Significant Bit

T1 provides the value of a delay in waiting for a response and/or acknowledgement. The delay is the binary count of clock CT where:

$$CT = \frac{16{,}384}{CLK}$$

Thus, if CLK = 1 MHz, then T1 may be set in increments 16.384 milliseconds, to a maximum delay of 16.78 seconds. All ones in T1 is maximum delay.

Once the CPU establishes T1 and N2, there is no need to write into T1 and N2 again unless a master reset (MR) has occurred, there is a power loss, or the CPU needs to change T1 or N2. If a time-out occurs, the WD2501 will still retain T1 and N2.

The conditions for starting, stopping, or restarting T1 are shown below: ("Re-start" means starting T1 before it ran-out).

START T1	RE-START T1	STOP T1
1. *I-frame sent if T1 not already in progress due to previous I-frame	*Acknowledgement received to some, but not all, I-frames.	Acknowledgement received for all I-frames.
2. —	*RNR received while link up.	
3. SARM or DISC sent. (N2 restarted at first occurrence)	—	UA Received for SARM, UA or DM Received for DISC.
4. Receiver Idle (REC IDLE = 1)		Detect REC IDLE = 0
5. RR or RNR sent because RRT1 = 1		—

*N2 is restarted where shown by asterisk (*).

NOTE: Reason 4 and 5 (above) are overridden by reasons 1, 2 and 3. On reason 2, T1 is stopped if the received RNR acknowledged all outstanding packets.

"A" FIELD REGISTERS

Registers E and F provide a programmable A-field. This allows the WD2501 to be a super-set of the X.25 document. That is, the WD2501 can handle a wider range of application than the DTE-DCE links defined in X.25. These wider ranges include: DTE-to-DTE connection, multipoint, and loop-back testing.

If the WD2501 is strictly in an X.25 DTE-DCE link, use the values shown below:

DTE Register E = 01
Register F = 03

DCE Register E = 03
Register F = 01

If performing a loop-back test either internal (CR03 = 1) or external, registers E and F should be the same.

V. LAP PROCEDURE

The Link Access Procedure (LAP) is described in CCITT Recommendation X.25 as the Level 2 protocol for the Asynchronous Response Mode (ARM). LAPB is the Level 2 protocol for the Asynchronous Balanced Mode (ABM). This section describes how the WD2501 performs LAP. (The WD2511 performs LAPB.)

Zero bit insertion/deletion, use of flags and FCS are part of Level 2, and have been discussed in this document.

The DTE is the Data Terminal Equipment and the DCE is the Data Circuit Termination Equipment, and is the network side of the DTE-DCE connection. X.25 defines the protocol in the DTE-DCE link. The link is symmetric in that both the "primary" and "secondary" functions are included within the DTE and DCE. Commands and responses are differentiated by the A-field as shown below:

The individual commands and responses are shown in Figure 7.

The Poll bit (P) is used in conjunction with Time-Out Recovery. Time T1 is started at the beginning of a transmitted command provided it has not been previously started. If T1 runs out (completes), the command will be retransmitted with P = 1. If T1 runs out again, the command will again be retransmitted, with P = 1 up to N2 times. At N2 + 1, an error interrupt will occur. If the command was an I-frame, the WD2501 will reset the link by transmitting a SARM. If a SARM (either for link set-up or link reset), the WD2501 will send a DISC, If a DISC, the WD2501 will continue to send a DISC.

If an RNR has been received, and at least one packet needs transmission, the WD2501 will send the next packet at T1 intervals with P = 1. As long as subsequent transmissions of the next packet receives an RNR, the WD2501 will not stop after N2 times.

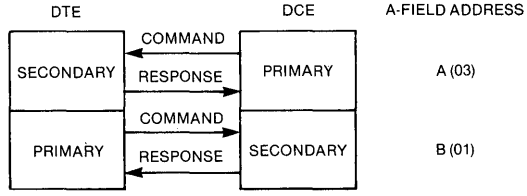

SYMMETRIC DTE-DCE LINK USED IN LAP

Figure 7. LAP COMMANDS AND RESPONSES

(Bit 0 is transmitted first)
Only the CMDR and I-frame contain I-fields.

NOTE: The WD2501 will accept a DM response to a transmitted DISC.

FRAME TYPE	COMMAND	RESPONSE	BIT #								
			7	6	5	4	3	2	1	0
INFORMATION (I)	I-FRAME OR PACKET		N(R)			P	N(S)			0
UNNUMBERED (U)	SARM		0	0	0	P	1	1	1	1
	DISC		0	1	0	P	0	0	1	1
		UA	0	1	1	F	0	0	1	1
		CMDR	1	0	0	F	0	1	1	1
SUPERVISORY (S)		RR	N(R)			F	0	0	0	1
		RNR	N(R)			F	0	1	0	1
		REJ	N(R)			F	1	0	0	1

TRANSMISSION OF ABORT

An ABORT (a 0 followed by at least seven 1's) is transmitted to terminate a frame in such a manner that the receiving station will ignore the frame. There are three conditions which will cause the WD2501 to transmit an ABORT:

1. An ABORT is sent when there is a Transmitter Under-Run
2. While transmitting a packet, if a REJ S-frame is received, the packet is ABORTed.
3. If T1 times-out while a packet is being transmitted, the packet is ABORTed. Caution: If a packet is longer (in time) than T1, the packet will always be ABORTed.

SELF-TESTS

There are two self-test features: 1) Internal RAM Register Test and 2) Loop-Back Test. Both tests are suitable for manufacturing testing, user incoming inspection testing, or system diagnostics and trouble-shooting.

INTERNAL RAM REGISTER TEST

There are eleven 8-bit registers internal to the WD2501 which are not directly accessible by the user's CPU. Seven of these registers can be tested by the Loop-Back Test. This test provides a means to check the other four registers.

The contents of Register A are placed in two even internal registers, and the contents of Register B in two odd internal registers. The four registers are then added together without carry and the result is placed in Registers 2, 5, 6 and 7. This test is initiated when RAMT (CR02) = 1. Use the following procedure:

1. Set-up Registers A and B.
2. Set RAMT.
3. Wait at least 50 times the CLK period.
4. Read Registers 2, 5, 6 and 7.

To repeat the test for new values in Registers A and B:

5. Clear RAMT.
6. Wait at least 100 times the CLK period.
7. Go back to step 1.

LOOP-BACK TEST

The loop-back may be internal (CR03 = 1) or external (CR03 = 0). Of course, if external, RD and TD must be tied together either directly or remotely.

If CR03 = 1, TD is internally tied to RD, and the pin at RD (16) is internally disconnected. Also, TC is internally tied to RC, and the pin at RC (17) is internally disconnected.

WD2501 ELECTRICAL SPECIFICATIONS:

ABSOLUTE MAXIMUM RATINGS:

Voltages referenced to V_{SS}

High Supply Voltage (V_{DD}) . . . . . −0.3 to 15V
Voltage at any Pin . . . . . −.03 to 15V
Operating Temperature Range . . . . . 0°C to +70°C
Storage Temperature Range . . . . −55°C to +125°C

NOTE:
Maximum limits indicate where permanent device damage occurs. Continuous operation at these limits is not intended and should be limited to those conditions specified in the DC Electrical characteristics.

Operating DC Characteristics: V_{SS} = 0V, V_{CC} = 5.0V ± .25, V_{SS} = 12.0V ± .6V T_A = 0° to 70°C

SYMBOL	PARAMETER	MIN	TYP	MAX	UNIT	CONDITIONS
I_{DD}	V_{DD} Supply Current		20	70	mA	
I_{CC}	V_{CC} Supply Current		200	280	mA	
V_{DD}	High Voltage Supply	11.4	12	12.6	V	
V_{CC}	Low Voltage Supply	4.75	5	5.25	V	
V_{IH}	Input High Voltage	2.4			V	
V_{IL}	Input Low Voltage			0.8	V	
V_{OH}	Output High Voltage	2.8			V	I_O = −0.1mA
V_{OL}	Output Low Voltage			0.4	V	I_O = 1.6mA
I_{LH}	Input Leakage Source or Sink			10	uA	
I_{OZH}	Input Leakage High Impedance			10	uA	Vin = V_{CC}
I_{OZL}	Output Leakage High Impedance			10	uA	Vin = V_{SS}

AC Timing Characteristics: V_{CC} = 5V ± .25; V_{DD} = RV ± .6V, V_{SS} = 0V, T_A = 0° to 70°C

SYMBOL	PARAMETER	MIN	TYP	MAX	UNIT	CONDITIONS
CLK	Clock Frequency	0.5		2.0	MHz	Note 1, 2
RC	Receive Clock Range	0			MHz	Note 3, 4
TC	Transmit Clock Range	0			MHz	Note 4
MR	Master Reset Pulse Width	10			mS	
T_{AR}	Input Address Valid to $\overline{RE}$	0			nS	
T_{RD}	Read Strobe (or $\overline{DACK}$ Read) to Data Valid			375	ns	Note 5
T_{HD}	Data Hold Time From Read to Strobe			80	nS	
T_{HA}	Address Hold Time From Read Strobe	80			ns	
T_{AW}	Input Address Valid to Trailing Edge of $\overline{WE}$	200			nS	
T_{WW}	Minimum $\overline{WE}$ Pulse Width	200			nS	
T_{DW}	Data Valid to Trailing Edge of $\overline{WE}$ or Trailing Edge of $\overline{DACK}$ for DMA Write	100			nS	
T_{AHW}	Address Hold Time After $\overline{WE}$	80			nS	
T_{DHW}	Data Hold Time After $\overline{WE}$ or After $\overline{DACK}$ for DMA Write	100			nS	
T_{DA1}	Time From $\overline{DRQO}$ (or $\overline{DRQI}$) to Output Address Valid if ADRV = 1			80	nS	
T_{DA0}	Time From $\overline{DACK}$ to Output Address Valid if ADRV = 0			400	nS	Note 5
T_{DD}	Time From Leading Edge of $\overline{DACK}$ to Trailing Edge of $\overline{DRQO}$ (or $\overline{DRQI}$)			400	nS	Note 5
T_{DAH}	Output Address Hold Time From $\overline{DACK}$			100	nS	
T_{DMW}	Data Hold Time From $\overline{DACK}$ For DMA Read			100	nS	
T_{RP1}	$\overline{REPLY}$ Response Time (leading edge)			240	nS	Note 5
T_{RP2}	$\overline{REPLY}$ Response Time (trailing edge)			260	nS	Note 5

NOTES:
1. Clock must have 50% duty cycle.
2. Buffer chaining is not guaranteed when CLK is greater than 1.5 MHz.
3. Residual bit detection logic not guaranteed when RC is greater than 50Kbps.
4. See "Ordering Information" for maximum serial rates.
5. C(load) = 100pf

DMA TIMING

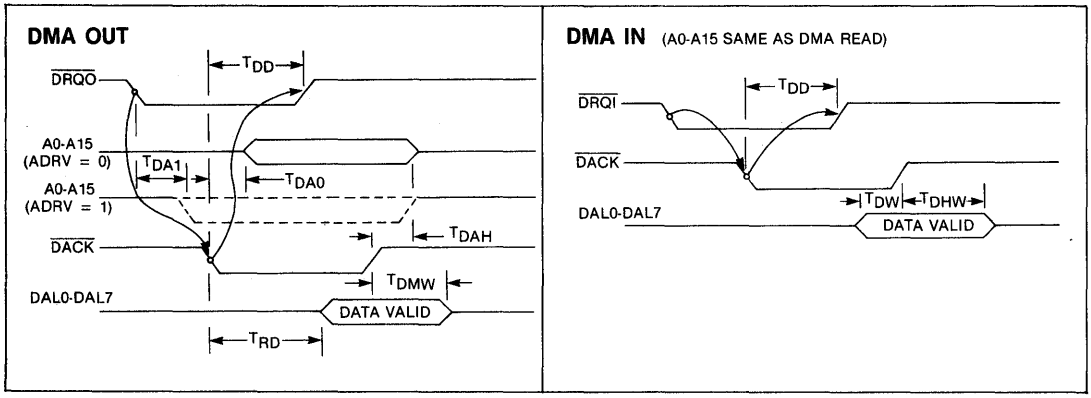

CPU READ/WRITE TIMING

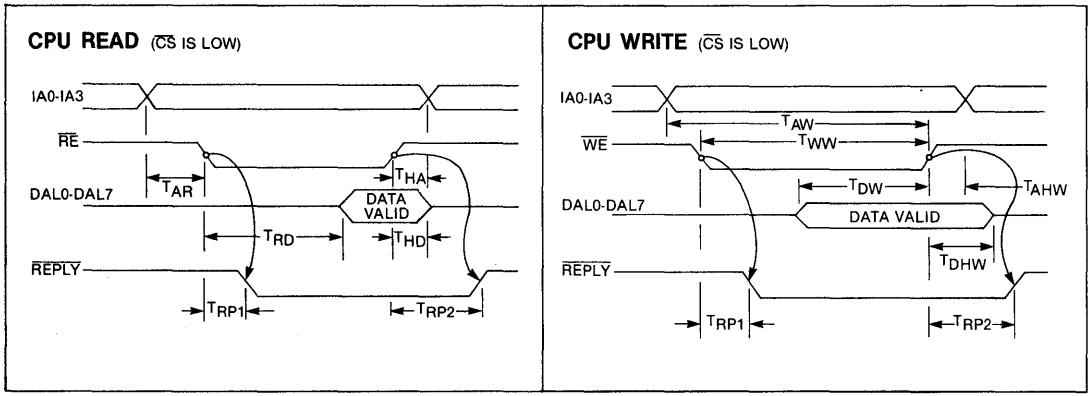

***NOTE:** There must not be a CPU read or write ($\overline{CS}$-$\overline{RE}$ or $\overline{CS}$-$\overline{WE}$) within 500 nanoseconds after the trailing (rising) edge of $\overline{DACK}$.

There must not be the leading (falling) edge of $\overline{DACK}$ allowed within 500 nanoseconds after the completion of a CPU write ($\overline{CS}$-$\overline{WE}$).

$\overline{RE}$ and $\overline{WE}$ must not be low at the same time.

APPENDIX A
STATES OF THE WD2501

APPENDIX A

LINK DOWN

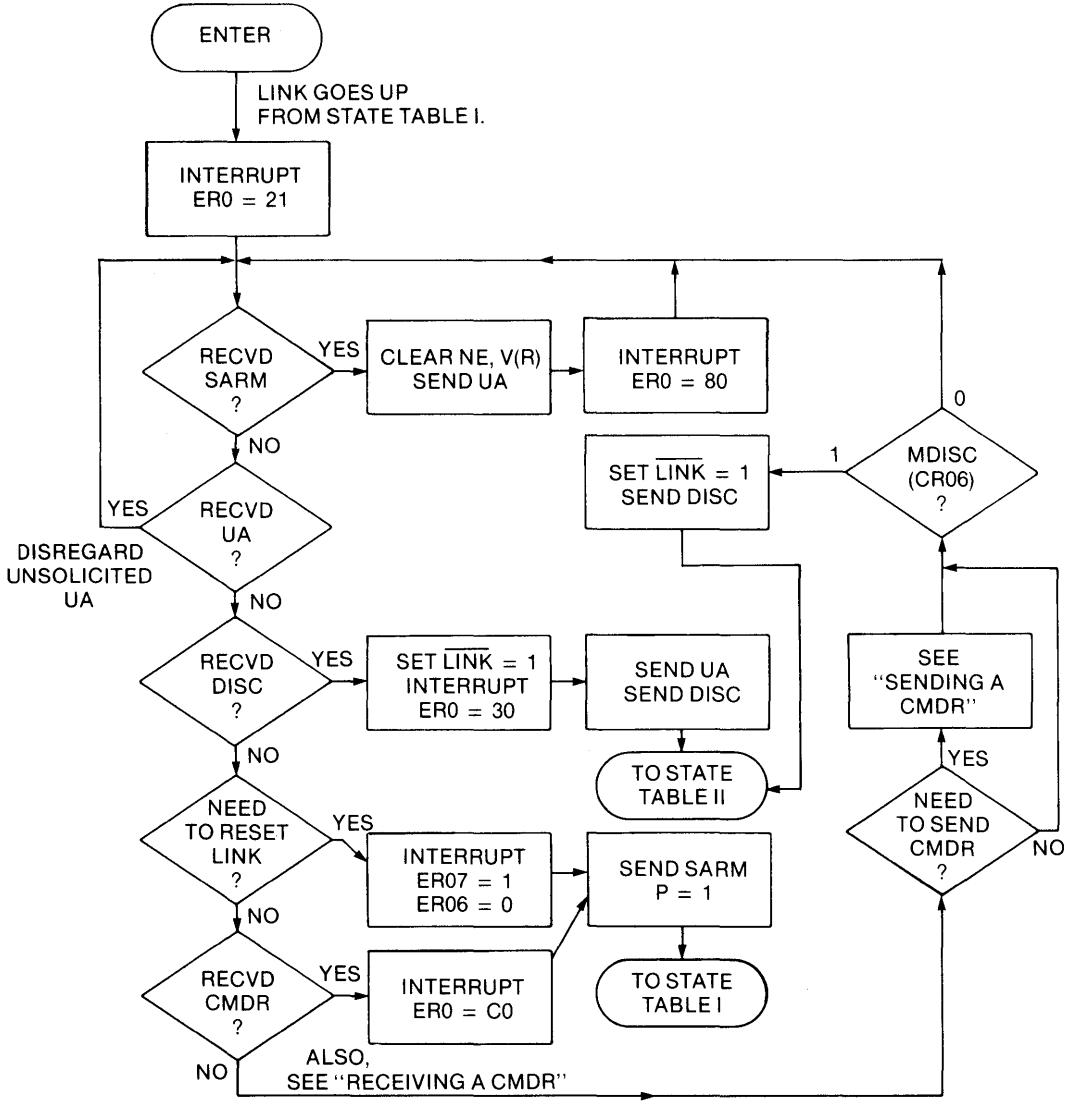

LINK UP
CONDITIONS FOR RESET OR DISCONNECT

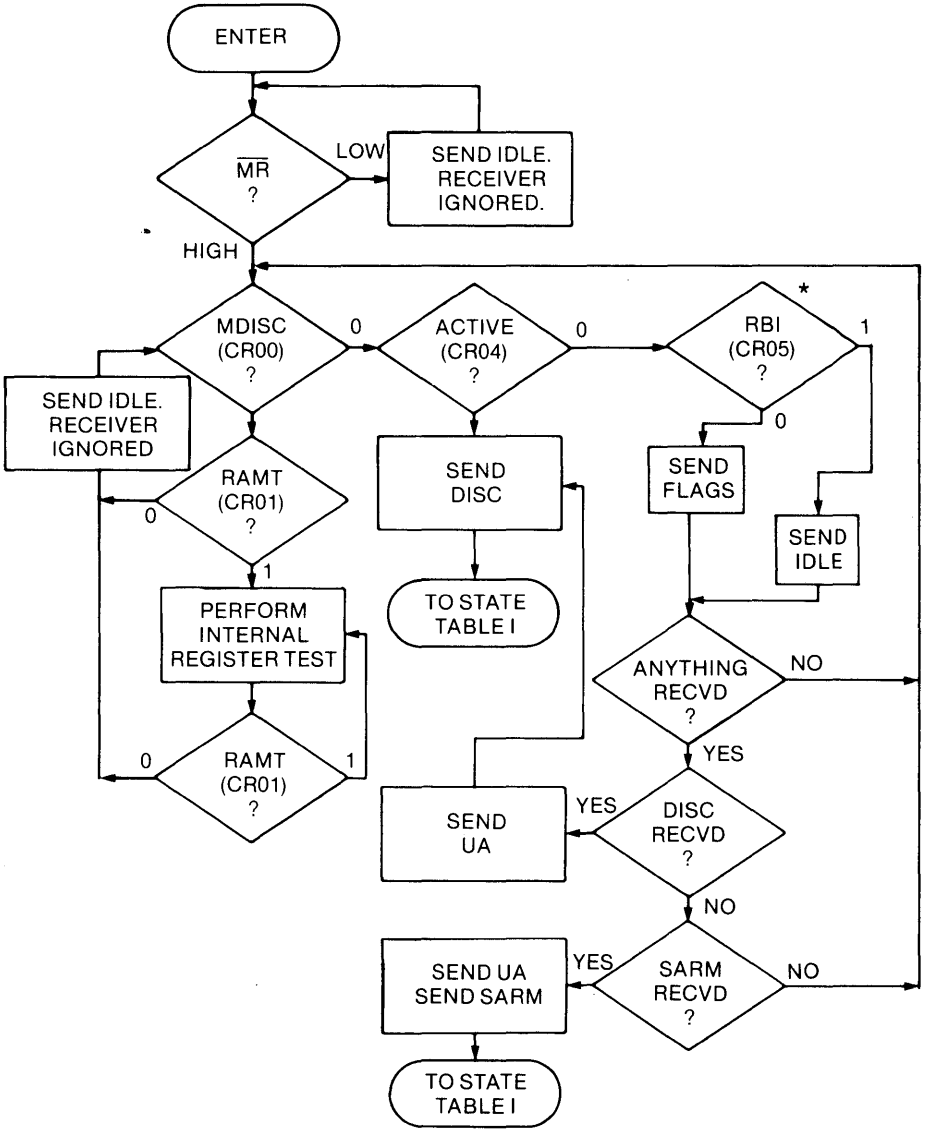

*ONCE A FRAME IS SENT, INTERFRAME TIME-FILL IS FLAGS EVEN IF RBI = 1. ALSO, ONCE FLAGS ARE STARTED, IDLE IS NOT SENT AGAIN UNTIL A COMPLETE DISCONNECT SEQUENCE OR MASTER RESET (MR LOW).

STATE TABLE I

LINK DOWN, BUT GOING UP

(Column 2 also applies to link-up, but reset transmitted.)

	1.	2.	3.
	DISC sent. Waiting for UA or DM.	SARM sent. Waiting for UA.	SARM sent. UA RECVD. Waiting for SARM.
T1 runs out, but not N2	Re-send DISC P = 1	Re-send SARM P = 1	—
N2 runs out	Re-send DISC P = 1	Send DISC. INTR. ERO = 24. To column 1.	Send DISC. INTR. Error Code 24. To column 1.
RECVD UA	Send SARM To column 2.	Clear NA, NB, V(S). If SARM RECVD and UA sent. To link-up flow. Else, to column 3.	Disregard
RECVD DISC	Send UA	Send UA	Send UA
RECVD SARM	Disregard	Send UA Clear NE, V(R).	Clear NE, V(R). Send UA. To link-up flow.
RECVD something, but not UA, DISC, SARM, nor DM.	Disregard	Disregard	Disregard
RECVD DM	Send SARM. To column 2.	Disregard	Disregard

STATE TABLE II

LINK GOING DOWN (WAS UP)

	1.	2.	3.
	DISC sent. Waiting for UA.	MDISC (CR00) set.	DISC sent. UA or DM RECVD. Sending flags. Link is down.
T1 runs out	Re-send DISC. P = 1	Send DISC to column 1	—
RECVD UA or DM	If DISC RECVD and UA send to link down flow. If not, to column column 3.	—	Disregard
RECVD SARM	Disregard	—	Send UA. Send SARM. To Table I.
RECVD DISC	Send UA	—	Send UA. To link down flow.
RECVD something but not UA, DISC, SARM, nor DM.	Disregard	—	Disregard

NOTE: If while the link is up, the MDISC bit is set, the WD2501 will send a DISC and wait for a UA. After receiving the UA, the WD2501 does not require a DISC to be received. The link is considered down.

If while the link is up, a DISC is received, but MDISC = 0, the WD2501 will send a UA immediately followed by a DISC being sent. Appropriate action shown in column 1 above is taken.

USE OF FLAGS BY THE WD2501

The WD2501 will send interframe time fill flags, and will send flags when the link is down if RBI = 0 and/or ACTIVE = 1 (see Link Down Flow). The WD2501 does not require the remote station to send flags nor interframe time fill of flags. The chip will accept idle or flags. However, if the receiver is idle for T1xN2, the WD2501 will disconnect the link.

The chip will send contiguous flags of:

01111110011111110011111...

The chip will accept either the above sequence as contiguous flags or:

01111110111111110111...

DEFINITIONS OF COMMAND AND RESPONSE

A transmitted or received command or response is a frame with the A-field as defined below:

FRAME	A-FIELD
Transmitted Command	Register E
Received Command	Register F
Transmitted Response	Register F
Received Response	Register E

Only command frames or response frames are transmitted. If a received frame is neither a command nor response, it is disregarded ("thrown-away") at the receiver.

SENDING I-FRAME (PACKETS) **NOTE:** In all subsequent pages, the link is considered UP (LINK = 0)

SEND (CR10)	BRDY OF TLOOK #NB	NA AND NB	RNRR(2) 0	T1 RUNS OUT(1) BUT NOT N2	RECVD REJ(2)	
1	0	Don't Care	0	No	No	Clear SEND (CR10)
1	1	Don't Care	0	No	No	Send next packet with N(S) = NB. After transmission complete, increment NB. Exception: if NB + 1 = NA, do not send next packet. There are 7 outstanding packets.
0	1 or 0	Equal	1	No	No	No packets to send.
0→1	1	Equal	1	No	No	Send next packet with N(S) = NB. After transmission complete, increment NB, but do not send next packet.
1	1 or 0	Not Equal	1	No	No	Do not send packet.
1 or 0	1 or 0	Not Equal	1 or 0	Yes	No	Retransmit packet with N(S) = NA and P = 1. When transmission complete, do not go to next packet.(4)
1 or 0	1 or 0	Not Equal	0	No	Yes	Make NA = received N(R). Start sequential retransmission of packets with N(S) = NA.

NOTES:

1. This table applies to T1 running out only if started due to at least one packet being outstanding.
2. Received S-frames shown in this table are taken to have valid N(R)'s.
3. When an acknowledgement is received, NA is made equal to the received N(R). All TLOOK segments from the

NOTES CONTINUED:

old value of NA up to N(R) – 1 are acknowledged and the appropriate TLOOK segments will have the ACKED bit set. After setting the ACKED bit(s), an XBA interrupt is generated.

4. Assuming appropriate TLOOK segments are ready, packets are sent sequentially without waiting for an acknowledgement with three exceptions:
   a. Seven outstanding (unacknowledged) packets have been transmitted.
   b. The remote station has indicated a busy condition as noted by the receiving of an RNR. The next unacknowledged packet is transmitted, if not previously transmitted, and T1 is started.
   c. If T1 expires (runs-out), the last unacknowledged packet is retransmitted with P = 1. This is the timer recovery condition and remains in effect until an S-frame with F = 1 and acknowledging that at least one packet is received, or until a link reset is transmitted, or until a DISC is received.

   If an acknowledgement is not received after N2 tries, a SARM is transmitted with P = 1. An interrupt is generated with ERO = 90.

RECEIVING A NULL PACKET

If an error-free packet is received with no I-field but with an in-sequence N(S), it will be treated the same as a packet with an I-field. RCNT will be all 0's.

The packet will be acknowledged at the next opportunity just as if the I-field was present.

The received N(R) of the null packet is recognized as an acknowledgement of a previously transmitted packet(s). If the N(R) is not valid, the link will be reset.

The WD2501 will not transmit a null packet. TCNT must not be all 0's.

SENDING A REJ

1. The reject condition is entered any time an error-free I-frame with an out-of-sequence N(S) is received.
2. A REJ frame with N(R) = V(R) is transmitted immediately if the WD2501 is not sending an I-frame, or immediately after the completion of the I-frame. There are two exceptions to this in notes 3 and 4 below.
3. If a link resetting SARM needs to be transmitted, the SARM is sent ahead of the REJ. Immediately after the SARM, the REJ is sent. As an example, if an I-frame is received with out-of-sequence N(S) and invalid N(R), a SARM will be sent followed by an REJ.
4. If the receiver is not ready (RNRX = 1), the REJ is not sent until RNRX is cleared.
5. Once the reject condition is entered, only one REJ is sent. Another REJ is not sent unless the reject condition is cleared and re-entered. The reject condition is cleared if an I-frame is received with an in-sequence N(S) [i.e., received N(S) = V(R)], or if a SARM is received.

   When the REJ is sent, error counter #6 is incremented.

RECEIVING A REJ

If a REJ is received error-free with no I-field:

1. If the N(R) is invalid, an interrupt is generated with ERO = 88, and a SARM is transmitted to reset the link. (See Table I, column 2.)
2. If the N(R) is valid and greater than NA, at least one transmitted I-frame is acknowledged. An XBA interrupt is generated with the appropriate ACKED bits in TLOOK set.
3. If the N(R) is valid and less than NB, or equal to NB but an I-frame is being transmitted, the WD2501 will begin sequential re-transmission starting with V(S) = received N(R). If an I-frame is in transmission when the REJ was received, it is aborted.
4. If the N(R) is valid and equal to NB and there is no I-frame being transmitted, no further action is taken. In this case, the REJ has the same effect as an RR.

DEFINITION OF VALID RECEIVED N(R)

Reference: X.25 paragraph 2.4.8.1 and paragraph 2.3.4.10

DEFINITION: A valid received N(R) is greater than or equal to NA and less than or equal to NB. An invalid received N(R) is one which does not conform to this rule.

NOTES:
1. This must be understood in a circular sense. 0 could be greater than 7, depending on the values of NA and NB.
2. If NA = NB, there is only valid possibility for N(R).
3. If NB + 1 = NA, there are seven outstanding packets, and it is impossible to receive an invalid N(R).
4. Basically, an invalid N(R) is one which acknowledges a packet which was never sent.

SENDING AN RNR (RNRX)

LINK	RECR (CR01)	REC RDY (NE)	ACTION
1	—	—	No S-frame transmitted when link down.
1→0	1	1	RLOOK ready. No S-frame sent immediately.
1→0	1	0	Interrupt: RLNR. RNR sent immediately. (2) RNRX set.
1→0	0	—	RNR sent immediately. RNRX set.
0	1	1	Receiver ready to accept I-frames.
0	1→0	1	Receiver ready to accept I-frames.
0	1	0	Interrupt: RLNR. RNR sent. RNRX set.
0	0→1	0	Next received I-frame not brought into memory. If P = 1, will send RNR with F = 1. RNRX was previously set.
0	0→1	1	Next received I-frame not brought into memory. Immediately after frame received, RNRX cleared. After that, an RR or REJ is sent.
0	1→0	0	Interrupt: RLNR. RNR sent. RNRX was previously set.
0	0	0	RNRX set. RNR sent. There is no RLNR interrupt.

NOTES ON RNRX:

1. The arrows (→) indicate a change in state from the value on the left to the value on the right.
2. If an RNR is sent "immediately," it means immediately after link is brought up.
3. After the link is up, if an RR, RNR or REJ is sent, it is sent at the next opportunity. This means immediately if no I-frame is being sent, or at the end of the current I-frame being transmitted.
4. The RNRX (SR00) status bit is set at the time the receiver-not-ready condition was established. The RNR may be sent later according to note 3 above.
5. When a received I-frame is brought into memory, it may be accepted provided the FCS and N(S) are good and the I-field is not too long. The N(R) may or may not be correct, but is checked separately. If the N(R) is bad, a link reset is transmitted.
6. Whenever RNRX = 1, the I-field of received I-frames is not brought into memory. The received N(S) and N(R) are checked as usual. (See note 9 below.)
7. If, while RNRX = 1, a link resetting SARM is received, a UA will be transmitted. However, an RNR will not be sent immediately after the UA. If an I-frame is received with P = 1, then an RNR will be sent with F = 1.
8. If a link resetting SARM is transmitted when RNRX = 1, RNRX will be cleared when the UA is received. If the condition which caused received-not-ready still exists, an RNR is sent and RNRX is set according to the RNRX table. If the receiver is ready, I-field data may be brought into memory.
9. If, while RNRX = 1, an out-of-sequence N(S) is received, the WD2501 will not immediately send an REJ. After RNRX is cleared, an REJ will be transmitted at the next opportunity.

RECEIVING AN RNR (RNRR)

SEND (CR10)	NA AND NB	RECVD ACK?	RECVD RNR	RECVD RR, REJ OR UA	T1 EXPIRES	
1	Not Equal	YES	YES	NO	NO	Set RNRR. Restart T1 and N2. Update NA.
0	Not Equal	YES	YES	NO	NO	Set RNRR. Restart T1 and N2. Update NA. If after update NA = NB, stop T1.
0	Equal	NO	YES	NO	NO	Set RNRR.
1 or 0	Not Equal	NO	NO	NO	YES	Re-transmit last un-acknowledged I-frame with P = 1. Start T1. If RNR subsequently received restart T1 and N2.
1 or 0	Not Equal	YES	NO	YES (but not UA)	NO	Clear RNRR. Restart T1 and N2. Update NA. If after update NA = NB, stop T1.
1 or 0	Don't Care	Don't Care	NO	YES	NO	Clear RNRR.
0 → 1	Equal	NO	NO	NO	NO	Send next I-frame. Increment NB after transmission. (Then, NB does not NA). Start T1 and N2.

NOTES:

1. If SEND = 1, it is assumed that BRDY (NA) is set. Otherwise, SEND is cleared.
2. If RNRR = 1, only one I-frame is sent per T1 interval.

CONDITIONS FOR SENDING LINK RESET

1. CMDR received.
2. Received response, other than CMDR, which contained an I-field.
3. Received I-frame or S-frame with invalid N(R).
4. Received response with unsolicited F-bit.
5. Have sent I-frame N2 times without acknowledgement. Exception: An RNR frame received will reset the N2 counter.

SENDING A CMDR

A CMDR may be sent for any of the reasons indicated in X.25. If an invalid N(R) is received, the link is reset whether the frame was a command or response.

Upon sending a CMDR, the WD2501 will not allow I-field data to come into the receiver DMA. However, it may send I-frames (packets). If an I-frame is received, the WD2501 will re-transmit the CMDR. If a SARM is received, the WD2501 will send a UA and clear NE, and I-field data may again be brought into memory by the receiver DMA (provided RECR = 1 and REC RDY of segment #0 is set to 1). If a DISC is received, the WD2501 will send a UA followed by a DISC and go to State Table II, column 1.

If a supervisory response is received, the WD2501 will accept any acknowledgements contained in the received frames N(R), and the CMDR is retransmitted.

RECEIVING A CMDR

Upon the reception of a CMDR:

1. The CMDR I-field will be in the memory referenced by the current NE segment, provided the receiver was ready.
2. The SEND (CR10) bit is cleared.
3. No more I-field data will be brought into memory by the receiver.
4. A link resetting SARM is transmitted and an error interrupt is generated at the beginning of the SARM. The hexidecimal value of ER0 is C0.

5. When a UA is received to the SARM, NA and NB are cleared to 0. I-field data may not come into memory provided RECR = 0. If RECR goes from 0 to 1, and after an I-frame is received, the next I-field may be brought into memory.

NOTE: The reason for step 5 is to allow the CPU time to initialize the transmitter in such a manner so as to not duplicate I-frames. The CPU should set-up the next transmission according to the received V(R) in the received I-field. The reason for step 2 is to prevent out-of-sequence transmissions.

Let Vc(R) = V(R) in I-field of received CMDR.

If Vc(R) > NA, a transmitted block or blocks have been acknowledged from NA and up to Vc(R) − 1. Vc(R) is found in the current RLOOK buffer.

NOTE: After receiving a CMDR, the CPU has a certain amount of time to read the current values of NA and NB. This minimum time is the time it takes to send a SARM and receive a UA. Immediately after the UA is received, NA and NB are cleared.

The CPU should take the packet referenced by Vc(R) and puts its address and count information in TLOOK segment #0. Vc(R) + 1 goes into segment #1 and so on up to NB − . The CPU should await for NA and NB to clear before setting the SEND (CR10) control bit.

EXAMPLE: Suppose four packets are outstanding with N(S) counts 3, 4, 5 and 6. Number 3 is a Data Packet on Logical Channel (LC) 14. Number 4 is a Call Request Packet on LC 220. Number 5 is a Reset Packet on LC 7. Number 6 is a Data Packet on LC 9. A CMDR is received, and the V(R) in the I-field is 5. Thus, numbers 3 and 4 are acknowledged. At this point NA = 3 and NB = 7. The CPU will put the Reset Packet on LC 7 in #0, and the Data Packet on LC 9 in #1. After NA and NB are cleared, the CPU may set SEND (CR10) and transmission will start with N(S) = 0, which is the Reset Packet on LC 7.

PROT COL SIGNIFICANCE TLOOK/RLOOK POINTERS

In addition to the use of the TLOOK/RLOOK pointers as used in the Memory Access Method, the pointers have a significance with the link level protocol.

RLOOK pointer NE = V(R) at all times

TLOOK pointer NB = V(S) when not in retransmission due to time-out recovery or received REJ.

TLOOK pointer NB = X when in retransmission

TLOOK pointer NA = last un-acknowledged packet. When a frame with a valid N(R) is received, NA is made equal to the N(R). If NA changed, an XBA interrupt is generated.

PROTOCOL SIGNIFICANCE OF TLOOK/RLOOK POINTERS

The NE, NA, and NB pointers have a relationship with the sequence counters used in the LAPB protocol.

The RLOOK pointer NE is equal to V(R) at all times if TRCV = 0. However, when TRCV = 1 when the link is UP, there is no guaranteed relationship between NE and V(R).

TLOOK pointer NB is the Next Block to be transmitted. If not in packet retransmission, NB is equal to the V(S) of the next new packet to be transmitted.

TLOOK pointer NA is the Next packet to be Acknowledged. It represents the V(S) number for the oldest packet in the retransmission buffer.

USE OF THE RECR BIT

The RECR (CR01) bit should be understood as an instruction to the WD2511 to initialize the receiver memory. The WD2511 will test RECR as soon as MDISC is cleared, and will test RECR after each link set-up and each link reset. Once the receiver memory is ready, the WD2511 will not test RECR again unless there is a link set-up, link reset, or a receiver-not-ready condition.

After the link is UP, and at least one packet has been received as indicated by the PKR interrupt, the user should clear RECR. This is an advantage, because if a link reset is either transmitted or received, the WD2511 will enter a receiver-not-ready condition. This will prevent packets received after the link reset from appearing to have arrived before some packets received prior to the link reset.

The receiver-not-ready condition is indicated by RNRX = 1. This condition is cleared after the user makes RECR = 1 with RECRDY = 1 (in RLOOK #0), and after either a packet or an S-frame is received.

If RECRDY of the next RLOOK is 0, but RECR = 0, there will not be an RLNR interrupt, but RNRX will be set. If RECR = 1, but the RECRDY bit of the next RLOOK segment is 0, there will be an RLNR interrupt (error code 10), and RNRX will be set.

ORDERING INFORMATION

Order Number	Maximum Rate
WD2501T-01	100 Kbps
WD2501T-05	500 Kbps
WD2501T-11	1.1 Mbps*

*Higher speeds available on special order.

The following devices have a tighter power supply specification and meet all other parameters in this data sheet.

Power Supply Range +5/−2%:

Order Number	Maximum Rate
WD2501T-91	100 Kbps
WD2501T-92	500 Kbps
WD2501T-93	1.1 Mbps

Printed in U.S.A.

WD2511 X.25 Packet Network Interface (LAPB)

FEATURES

- Packet switching controller, compatible with CCITT recommendation X.25, level 2, LAPB (WD2511)
- Programmable primary timer (T1) and retransmission counter (N2)
- Programmable A-field which provides a wider range of applications than defined by X.25. These include: DTE-to-DTE connection, multipoint and loop-back testing.
- Direct memory access (DMA) transfer: two channels; one for transmit and one for receive. Send/receive data accessed by indirect addressing method. Sixteen output address lines.
- Zero bit insert and delete
- Automatic appending and testing of FCS field
- Computer bus interface structure: 8 bit bi-directional data bus. CS, WE, RE and four input address lines
- DC to 1.1 MBPS data rate
- TTL compatible
- 48 pin dual in-line packages
- Higher bit rates available by special order

APPLICATIONS

X.25 PACKET SWITCHING CONTROLLER
PART OF DTE OR DCE
PRIVATE PACKET NETWORKS
LINK LEVEL CONTROLLER

GENERAL DESCRIPTION

The WD2511 is a MOS/LSI device which handles bit-oriented, full-duplex serial data communications with DMA, which conforms to CCITT X.25 with programmable enhancements.

The device is fabricated in N-Channel silicon gate MOS technology and is TTL compatible on all inputs and outputs.

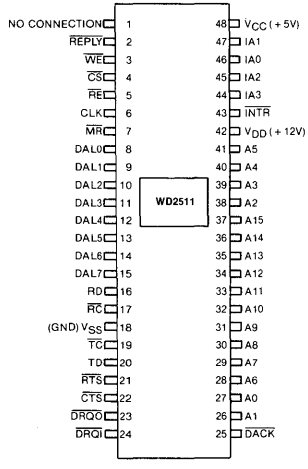

PIN DESIGNATION

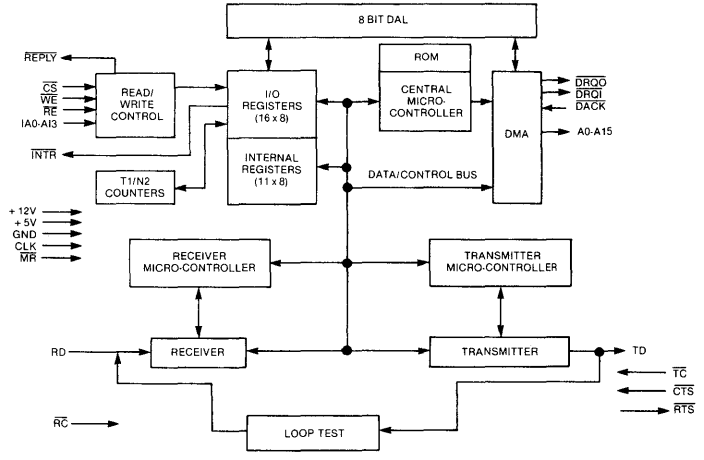

BLOCK DIAGRAM

INTERFACE SIGNALS DESCRIPTION (All signals are TTL compatible.)

PIN NUMBER	SYMBOL	PIN NAME	FUNCTION
1		No Connection	Leave pin open
2	$\overline{REPLY}$	Reply	An active low output to indicate the WD2511 has either a $\overline{CS}$•$\overline{RE}$ or a $\overline{CS}$•$\overline{WE}$ input.
3	$\overline{WE}$	Write Enable	The data on the DAL are written into the selected register when $\overline{CS}$ and $\overline{WE}$ are low.
4	$\overline{CS}$	Chip Select	Active low chip select for CPU control of I/O registers.
5	$\overline{RE}$	Read Enable	The contents of the selected register are placed on DAL when $\overline{CS}$ and $\overline{RE}$ are low.
6	CLK	Clock	Clock input used for internal timing. Must be square wave, and greater than 250 KHz.
7	$\overline{MR}$	Master Reset	Initialize on active low. All registers reset to zero, except control bits MDISC and $\overline{LINK}$ are set to 1. $\overline{DACK}$ must be stable high before $\overline{MR}$ goes high.
8-15	DAL0-DAL7	Data Access Lines	An 8-bit bi-directional three-state bus for CPU and DMA controlled transfers.
16	RD	Receive Data	Receive serial data input.
17	$\overline{RC}$	Receive Clock	This is a 1x clock input, and RD is sampled on the rising edge of $\overline{RC}$. RD changes occur on the falling edge of $\overline{RC}$.
18	V_{SS}	Ground	Ground.
19	$\overline{TC}$	Transmit Clock	A 1x clock input. TD changes on the falling edge of $\overline{TC}$.
20	TD	Transmit Data	Transmitted serial data output.
21	$\overline{RTS}$	Request-To-Send	An open collector (drain) output which goes low when the WD2511 is ready to transmit either flags or data.
22	$\overline{CTS}$	Clear-To-Send	An active low input which signals the WD2511 that transmission may begin. If high, the TD output is forced high. May be hard-wired to ground.
23	$\overline{DRQO}$	DMA Request Out	An active low output signal to initiate CPU bus request so the WD2511 can output onto the bus.
24	$\overline{DRQI}$	DMA Request In	An active low output signal to initiate CPU bus request so that data may be input to the WD2511. $\overline{DRQO}$ and $\overline{DRQI}$ will not be low at the same time.
25	$\overline{DACK}$	DMA Acknowledge	An active low input from the CPU in response to $\overline{DRQI}$ or $\overline{DRQO}$. $\overline{DACK}$ must not be low if $\overline{CS}$ and $\overline{RE}$ are low or if $\overline{CS}$ and $\overline{WE}$ are low.
26-41	A0-A15	Address Lines Out (See front page for Pin Assignments)	Sixteen address outputs from the WD2511 for DMA operation. If the control bit ADRV is 1, the outputs are TTL drives at all times. If ADRV is 0, the outputs are three-state, and are HI-Z whenever $\overline{DACK}$ is high. (ADRV is in Control Register #1.)
42	V_{DD}	Power Supply	+ 12VDC power supply input.

INTERFACE SIGNALS DESCRIPTION CONTINUED (All signals are TTL compatible.)

PIN NUMBER	SYMBOL	PIN NAME	FUNCTION
43	$\overline{INTR}$	Interrupt Request	An active low interrupt service request output, and returns high when Status Register #1 is read.
44-47	IA0-IA3	Address Lines In (See front page for Pin Assignments)	Four address inputs to the WD2511 for CPU controlled read/write operation with registers in the WD2511. If ADRV = 0, these may be tied to A0-A3. (ADRV is in Control Register #1.)
48	V_{CC}	Power Supply	+5VDC power supply input.

ORGANIZATION

A detailed block diagram of the WD2511 is shown in Figure 1.

Mode control and monitor of status by the user's CPU is performed through the Read/Write Control circuit, which reads from or writes into registers addressed by IA0-IA3.

Transmit and receive data are accessed through the DMA control. Serial data is generated and received by the bit-oriented controllers.

Internal Control of the WD2511 is by means of three internal microcontrollers; one for transmit, one for receive, and one for overall control.

Parallel transmit data are entered into the Transmitter Holding Register (THR), and then presented to the Transmitter Register (TR) which converts the data to a serial bit stream. The Cyclic Redundancy Check (CRC) is computed in the 16-bit CRC register, and the results become the transmitted Frame Check Sequence (FCS).

Parallel receive data enter the Receiver Holding Register (RHR) from the 24-bit serial Receiver Register (RR). The 24-bit length of RR prevents received FCS data from entering the RHR (unless residual bits are received). The receiver CRC register is used to test the validity of the received FCS. A 3-stack FIFO is included in the receiver.

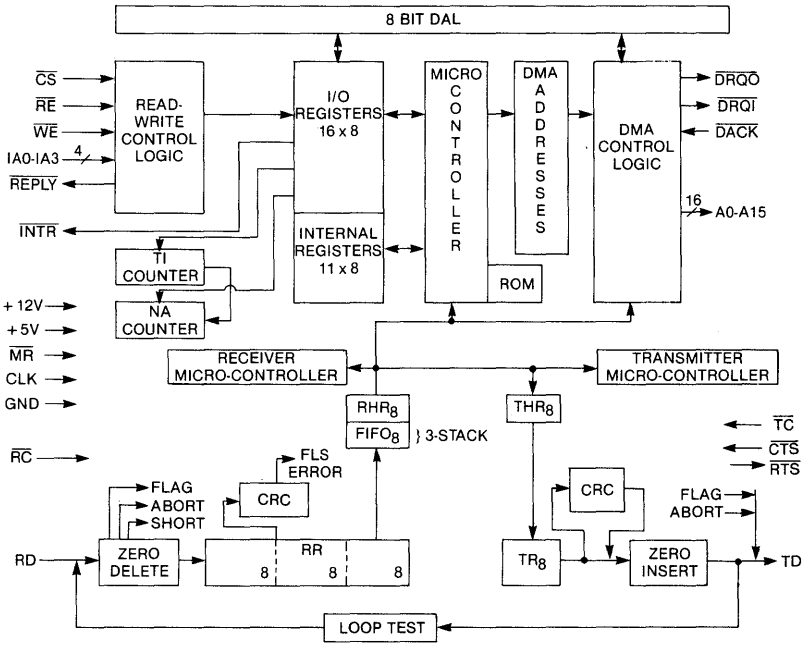

Figure 1 WD2511 BLOCK DIAGRAM — DETAILED

FRAME FORMAT

The WD2511 performs "bit-oriented" data communications control. According to the bit-oriented procedures (HDLC, SDLC, ADCCP), each serial block of data is called a frame.

Each frame starts and ends with a Flag (01111110). A single flag may be used both as the closing flag of one frame and the opening flag of the next frame. In between flags, data transparency is provided by the insertion of an 0 bit after all sequences of 5 contiguous 1 bits. The receiver will strip the inserted 0 bits. The last 16-bits before the closing flag is in the Frame Check Sequence (FCS). Each frame also includes address and control fields (A and C fields).

The FCS calculation includes all data between the opening flag and the first bit of the FCS, except for 0's inserted for transparency. The 16-bit FCS has the following characteristics:

Polynomial = $X^{16} + X^{12} + X^5 + 1$
Transmitted Polarity —Inverted
Transmitted Order —High Order Bit First
Preset Value —All I's

After the frame is received, if there were no errors, then the remainder in the CRC register (internal in the WD2511) will be:

1111000010111000 F0B8

The WD2511 generates and tests the Flag, FCS, A-Field, C-Field, and performs zero bit insertion and deletion. According to the X.25 protocol, there are three types of frames: supervisory (S-frame), unnumbered (U-frame), and information (I-frame).

The WD2511 performs frame level (level 2) link access control. All S- and U-frames are automatically generated and tested by the WD2511. The user need only be concerned with the I-frames, which are packets.

The WD2511 will transmit contiguous flags for interframe time fill.

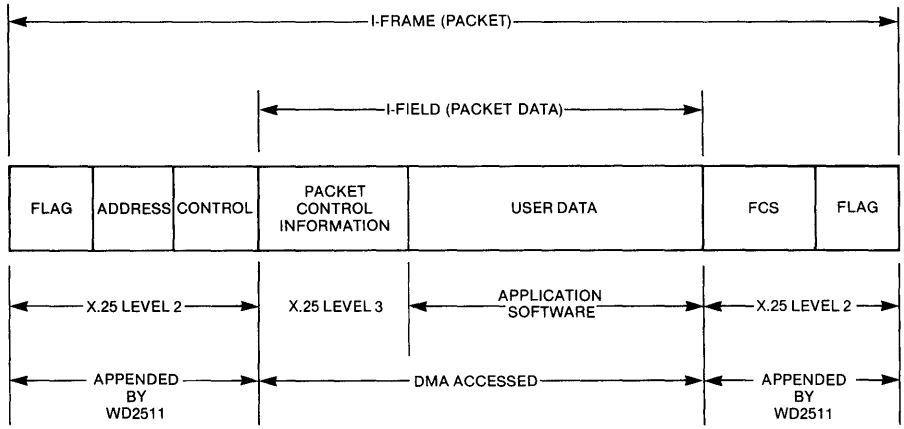

X.25 MODE

NOTE: X.25 Level 1, is the Physical Interface

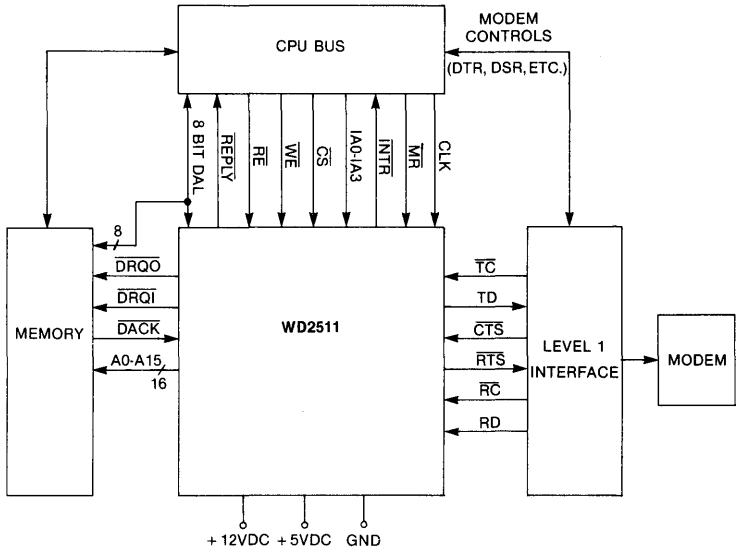

Figure 2. SYSTEM CONNECTION

II. PROGRAMMING THROUGH REGISTERS

The WD2511 is controlled and monitored by sixteen I/O registers.

Control, status, and error bits will be referred to as CR, SR, or ER, respectively, along with two digits. For example, SR16 refers to status register #1 and bit 6, which is "XBA."

REGISTER DEFINITION

REG #	IA3	IA2	IA1	IA0	REGISTER	REGISTER GROUPING
0	0	0	0	0	CR0	OVERALL CONTROL AND MONITOR
1	0	0	0	1	CR1	
2	0	0	1	0	*SR0	
3	0	0	1	1	*SR1	
4	0	1	0	0	*SR2	
5	0	1	0		*ER0	
6	0	1	1	0	*CHAIN MONITOR	RECEIVER MONITOR
7	0	1	1	1	*RECEIVED C-FIELD	
8	1	0	0	0	T1	TIMER
9	1	0	0	1	N2/T1	
A	1	0	1	0	TLOOK HI	DMA SET-UP
B	1	0	1	1	TLOOK LO	
C	1	1	0	0	CHAIN/LIMIT	
D	1	1	0	1	NOT USED	
E	1	1	1	0	XMT COMMAND "E"	"A" FIELD
F	1	1	1	1	XMT RESPONSE "F" (Note 1)	

*CPU READ ONLY. (Write Not Possible)

NOTE:
1. Registers E and F should be set-up while MDISC = 1.

CONTROL, STATUS, ERROR REGISTERS

REGISTER	7	6	5	BIT # 4	3	2	1	0
CR0	$\overline{\text{ADISC}}$	0	H/$\overline{\text{F}}$	ACTIVE/ $\overline{\text{PASSIVE}}$	LOOP TEST	2RAMT	RECR	MDISC
CR1	TXMT	TRCV	$\overline{\text{XI}}$	ADRV	0	0	0	SEND
SR0	NA2	NA1	NA0	RNRR	NB2	NB1	NB0	RNRX
SR1	1PKR	1XBA	1ERROR	0	NE2	NE1	NE0	0
SR2	T1OUT	$\overline{\text{IRTS}}$	REC IDLE	0	0	0	0	$\overline{\text{LINK}}$
ER0	ER07	ER06	ER05	ER04	ER03	ER02	ER01	ER00

1 Causes Interrupt ($\overline{\text{INTR}}$ Goes Low).

2 RAMT is used to initiate a self-test feature (see section on Self-Test). For normal operation, program to 0.

CONTROL REGISTER 0

REGISTER	CR07	CR06	CR05	CR04	CR03	CR02	CR01	CR00
CR0	$\overline{\text{ADISC}}$	0	H/$\overline{\text{F}}$	ACTIVE/ $\overline{\text{PASSIVE}}$	LOOP TEST	RAMT	RECR	MDISC

BIT	DESCRIPTION
CR00	MDISC is a mandatory disconnect command. MDISC will cause a logical disconnect in the link. No DMA accessed data will be transferred as long as MDISC = 1. After Master Reset (MR pin transition from low to high), MDISC will be set. The WD2511 will neither transmit nor accept received data until MDISC = 0.
CR01	This bit is RECR which defines the CPU's receiver buffer as initially Ready (CR01 = 1). If RECR = 1, this bit indicates that the WD2511 may begin receiving I-frames. (See SR00)
CR02	RAMT — This bit activates an internal register test. See Section on Self tests for description.
CR03	The LOOP TEST bit will connect the transmitted data output to the receiver input. The receiver input pins RD and $\overline{\text{RC}}$ are logically disconnected. The "E" and "F" registers of the A-field should be equal.
CR04	This bit will cause the WD2511 to initiate link set-up if CR04 = 1, or to wait for a link set-up from the remote device if CR04 = 0.
CR05	H/$\overline{\text{F}}$ selects full duplex if CR05 = 0, and half duplex if CR05 = 1. (See Appendix A).
CR06	Unused control bits, like CR06, should remain at 0.
CR07	$\overline{\text{ADISC}}$ is used when CR04 = 1 (ACTIVE). When the WD2511 actively initiates link set-up, a DISC will be transmitted and acknowledged prior to transmission of the SABM if CR07 = 0. Otherwise, the WD2511 will send the SABM for link set-up, but not preceed the SABM with a DISC if CR07 = 1.

CONTROL REGISTER 1

REGISTER	CR17	CR16	CR15	CR14	CR13	CR12	CR11	CR10
CR1	TXMT	TRCV	$\overline{XI}$	ADRV	0	0	0	SEND

BIT	DESCRIPTION
CR10	The SEND bit (CR10) is used to command the WD2511 to send the next packet or packets. If SEND = 1, the WD2511 will read from TLOOK the BRDY bit of the next segment for transmission. If BRDY = 0, the WD2511 will clear SEND and no action occurs. If BRDY = 1, the WD2511 will then read TSADR and TCNT, followed by the transmission of that buffer. After transmission, the WD2511 clears BRDY of the segment just transmitted, and reads BRDY of the next segment. If 1, the next segment is transmitted. If 0, the SEND bit is cleared, and transmission of packets is stopped. As a matter of good practice, the CPU should set SEND each time one of the BRDY bits is set.
CR11-13	Unused bits, write in 0's.
CR14	The ADRV bit (CR14) is the control for the 16 bit output addresses (A0-A15). If ADRV = 0, the outputs are tri-state and are in HI-Z, except when DACK goes low. If ADRV = 1, the outputs are always low impedance (TTL), and are high-level when DACK is high.
CR15	$\overline{XI}$ — Used when TXMT = 1 $\overline{XI}$ = 0 Frame ≥ 3 bytes excluding FCS $\overline{XI}$ = 1 Frame < 3 bytes excluding FCS
CR16	TRCV — Receive all frames including unknown frames. See Appendix A.
CR17	TXMT — Transparent transmit. See Appendix A.

STATUS REGISTER 0

REGISTER	SR07	SR06	SR05	SR04	SR03	SR02	SR01	SR00
SR0	NA2	NA1	NA0	RNRR	NB2	NB1	NB0	RNRX

BIT	DESCRIPTION
SR00	RNRX. An RNR has been transmitted, or will be at next opportunity. The CPU should set RECR, when receive buffers are available.
SR03-SR01	NB2-NB0. Next block to be transmitted.
SR04	RNRR. This bit is set when an RNR frame is received. Once set, it is cleared when an RR, REJ, or UA is received.
SR07-SR05	NA2-NA0. Next block of transmitted data to be Acknowledged.

STATUS REGISTER 1

REGISTER	SR17	SR16	SR15	SR14	SR13	SR12	SR11	SR10
SR1	↑PKR	↑XBA	↑ERROR	0	NE2	NE1	NE0	0

BIT	DESCRIPTION
SR10	0
SR13-SR11	NE2-NE0. Next Expected packet segment number of RLOOK.
SR14	0
SR15↑	The ERROR bit indicates: 1) An error has occurred which is not recoverable by the WD2511, or 2) A significant event has occurred. The "significant events" are: change in link status (link-up or down), the WD2511 is progressing to the next segment in a chained receive buffer, or one direction of the link has been reset. The exact nature of the reason for the ERROR bit is given in ER0.
SR16↑	The XBA bit means that a previously transmitted Block, or Blocks, have been acknowledged by the remote device. Upon acknowledgement, the ACK'ED bit is set to "1" for each segment in TLOOK which was acknowledged.
SR17↑	The PKR bit stands for Packet Received. This means that a packet has been received error-free and in correct sequence according to the received N (S) count. The data (I-field) has been placed in the CPU's RAM memory. NE is advanced.

NOTE 1: The three interrupt-causing bits are SR17, SR16, and SR15. Any of the three will cause an interrupt request ($\overline{INTR}$ goes low). After SR1 is read, all three bits are reset to 0, and $\overline{INTR}$ returns high.

STATUS REGISTER 2

REGISTER	SR27	SR26	SR25	SR24	SR23	SR22	SR21	SR20
SR2	T1OUT	$\overline{IRTS}$	REC IDLE	0	0	0	0	$\overline{LINK}$

BIT	DESCRIPTION
SR20	If the link is established, $\overline{LINK}$ = 0. If the link is logically disconnected, $\overline{LINK}$ = 1.
SR24-21	Unused Bits — 0
SR25	REC IDLE indicates that the WD2511 has received at least 15 contiguous 1's.
SR26	$\overline{IRTS}$ stands for the Internal Request-To-Send bit, and indicates that the transmitter is attempting (successful or not) to send either data or flags. If the $\overline{RTS}$ pin is not tied to ground or WIRE-ORED with another signal, then $\overline{IRTS}$ = RTS.
SR27	T1OUT bit means that timer T1 has timed-out. This bit returns to 0 when T1 is re-started. When T1OUT = 1, T1 is not running. NOTE: This bit could be a 1 for a few micro-seconds in between intervals when T1 stops and is restarted.

ERROR REGISTER (ER0)

ER07	ER06	ER05	
0	0	0	ER01 = ROR ER02 = TUR ER04 = RLNR
0	0	1	**ER04 ER03 ER02 ER01 ER00** 0 0 0 0 1 — LINK is up. (Was down) 1 0 0 0 0 — Received DISC or DM while LINK up (See Note 3) 0 0 1 0 0 — DISC sent, since sent N2 times without UA. (Note 3). 0 0 0 1 0 — DISC sent, REC IDLE for T1xN2.
0	1	0	**CHAIN STATUS** ER00 = GNCS ER01 = CNR
1	0		LINK RESET RECEIVED if ER05-ER00 = 000000 LINK RESET TRANSMITTED if ER05-ER00 = non-zero ER00 similar to W ER01 similar to X ER02 similar to Y ER03 similar to Z ER05 means received F = 1, but did not send P = 1 ER04 — I-frame was sent N2 times without acknowledge
1	1		FRAME REJECT (See Note 1) RECEIVED if ER05-ER00 = 000000 TRANSMITTED if ER05-ER00 = non-zero ER00 = W ER01 = X ER02 = Y ER03 = Z (See Note 2)

NOTES: 1. Whenever a frame reject (FRMR) is received, the I-field will have been placed in appropriate memory by DMA, and a link reset SABM will be transmitted. The NB is not advanced.
2. Definitions of W, X, Y, Z as stated in CCITT X.25.
3. DISC transmitted if DM received while link up.

Assigned bits are always 0.

TERMS USED IN ERROR REGISTER

GNCS	Going to Next Chain Segment.
RLNR	RLOOK Not Ready. REC RDY bit of next segment is 0, but RECR = 1. This interrupt will not occur if RECR = 0.
ROR*	Receiver Over-Run. The Receiver Register (RR) had a character to load into the FIFO, but the FIFO was full.
TUR*	Transmitter Under-Run. The Transmitter Register (TR) needed a character from the Transmitter Holding Register (THR), but the THR was not ready. The frame being transmitted is ABORTed.
CNR	Next Chain Segment of Receiver was Not Ready.

*TUR and ROR means that the bit rate clocks (TC and RC) are either too fast for the WD2511, or the DACK response to DRQI and DRQO is too slow, or both.

W, X, Y, Z OF FRMR

A frame reject (FRMR) contains a three byte I-field. The first byte is the rejected frame control field. The second byte contains the current N(S) and N(R) counts of the terminal reporting the rejection condition. The third byte contains W-X-Y-Z-Z1-0-0-0 where W is the LSB.

W set to 1 indicates that the control field received and returned in the first I-frame byte was invalid.

X set to 1 indicates the rejected frame contained an I-field which is not permitted with this command. X = 1 in conjunction with W = 1.

Y set to 1 indicates the I-field exceeded the total limit established. Y is mutually exclusive with W.

Z set to 1 indicates the control field contained an invalid N(R). Z is mutually exclusive with W.

Upon receiving a FRMR, the WD2511 will place the 3 byte I-field in memory by DMA, just as if the FRMR were a packet.

When the WD2511 transmits a FRMR, the frame reject condition is entered. Only a received SABM or DISC will clear this condition. If any other command is received, the WD2511 will re-transmit the FRMR. However, valid responses received will be acted upon accordingly. Also, the WD2511 may **not** transmit packets while in the frame reject condition.

In the FRMR I-field, bit #4 of the second byte is a "1" if the rejected frame was a response, and a "0" if the frame was a command.

MEMORY ACCESS METHOD

The WD2511 memory access is accomplished by the use of DMA and two look-up tables. These tables are set-up to allow up to 7 I-frames to be outstanding in each direction of the communications link. The look-up tables are divided into two areas (T-LOOK and R-LOOK) and are in memory external to the WD2511.

TLOOK
RLOOK

These tables contain address and control information for individual send/receive packets.

To access TLOOK and RLOOK only the starting address of TLOOK is loaded into the WD2511 in registers A' and B'.

REG A	A15	A14	A13	A12	A11	A10	A9	A8
REG B	A7	A6	A5	A4	A3	A2	A1	A0

A0-A15 16 bit TLOOK starting address

The TLOOK and RLOOK tables are each divided into 8 segments and each segment contains 8 bytes. Figure 3 illustrates the segmentation of TLOOK and RLOOK. Figure 5 and 6 illustrate the contents of a single TLOOK and RLOOK segment.

In transmit, the WD2511 will have read from TLOOK the starting address and length of the first packet to be transmitted. The WD2511 will automatically transmit the flag, address, and control fields. Next, the information field data will be transmitted using DMA from the "SEND #0 PACKET" memory. At the end of the information field, the WD2511 will automatically send the FCS and closing Flag. The WD2511 will then move on to the next packet.

If retransmission of one or more (up to seven) packets becomes necessary, the WD2511 will automatically retrace the previous transmissions through the TLOOK table. The user's CPU software does not become involved in the retransmission. However, an ERROR COUNTER is incremented. (See Error Counter Section.)

To receive, each frame is checked for correct address and FCS fields and for type of control field. If the frame is a packet, the information field is placed in the assigned memory location in a method similar to that used in transmit. After the packet is received error-free and in proper N(S) sequence count, an interrupt is generated, and the WD2511 is ready for the next packet which will be placed in the next location.

Figure 4 shows a "store-and-forward" example that is useful in a network node.

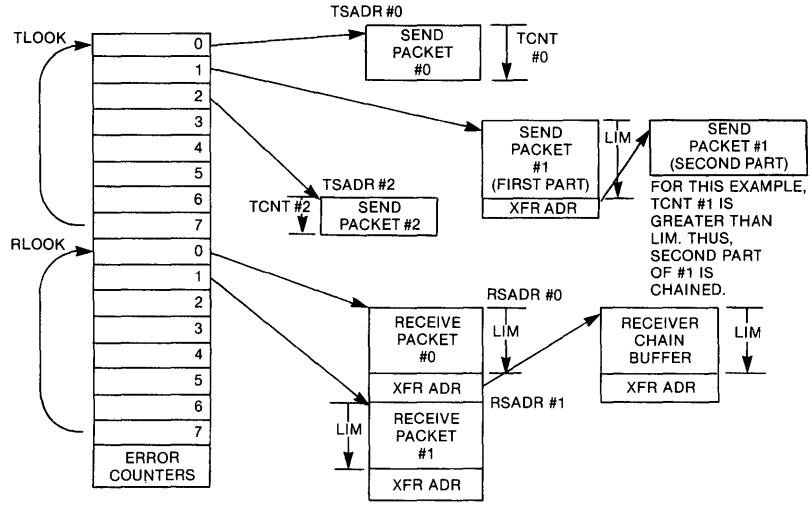

Figure 3. MEMORY ACCESS SCHEME

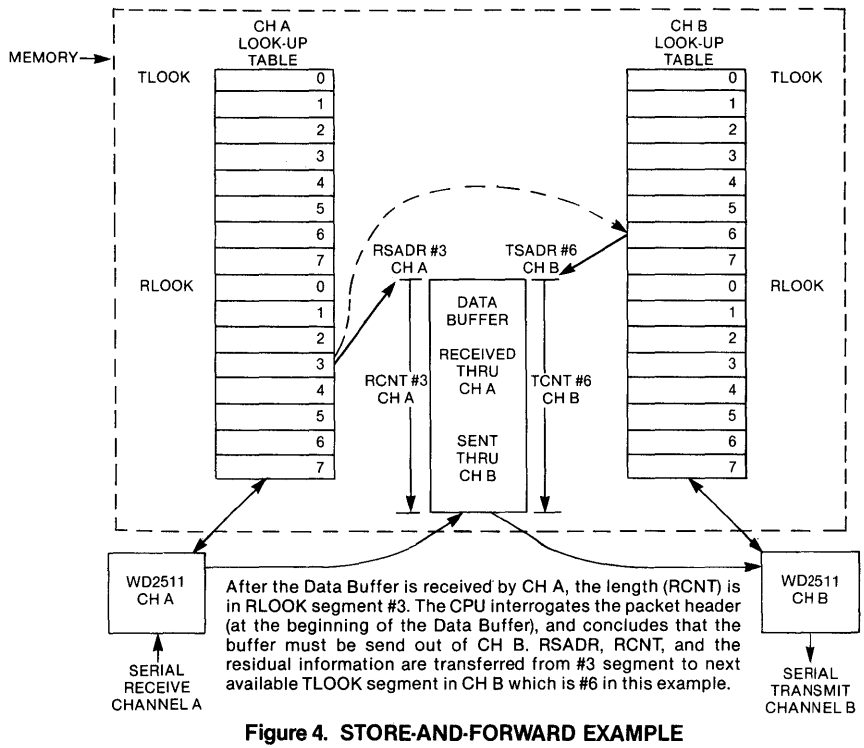

Figure 4. STORE-AND-FORWARD EXAMPLE

TLOOK AND RLOOK

Figures 5 and 6 detail the individual segments for TLOOK and RLOOK.

BRDY means that the transmit buffer is ready. The WD2511 will send the block only after the CPU makes BRDY = 1. (BRDY is used in conjunction with the SEND bit.) At the completion of the transmission, the WD2511 will make BRDY = 0, and then read the BRDY of the next segment.

After transmitting a packet, an acknowledgement must be received from the remote device. The acknowledgement is contained in the received N(R) count of an I-frame or S-frame. Upon acknowledgement, the WD2511 will make ACK'ED = 1, and generate a block-acknowledged interrupt. Before assigning a new block to a segment in TLOOK, the CPU must make sure that the previous block which used that segment number has been acknowledged.

REC RDY informs the WD2511 that the receive buffer is ready. The WD2511 will not receive a packet into a buffer referenced by a particular segment until REC RDY = 1. If the WD2511 progresses to a segment which has REC RDY = 0, an error interrupt will be generated.

After receiving an error-free packet with correct N(S), the WD2511 will, in order: 1) Set FRCML (Frame Complete), clear REC RDY, and store received residual count. 2) Store the received length, in characters, of the I-field in RCNT HI and RCNT LO. 3) Advance the NE count, and generate a packet received interrupt. 4) Acknowledge the received packet at the first opportunity.

The addresses (TSADR and RSADR) are 16-bit binary addresses. HI represents the upper 8-bits and LO represents the lower 8-bits. The counts (TCNT and RCNT) are 12-bit binary numbers for the number of characters in the I-field.

TSADR is the starting address of the buffer to transmit, and TCNT is the binary count of the number of bytes to transmit.

RSADR is the starting address of the receive buffer. After successfully receiving the packets, the WD2511 will write the value of RCNT which is the binary length of the received packet.

Whether the WD2511 accesses a look-up table or a memory block, a DMA Cycle is required for each access.

Figure 5. TLOOK SEGMENT

BYTE # IN SEGMENT	BIT # 7	6	5	4	3	2	1	0
1	ACK'ED	NS	NS	NS	NS	NS	NS	BRDY
2	TSADR HI							
3	TSADR LO							
4	SPARE				TCNT HI			
5	TCNT LO							
6	SPARE FOR USER DEFINITION							
7	SPARE							
8	SPARE							

NS = Not Spare

Figure 6. RLOOK SEGMENT

BYTE # IN SEGMENT	BIT # 7	6	5	4	3	2	1	0
1	FRCML*	NS	NS	NS	RES2	RES1	RES0	REC RDY
2	RSADR HI							
3	RSADR LO							
4	NOT SPARE				RCNT HI			
5	RCNT LO							
6	SPARE FOR USER DEFINITION							
7	SPARE							
8	SPARE							

NS = Not Spare (NOTE: The "not spare" bits may be either 1 or 0).

*FRCML = Frame Complete

The control bits in TLOOK (BRDY and ACK'ED) and in RLOOK (FRCML and REC RDY) define various states for each segment. These states are shown below:

TLOOK STATES

ACK'ED	BRDY	STATE
0	1	Ready To Transmit (CPU set BRDY, cleared ACK'ED)
0	0	*Transmitted and Awaiting Acknowledge (WD2511 cleared BRDY)
1	0	Received Acknowledge (WD2511 set ACK'ED)
1	1	This state not allowed

*State 0-0 could also occur whenever there is no data ready to send.

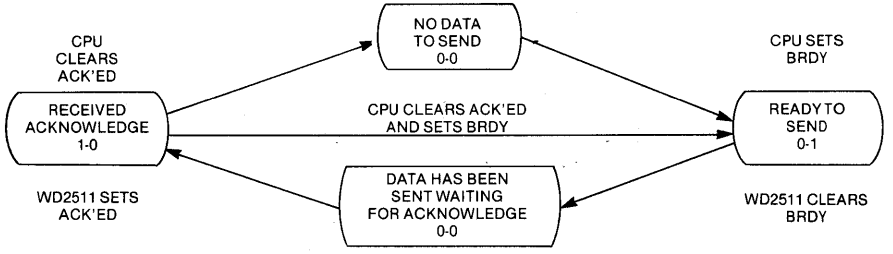

TLOOK SEGMENT STATE FLOW

Notice that in a TLOOK segment, the 0-0 state could have two meanings. Due to control internal to the WD2511, this will not pose an ambiguity to the WD2511. However, if it is a difficulty to the CPU, the CPU could do two things:

1. Avoid ever entering the "NO DATA TO SEND"
2. At start-up, the CPU should set all ACK'ED bits. (Since this would only be a start-up procedure, this would not violate the "deadly embrace" rule.)

As another alternative, the CPU could use one or more of the SPARE bits for additional state control.

In the "WAITING FOR ACKNOWLEDGE" state, one or more re-transmissions could occur.

REG = C	CHAIN				LIMIT			
CHAIN / LIMIT	Bit 7	Bit 6	Bit 5	Bit 4	Bit 3	Bit 2	Bit 1	Bit 0

CHAINING

The WD2511 includes a chained-block feature which allows the user more efficient use of memory particularly in situations where the maximum packet size is much larger than the average packet size.

Register C is used to program the chaining feature. The upper 4 bits define CHAIN which is the number of chain segments allowed in addition to the first segment. (If this feature is not used, make CHAIN all 0's).

The lower 4 bits of Register C define limit which is the size of the buffer in multiples of 64 bytes including the transfer address (XFR ADR). If LIMIT is 0000, the size is 64. For 0001, the size is 128, and so on.

For example, suppose that the LIMIT defines a segment size of 128 and that CHAIN defines 8 additional segments in addition to the first. (Note: Register C would be hex 81 in this example.) When 126 bytes of I-field have been received, the WD2511 will take the next two bytes as a transfer address (XFR ADR) pointing to another segment. At the end of that segment is another XFR ADR, and so on, up to a maximum of 9 total segments.

On the receiver, a XFR ADR of all 0's will mean that the next segment is not ready. If the WD2511 reaches a XFR ADR on the receiver, there will be an Error Interrupt code 42 if XRF ADR = all 0's. Otherwise, there will be an Interrupt code 41 which is a status indication that the WD2511 is going to the next segment.

On the receiver, Register 6 upper 4 bits gives a status of which chain segment is currently being used.

The transmitter chaining works just like the receiver with the following exceptions:

1. XFR ADR = all 0's will not indicate next segment not ready.
2. There is no interrupt when going from one segment to another.
3. There is no status of the current segment being used.

The total receiver limit is given by the expression:

$$\text{TOTAL RECEIVER LIMIT} = (64 \times (1 + \text{LIMIT}) - 2) \times (1 + \text{CHAIN})$$

Also, note that the transmitter and receiver counts are modified by 2 for each time a chain boundary is crossed. For example, if LIMIT = 0001 (segment size = 128 including XFR ADR), and if an I-field of 270 bytes is to be transmitted, then there will be two times that a chain boundary is crossed. The TCNT must be made 274 to send 270 bytes. The same is true for RCNT.

RLOOK STATES

FRCML	REC RDY	STATE
0	1	Ready To Receive (CPU set REC RDY, cleared FRCML)
1	0	Received Packet (WD2511 set FRCML, cleared REC RDY)
0	0	Not Ready (CPU cleared FRCML)
1	1	This state not allowed

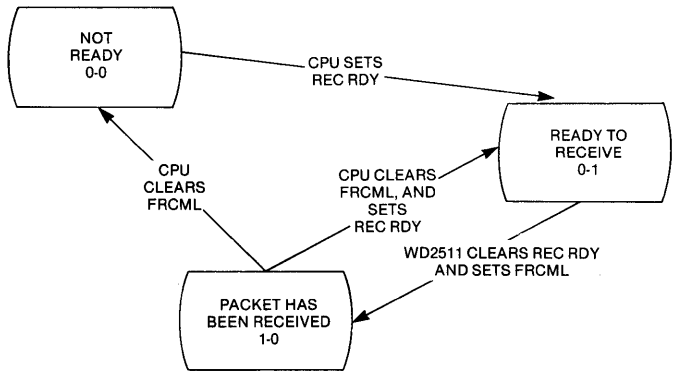

RLOOK SEGMENT STATE FLOW

"DEADLY EMBRACE" PREVENTION

A "deadly embrace" can occur when two processors reach a state where each is waiting for the other. In this case, the two processors are the user's CPU and the micro-controller inside the WD2511. Therefore, to prevent the "deadly embrace," the following rule is obeyed by the WD2511 and should also be obeyed by the user's CPU. This rule applies to TLOOK, RLOOK and to the I/O registers. The Error Counters do not apply to this rule.

> RULE: If a bit is set by the CPU, it will not be set by the WD2511, and vice versa. If a bit is cleared by the WD2511, it will not be cleared by the CPU, and vice versa.

As an example, the BRDY bit in the TLOOK segments is set by the CPU, only, but cleared by the WD2511, only.

TLOOK AND RLOOK POINTERS

There are three 3-bit counters for the status of the segments in TLOOK and RLOOK. Status Register #0 (SR0) contains counters NA and NB which are used in conjunction with TLOOK. NB is the segment number of the next block to be transmitted, and is advanced at the end of each block transmission. NA is the value of the segment of the next block which will be acknowledged. If all transmitted blocks have been acknowledged, then NA = NB.

In SR1 is a 3-bit counter, NE, used with RLOOK. NE is the value of the segment number where the next received packet will be placed.

NA = Next to be Acknowledged

NB = Next Block to be Transmitted

NE = Next Expected to be Received

VARIABLE BIT LENGTH AND RESIDUAL BITS

The WD2511 will send 8 bits per character, only. Also all transmitted frames will be integer multiples of bytes.

The WD2511 may receive a packet with, or without, an integer multiple of bytes. The "RES" bits in the RLOOK tables shows the number of residual bits. The residual bits occupy the lower portion of the last received character.

RES 2	RES 1	RES 0	Received Residual Bits
0	0	0	0
0	0	1	7
0	1	0	6
0	1	1	5
1	0	0	4
1	0	1	3
1	1	0	2
1	1	1	1

ERROR COUNTERS

Following contiguously after RLOOK are six 8-bit error counters. The WD2511 will increment each counter at the occurrence of the defined event. However, the WD2511 will not increment past 255 (all 1's). The CPU has the responsibility of clearing each counter. The first counter past RLOOK is #1, etc.

ERROR COUNTER	COUNT
1	*Received Frames with FCS Error (includes frames ABORTed in the I-field).
2	Received Short Frames (less than 32-bits)
3	**Number of times T1 ran-out (completed)
4	Number of I-Frame Retransmissions due to T1 completing
5	*REJ Frames Received
6	REJ Frames Transmitted

*These counters are incremented only if the received A-field is equal to either Register E or F.

**Incremented only when attempting to transmit a command.

The Error Counters are accessed by the WD2511 transmitter DMA channel. Therefore, if multiple errors are received while the WD2511 is transmitting a long frame, only the last error will be counted. The only Counters which could miss counts because of this are Counters #1, #2, and #5. The error Counters are incremented only when the link is up (LINK = 0).

OTHER I/O REGISTERS

RECEIVED C-FIELD

Register 7 is the C-field of the last received frame, provided the A-field of the frame was equal to either register E or F, the FCS was good, and the frame contained 32 or more bits, and the WD2511 is not waiting for a SARM or DISC in response to a transmitted CMDR.

TIMER

Registers 8 and 9 define a 10-bit timer (T1), and a 6-bit Maximum Number of Transmissions and Retransmissions counter (N2).

REGISTER	7	6	5	4	3	2	1	0
8	T1							LSB
9	N2							
	MSB					LSB	MSB	

MSB = Most Significant Bit
LSB = Least Significant Bit

T1 provides the value of a delay in waiting for a response and/or acknowledgement. The delay is the binary count of clock CT where:

$$CT = \frac{16{,}384}{CLK}$$

Thus, if CLK = 1 MHz, then T1 may be set in increments 16.384 milliseconds, to a maximum delay of 16.78 seconds. All ones in T1 is maximum delay.

Once the CPU establishes T1 and N2, there is no need to write into T1 and N2 again unless a master reset (MR) has occurred, there is a power loss, or the CPU needs to change T1 or N2. If a time-out occurs, the WD2511 will still retain T1 and N2.

The conditions for starting, stopping, or restarting T1 are shown below: ("Re-start" means starting T1 before it ran-out).

START T1	RE-START T1	STOP T1
1. *I-frame sent if T1 not already in progress due to previous I-frame	*Acknowledgement received to some, but not all, I-frames.	Acknowledgement received for all I-frames.
2. —	*RNR received while link up.	
3. SARM or DISC sent. (N2 restarted at first occurrence)	—	UA or DM Received
4. Receiver Idle (REC IDLE = 1)		Detect REC IDLE = 0
5. S — command sent		—

*N2 is restarted where shown by asterisk (*).

NOTE: Reason 4 and 5 (above) are overridden by reasons 1, 2 and 3. On reason 2, T1 is stopped if the received RNR acknowledged all outstanding packets.

"A" FIELD REGISTERS

Registers E and F provide a programmable A-field. This allows the WD2511 to be a super-set of the X.25 document. That is, the WD2511 can handle a wider range of application than the DTE-DCE links defined in X.25. These wider ranges include: DTE-to-DTE connection, multipoint, and loop-back testing.

If the WD2511 is strictly in an X.25 DTE-DCE link, use the values shown below:

DTE Register E = 01
Register F = 03

DCE Register E = 03
Register F = 01

If performing a loop-back test either internal (CR03 = 1) or external, registers E and F should be the same.

V. LAPB PROCEDURE

The Link Access Procedure Balanced (LAPB) is described in CCITT Recommendation X.25 as the Level 2 protocol for the Asynchronous Balanced Mode (ABM).

Zero bit insertion/deletion, use of flags and FCS are part of Level 2, and have been discussed in this document.

The DTE is the Data Terminal Equipment and the DCE is the Data Circuit Termination Equipment, and is the network side of the DTE-DCE connection.

The DTE and DCE are each "combined" stations in that each can transmit and receive commands and responses. Commands and responses are defined by the A-field. Commands from DCE to DTE and responses from DTE to DCE are address A (hex 03). Commands from DTE to DCE and responses from DCE to DTE are address B (hex 01).

The individual commands and responses are shown in Figure 7.

One use of the Poll bit (P) is in conjunction with Time-Out Recovery. Timer T1 is started at the beginning of a transmitted command provided it has not been previously started. If T1 runs out (completes), the command will be retransmitted with P = 1. If T1 runs out again, the command will again be retransmitted, with P = 1 up to N2 times. At N2 + 1, an error interrupt will occur. If the command was an I-frame, the WD2511 will reset the link by transmitting a SABM. If a SABM (either for link set-up or link reset), the WD2511 will send a DISC. If a DISC, the WD2511 will continue to send a DISC.

Figure 7

LAPB Commands and Responses (Bit 0 is transmitted first)
Only the CMDR and I-frame contain I-fields

FRAME TYPE	COMMAND	RESPONSE	BIT #							
			7	6	5	4	3	2	1	0
INFORMATION (I)	I-FRAME OR PACKET		N(R)			P	N(S)			0
UNNUMBERED (U)	SABM		0	0	1	P	1	1	1	1
	DISC		0	1	0	P	0	0	1	1
		UA	0	1	1	F	0	0	1	1
		FRMR	1	0	0	F	0	1	1	1
		DM	0	0	0	F	1	1	1	1
SUPERVISORY (S)	RR	RR	N(R)			P/F	0	0	0	1
	RNR	RNR	N(R)			P/F	0	1	0	1
	*REJ	REJ	N(R)			P/F	1	0	0	1

*The WD2511 will not send a REJ command (will send REJ response, only), but may receive either a REJ command or REJ response.

TRANSMISSION OF ABORT

An ABORT (a 0 followed by at least seven 1's) is transmitted to terminate a frame in such a manner that the receiving station will ignore the frame. There are three conditions which will cause the WD2511 to transmit an ABORT:

1. An ABORT is sent when there is a Transmitter Under-Run
2. While transmitting a packet, if a REJ S-frame is received, the packet is ABORTed.
3. If T1 times-out while a packet is being transmitted, the packet is ABORTed. Caution: If a packet is longer (in time) than T1, the packet will always be ABORTed.

SELF-TESTS

There are two self-test features: 1) Internal RAM Register Test and 2) Loop-Back Test. Both tests are suitable for manufacturing testing, user incoming inspection testing, or system diagnostics and trouble-shooting.

INTERNAL RAM REGISTER TEST

There are eleven 8-bit registers internal to the WD2511 which are not directly accessible by the user's CPU. Seven of these registers can be tested by the Loop-Back Test. This test provides a means to check the other four registers.

The contents of Register A are placed in two even internal registers, and the contents of Register B in two odd internal registers. The four registers are then added together without carry and the result is placed in Registers 2, 5, 6 and 7. This test is initiated when RAMT (CR02) = 1. Use the following procedure:

1. Set-up Registers A and B.
2. Set RAMT.
3. Wait at least 50 times the CLK Period.
4. Read Registers 2, 5, 6 and 7.

To repeat the test for new values in Registers A and B:

5. Clear RAMT.
6. Wait at least 100 times the CLK period.
7. Go back to step 1.

LOOP-BACK TEST

The loop-back may be internal (CR03 = 1) or external (CR03 = 0). Of course, if external, RD and TD must be tied together either directly or remotely.

If CR03 = 1, TD is internally tied to RD, and the pin at RD (16) is internally disconnected. Also, $\overline{TC}$ is internally tied to $\overline{RC}$, and the pin at $\overline{RC}$ (17) is internally disconnected.

WD2511 ELECTRICAL SPECIFICATIONS:

ABSOLUTE MAXIMUM RATINGS:

Voltages referenced to V_{SS}

High Supply Voltage (V_{DD}) . . . . . −0.3 to 15V
Voltage at any Pin . . . . . −.03 to 15V
Operating Temperature Range . . . . . 0°C to +70°C
Storage Temperature Range . . . . −55°C to +125°C

NOTE:
Maximum limits indicate where permanent device damage occurs. Continuous operation at these limits is not intended and should be limited to those conditions specified in the DC Electrical characteristics.

Operating DC Characteristics: V_{SS} = 0V, V_{CC} = 5.0V ± .25, V_{SS} = 12.0V ± .6V T_A = 0° to 70°C

SYMBOL	PARAMETER	MIN	TYP	MAX	UNIT	CONDITIONS
I_{DD}	V_{DD} Supply Current		20	70	mA	
I_{CC}	V_{CC} Supply Current		200	280	mA	
V_{DD}	High Voltage Supply	11.4	12	12.6	V	
V_{CC}	Low Voltage Supply	4.75	5	5.25	V	
V_{IH}	Input High Voltage	2.4			V	
V_{IL}	Input Low Voltage			0.8	V	
V_{OH}	Output High Voltage	2.8			V	I_O = −0.1mA
V_{OL}	Output Low Voltage			0.4	V	I_O = 1.6mA
I_{LH}	Input Leakage Source or Sink			10	uA	
I_{OZH}	Input Leakage High Impedance			10	uA	Vin = V_{CC}
I_{OZL}	Output Leakage High Impedance			10	uA	Vin = V_{SS}

AC Timing Characteristics: V_{CC} = 5V ± .25; V_{DD} = RV ± .6V, V_{SS} = 0V, T_A = 0° to 70°C

SYMBOL	PARAMETER	MIN	TYP	MAX	UNIT	CONDITIONS
CLK	Clock Frequency	0.5		2.0	MHz	Note 1, 2
RC	Receive Clock Range	0			MHz	Note 3, 4
TC	Transmit Clock Range	0			MHz	Note 4
MR	Master Reset Pulse Width	10			mS	
T_{AR}	Input Address Valid to $\overline{RE}$	0			nS	
T_{RD}	Read Strobe (or $\overline{DACK}$ Read) to Data Valid			375	ns	Note 5
T_{HD}	Data Hold Time From Read to Strobe			80	nS	
T_{HA}	Address Hold Time From Read Strobe	80			ns	
T_{AW}	Input Address Valid to Trailing Edge of $\overline{WE}$	200			nS	
T_{WW}	Minimum $\overline{WE}$ Pulse Width	200			nS	
T_{DW}	Data Valid to Trailing Edge of $\overline{WE}$ or Trailing Edge of $\overline{DACK}$ for DMA Write	100			nS	
T_{AHW}	Address Hold Time After $\overline{WE}$	80			nS	
T_{DHW}	Data Hold Time After $\overline{WE}$ or After $\overline{DACK}$ for DMA Write	100			nS	
T_{DA1}	Time From $\overline{DRQO}$ (or $\overline{DRQI}$) to Output Address Valid if ADRV = 1			80	nS	
T_{DA0}	Time From $\overline{DACK}$ to Output Address Valid if ADRV = 0			400	nS	Note 5
T_{DD}	Time From Leading Edge of $\overline{DACK}$ to Trailing Edge of $\overline{DRQO}$ (or $\overline{DRQI}$)			400	nS	Note 5
T_{DAH}	Output Address Hold Time From $\overline{DACK}$			100	nS	
T_{DMW}	Data Hold Time From $\overline{DACK}$ For DMA Read			100	nS	
T_{RP1}	$\overline{REPLY}$ Response Time (leading edge)			240	nS	Note 5
T_{RP2}	$\overline{REPLY}$ Response Time (trailing edge)			260	nS	Note 5

NOTES:
1. Clock must have 50% duty cycle.
2. Buffer chaining is not guaranteed when CLK is greater than 1.5 MHz.
3. Residual bit detection logic not guaranteed when $\overline{RC}$ is greater than 50Kbps.
4. See "Ordering Information" for maximum serial rates.
5. C(load) = 100pf

DMA TIMING

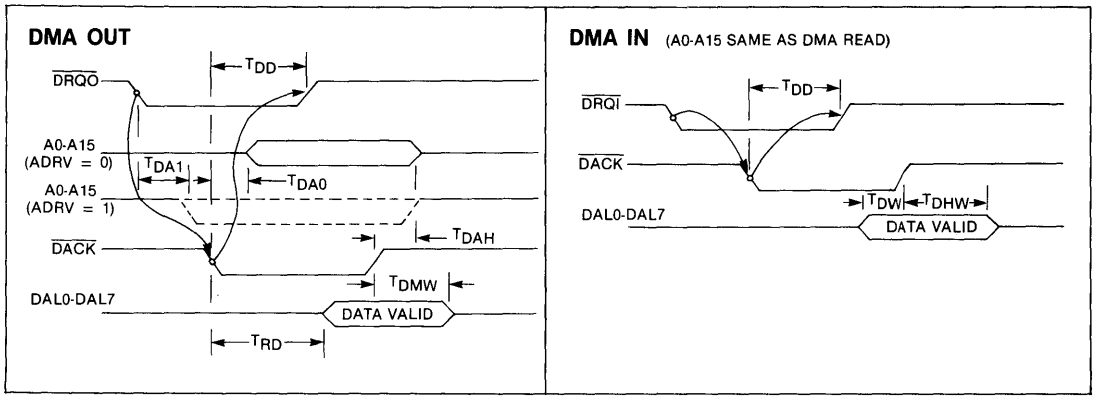

CPU READ/WRITE TIMING

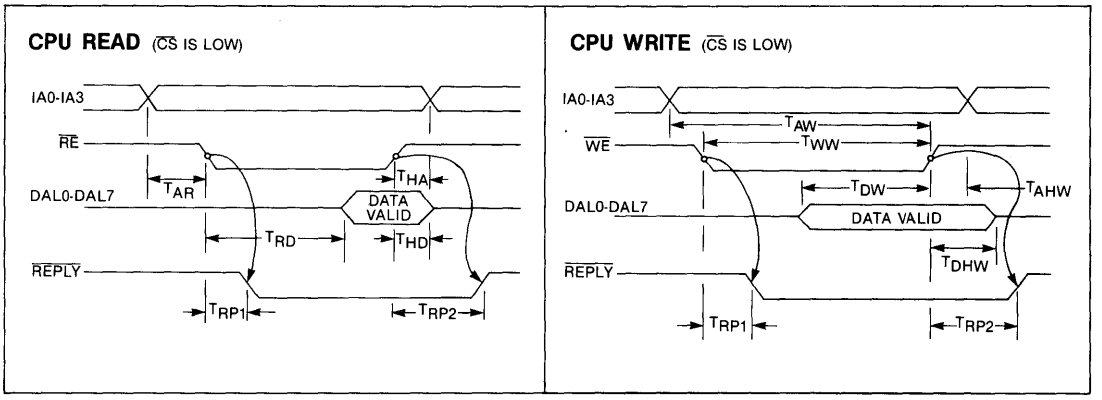

***NOTE:** There must not be a CPU read or write ($\overline{CS}$-$\overline{RE}$ or $\overline{CS}$-$\overline{WE}$) within 500 nanoseconds after the trailing (rising) edge of $\overline{DACK}$.

There must not be the leading (falling) edge of $\overline{DACK}$ allowed within 500 nanoseconds after the completion of a CPU write ($\overline{CS}$-$\overline{WE}$).

$\overline{RE}$ and $\overline{WE}$ must not be low at the same time.

APPENDIX A

APPENDIX B

APPENDIX C

APPENDIX A

TRANSPARENT MODES

The WD2511 was originally intended to be a link level controller meeting the requirements of X.25 LAPB, and this has been accomplished. However, there has been an increasing demand from potential WD2511 users for additional frame types not included in the LAPB frame type repertoire.

For example, the Bell System standard, BX.25, calls for the use of XID (exchange identification) in LAPB connections of DTE-to-DTE. (Of course, DTE-to-DTE is not X.25 in the strictest sense.) Also, Western Digital has received several requests for the use of a SIM (set initialization mode). Also, there has been one request to allow "unknown" frames to pass thru the chip for the purpose of teleloading.

Therefore, to the WD2511 we have added two selectable modes: transparent transmit and transparent receive. Basically, these two modes allow the user the option to pass non-LAPB frames thru the chip without controlling these frames according to the LAPB protocol.

FEATURES OF THE TRANSPARENT MODES

- May transmit any A and C field under transparent control.
- May receive any U-frame not part of the LAPB repertoire if transparent-receive enabled.
- Transparent modes are link state independent.
- Software to control WD2511 is backward compatible with the WD2501 (LAP chip).

1.0 HOW THE TRANSPARENT MODES WORK

Two control bits have been added. TXMT (CR17) is the bit to enable the Transparent Transmit, and TRCV (CR16) will enable the Transparent Receive.

1.1 TRANSPARENT TRANSMIT

When TXMT = 1, the WD2511 will transmit the frame in the next TLOOK segment provided SEND (CR10) = 1, and BRDY of that TLOOK segment is 1. The link may be either UP or DOWN. The WD2511 will not add the A and C fields to the Transparent Transmitted frame. The user's CPU must add these fields as the first two bytes in the transmit buffer. Thus, the significance of the transmit count (TCNT) is different from packet transmission. In packet transmission, TCNT is the count of the I-field. In transparent transmission, TCNT is the I-field plus the A and C fields (I-field plus two bytes).

The timer, T1, will be disabled in transparent transmission. Therefore, if using this feature while the link is UP, it is advised that TXMT be set only when there are no outstanding (unacknowledged) packets which is indicated whenever NA = NB.

At the end of the transparent transmission, there will be an interrupt with XBA = 1. The SEND bit will be cleared, but the BRDY bit will not be cleared. The NB pointer will not be incremented. To send another transparent frame, set SEND. To resume packet transmission, clear TXMT and set SEND. (Of course, the TLOOK segment must be set-up prior to setting SEND.)

If SEND is set while the link is down, a transmission will occur even if TXMT = 0. Under this condition, a packet will be transmitted from current TLOOK segment, and NB and V(S) will be incremented, and the chip will go on to the next TLOOK segment just as if the link were UP. However, the WD2511 will expect no acknowledgment to the packet(s). If the link is brought UP later, NB and V(S) are cleared to 0 at the time the link comes UP.

The bit $\overline{XI}$ (CR15) is used only when TXMT = 1. $\overline{XI}$ stands for Transmit I-field. If the frame contains three, or more bytes, not counting FCS, make $\overline{XI} = 0$. If the frame contains two bytes not counting FCS, make $\overline{XI} = 1$. When $\overline{SI} = 1$, only two frame bytes will be transmitted regardless of TCNT. DO NOT attempt to transmit a frame with TXMT = 1 and $\overline{XI} = 0$ if TCNT is 2, 1, or 0.

1.2 TRANSPARENT RECEIVE

For the purposes of this discussion, it is necessary to define an "unknown frame." That is, a frame which is "unknown" to the WD2511.

Unknown Frame: A U-frame (unnumbered) frame which is not part of the LAPB repertoire. The U-frame repertoire in LAPB is SABM, DISC, DM, UA, and FRMR. For the purposes of this discussion, "UF" will refer to an unknown frame without an I-field, and "UFI" will refer to an unknown frame with a I-field.

A received SREJ (Selective REJect), which is an S-frame, is not considered an unknown frame by the WD2511. If the link is DOWN and an SREJ command is received, a DM response will be sent. If the link is DOWN and a SREJ response is received, the SREJ is disregarded. If the link is UP and a SREJ command or response is received, a FRMR will be sent with W = 1. The WD2511 will treat a received SREJ the same whether TRCV is 0 or 1.

A received packet (I-frame) response is not considered an unknown frame by the WD2511. If the link is DOWN, the frame is disregarded. If the link is UP, a FRMR will be sent with W = 1 and X = 1. The received packet response is treated the same whether TRCV is 0 or 1.

Whether TRCV is 0 or 1, the WD2511 will check all received frames to insure that the A-field equals either Register E or F, and that the FCS is correct, and that the frame contains 32 bits, or more. If TRCV = 0, and if a UF or UFI is received, and if the link is UP, the WD2511 will send a FRMR with W = 1 (W and X are 1 in the case of a UFI). See "States of the WD2511."

When TRCV = 1, the WD2511 will be enabled to receive all frames. If the frame is "known" by the

WD2511, it will be treated according to the protocol just as if TRCV = 0. However, if the frame is a UF or UFI, it will be passed on to the user's CPU.

When an unknown frame is received, and when TRCV = 1, there will be an interrupt with ERROR = 1, and the Error Register (ER0) will contain one of the following hexidecimal values:

ER0	FRAME RECEIVED
60	UFI Response
61	UFI Command
62	UF Response
63	UF Command

The C-field of the received frame is contained in Register #7. If the frame had an I-field, the frame will be placed in the next RLOOK segment, and the value of RCNT will represent the count of bytes in the I-field (not including the A and C fields). The RLOOK pointer, NE, will be incremented. Therefore, the relationship between NE and V(R) will not be guaranteed if transparent receive is used while the link is UP. However, this will not cause a sequence problem in the protocol since the actual V(R) is maintained in an internal register in the WD2511. Note that NE is cleared when the link is brought UP. Thus, if transparent receive is used only when the link is DOWN, then NE will be equal to V(R).

A word of caution. If the next RLOOK segment is not ready when a UFI is received, the Error Register (60 or 61) will be overwritten almost immediately with an error code 10 (RLNR), and the user will not know if the received UFI was a command or response.

If RECR is set while the link is DOWN, the WD2511 will prepare to receive I-fields, whether TRCV is 0 or 1. If a packet command is received, there will be a PKR interrupt, and the NE and V(R) will be incremented. Of course, NE and V(R) are cleared once the link is brought up.

The following tables show what action the WD2511 will take when various frames are received.

TABLE I. PACKET RECEIVED (command, not response)

LINK	RLOOK READY	TRCV	ACTION BY WD2511
DOWN	NO	0 or 1	DISREGARD
DOWN	YES	0 or 1	If N(S) = V(R), PKR interrupt. V(R) and NE incremented. No ack transmitted. If N(S) not = V(R), DISREGARD.
UP	NO	0 or 1	If N(S) = V(RT), RNR sent. Else, REJ condition entered.
UP	YES	0 or 1	If N(S) = V(R), PKR interrupt. V(R) and NE incremented. Acknowledgement sent at next opportunity. If N(S) not = V(R), enter REJ condition.

TABLE II. UFI RECEIVED

LINK	RLOOK READY	TRCV	ACTION BY WD2511
DOWN	NO	0 or 1	DISREGARD
DOWN	YES	0	DISREGARD
DOWN	YES	1	Error interrupt 60 or 61. NE incremented.
UP	NO	0	FRMR sent. W = 1 X = 1
UP	NO	1	DISREGARD
UP	YES	0	FRMR sent. W = 1 X = 1
UP	YES	1	Error interrupt 60 or 61. NE incremented.

If TRCV = 1 and UF (no I-field) is received, there will be an Error interrupt 62 or 63, independent of the link state or the readiness of RLOOK.

Of course, the received C-field of any frame will be in Register #7 provided the A-field matched either Register E or F, the FCS was good, and the frame contained 32, or more, bits.

APPENDIX B

HALF DUPLEX OPTION

The WD2511 is basically a full duplex device. The receiver is maintained in an "always ready" condition even if the receive buffer is not ready. Thus, whether the received frame came from a full or half duplex system is of no consequence to the WD2511.

Therefore, the half duplex option affects only the WD2511 transmitter. Half duplex is enabled when H/$\overline{F}$ (CR05) = 1.

The WD2511 will transmit one frame at a time according to the following procedure:

A. Enable $\overline{RTS}$ ($\overline{RTS}$ goes low).

B. Wait for $\overline{CTS}$ ($\overline{CTS}$ input goes low).

C. Transmit frame.

D. Remove $\overline{RTS}$ ($\overline{RTS}$ goes high 2½ bits of time after the last 0 of the trailing flag.)

NOTES:
The leading flag will be transmitted somewhere between 5 and 13 bits after $\overline{CTS}$ goes low.

RTS returns high after about 3 bits of the second trailing flag has been transmitted.

If T1 is started, it is started when $\overline{RTS}$ goes low.

After $\overline{RTS}$ goes low, the frame will not begin transmission until $\overline{CTS}$ goes low. After the frame has started, the transmission of that frame is completed even if $\overline{CTS}$ returns high during the frame.

APPENDIX C

STATE DESCRIPTIONS

The State Descriptions consists of flow charts, tables, and verbal definitions to show how the WD2511 will react to various external conditions. The State Descriptions contain the following topics:

Link Down Flow
Link Up Flow
State Table I. Link Down, but Going Up
State Table II. Link Going Down
Use of Flags by the WD2511
Definitions of Command and Response
State Table III. Sending I-Frames (Packets) and S-Commands
Receiving and Transmitting a Null Packet
Sending a REJ
Receiving a REJ
Definition of Valid Received N(R)
State Table IV. Sending an RNR (RNRX)
State Table V. Receiving an RNR (RNRR)
Sending S-Frame Commands
Conditions for Sending Link Reset
Unsolicited UA and Unsolicited F Bit
Sending an FRMR
Receiving an FRMR
Protocol Significance of TLOOK/RLOOK Pointers
Use of the RECR Bit

LINK DOWN

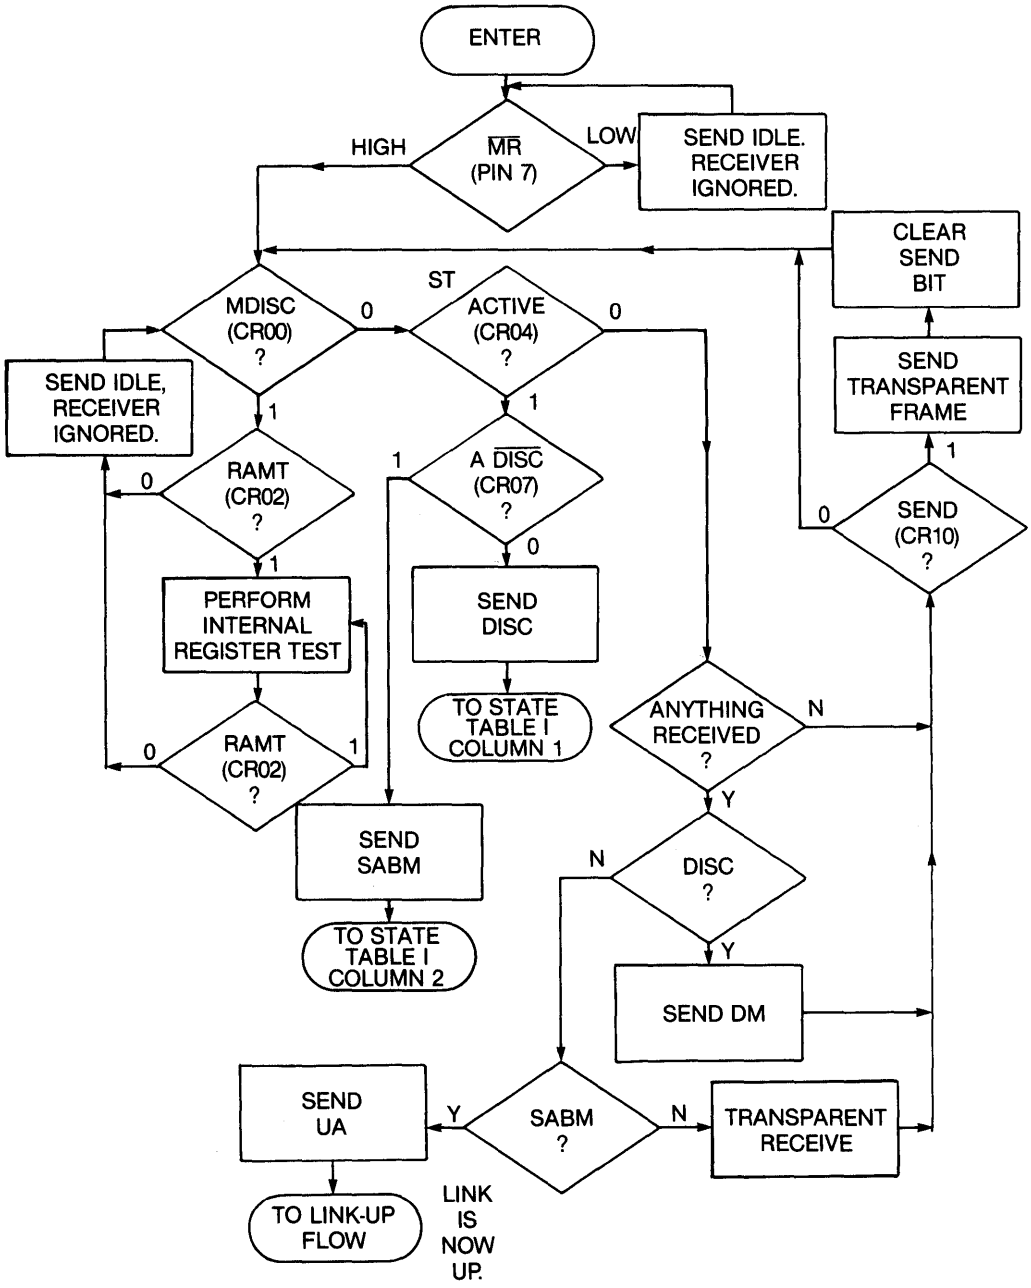
ENTER
HIGH
MR (PIN 7)
LOW
SEND IDLE. RECEIVER IGNORED.
CLEAR SEND BIT
ST
MDISC (CR00) ?
0
ACTIVE (CR04) ?
0
SEND IDLE, RECEIVER IGNORED.
1
1
SEND TRANSPARENT FRAME
1
RAMT (CR02) ?
0
1
A DISC (CR07) ?
0
SEND (CR10) ?
0
1
PERFORM INTERNAL REGISTER TEST
SEND DISC
TO STATE TABLE I COLUMN 1
ANYTHING RECEIVED ?
N
RAMT (CR02) ?
0
1
Y
SEND SABM
TO STATE TABLE I COLUMN 2
N
DISC ?
Y
SEND DM
SEND UA
Y
SABM ?
N
TRANSPARENT RECEIVE
TO LINK-UP FLOW
LINK IS NOW UP.

LINK UP
CONDITIONS FOR RESET OR DISCONNECT

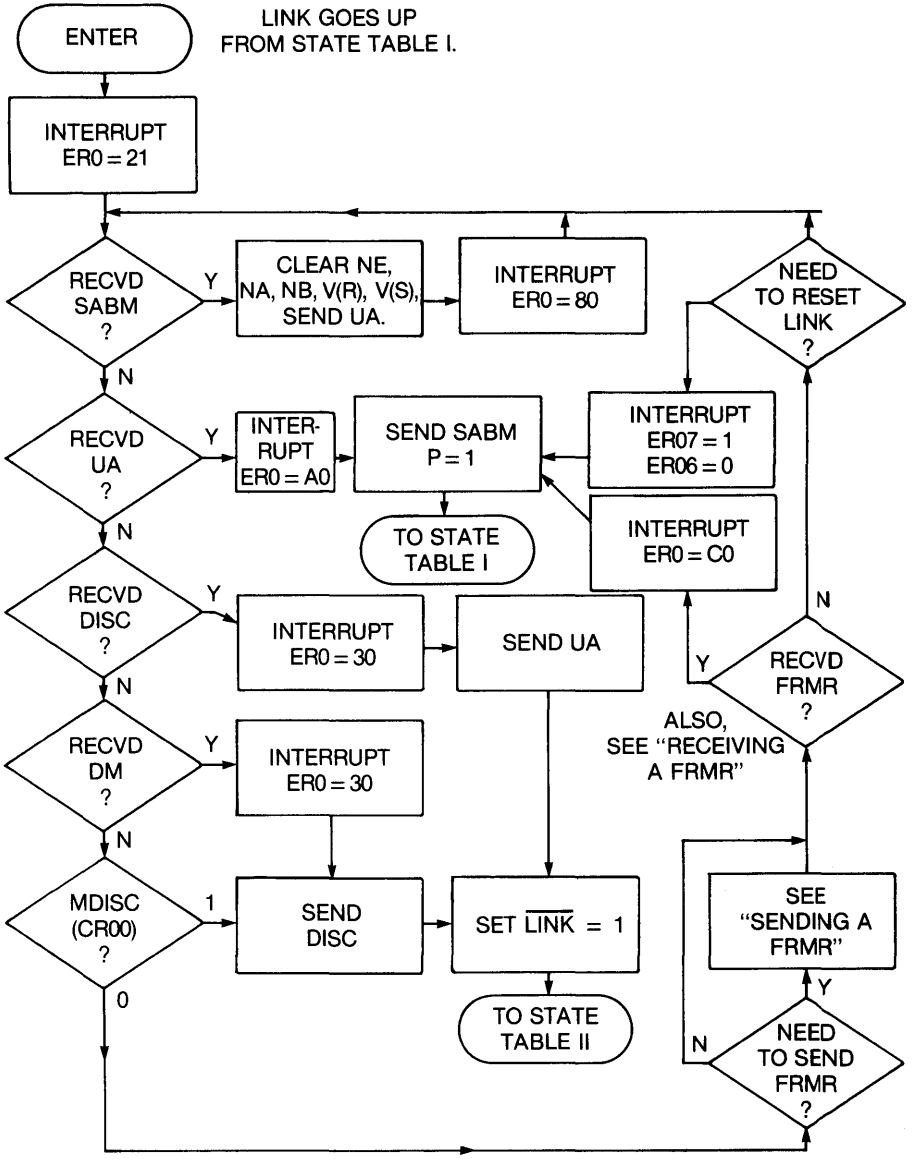

STATE TABLE I
LINK DOWN, BUT GOING UP
(Column 2 also applies to link reset)

ACTION BY WD2511

Stimulus:	Column 1: DISC sent. Waiting for UA or DM.	Column 2: SABM sent. Waiting for UA.
T1 runs out	Re-send DISC. P = 1.	Re-send SABM. P = 1.
T 1 and N2 run out	Re-send DISC. P = 1.	Send DISC. Interrupt ER0 = 24. Go to column 1
Received UA	Send SABM. Go to column 2.	Clear NA, NB, NE, V(R), V(S). Go to link up flow.
Received DISC	Send DM.	Send DM.
Received SABM	Send DM.	Send UA. Clear NA, NB, NE, V(R), V(S). Keep waiting for UA.
Received DM	Send SABM. Go to column 2.	Send DISC. Go to column 1.
Received something other than UA, DM, DISC, or SABM	Disregard.	Disregard.

STATE TABLE II
LINK GOING DOWN (WAS UP)

ACTION BY WD2511

Stimulus:	Column 1: DISC sent. Waiting for UA.	Column 2: MDISC (CR00) set.
T1 runs out	Re-send DISC. P = 1.	Send DISC. Go to column 1.
Received UA or DM	Go to Link Down Flow	
Received SABM	Disregard	
Received DISC	Send UA. Go to Link Down Flow	
Received something other than DISC, SABM, UA, or DM	Disregard.	

USE OF FLAGS BY THE WD2511

The WD2511 will send interframe flags whenever full duplex is selected (CR05 = 0), and point ST of the Link Down Flow point has been entered. In half duplex (CR05 = 1), interframe fill will be all 1's.

The WD2511 does not require the interframe time fill flags. Either idle (all 1's) or flags will be accepted. However, if the receiver detects idle for time T1 X N2, the WD2511 will send a DISC.

When sending continuous flags, the WD2511 will send:

0 1 1 1 1 1 1 0 0 1 1 1 1 1 1 0 0 1 1 1 1 1 1 1 0 0 1 1 . . .

The WD2511 will accept either the above sequence as continuous flags, or the "shared zero" pattern:

0 1 1 1 1 1 1 0 1 1 1 1 1 1 0 1 1 1 1 1 1 1 0 1 1 1 1 1 . . .

DEFINITIONS OF COMMAND AND RESPONSE

A transmitted or received command or response is a frame with the A-field defined below:

FRAME	A-FIELD =
Transmitted Command	Register E
Received Command	Register F
Transmitted Response	Register F
Received Response	Register E

For non-transparent transmitted frames, only commands or responses are transmitted. A transparent transmitted frame (TXMT = 1) may have any A-field the user chooses.

All received frames must be either commands or responses, or the frame is disregarded ("thrown away"), even if transparent receive is enabled (TRCV = 1).

STATE TABLE III
SENDING I-FRAMES (PACKETS) AND S-COMMANDS

NOTES:
In all subsequent pages, the link is considered Up (LINK = 0) unless otherwise stated. X = don't care. TXMT = 0 for Table III.

SEND	BRDY	NA AND NB	RNRR	T1 EXPIRES	RCVD REJ	ACTION BY WD2511
1	0	X	0	No	No	Clear SEND (CR10)
1	1	X	0	No	No	Send next packet with N(S) = NB. After transmission complete. Increment NB. Exception: If NB + 1 = NA, do not send next packet. There are 7 outstanding.
X	X	X	1	Yes	No	Send S-command, P = 1.
X	X	not =	X	Yes	No	Send S-command, P = 1.
X	X	not =	0	No	Yes	Make NA = received N(R). Start sequential retransmission of packets beginning with N(S) = NA. See Note 3.

NOTES ON STATE TABLE III

1. Received S-frames in Table III are assumed to have valid N(R)'s.
2. When an acknowledgement of one, or more, previously transmitted packets is received, NA is made equal to the received N(R). All TLOOK segments from the old value of NA up to N(R) − 1 are acknowledged, and the appropriate ACKED bits in the TLOOK segments will be set. After setting the ACKED bits, an XBA interrupt is generated.
3. Assuming appropriate TLOOK segments are ready, packets are transmitted sequentially without waiting for an acknowledgement with three exceptions:
   a. There are seven outstanding (unacknowledged) packets. (NB + 1 = NA)
   b. The remote station has indicated a busy condition, because an RNR was received. T1 is started, and an S-command will be transmitted with P = 1 when T1 expires.
   c. If T1 expires, and there are one, or more, outstanding packets, an S-command will be transmitted with P = 1.
4. If an S-frame command is received, the WD2511 will transmit an S-frame response at the next opportunity.
5. If SEND = 1 and TXMT = 1, a frame will be transmitted from the next TLOOK segment if BRDY = 1. After transmission, SEND is cleared by the WD2511.

RECEIVING AND TRANSMITTING A NULL PACKET

If an error-free (FCS good) packet is received with a correct N(S), but has no I-field, that packet will be treated the same as a packet with an I-field. The fact that there was no I-field is shown by RCNT equal to all 0's.

The WD2511 will not transmit a null packet. TCNT must not be allowed to be all 0's.

SENDING A REJ (RESPONSE)

1. The REJ condition is entered any time an error-free packet is received with an out-of-sequence N(S). Exception: If the received N(S) + 1 = V(R), then the received N(S) has been acknowledged, and either an RR or RNR is transmitted.
2. When the REJ condition is entered, the REJ frame with N(R) = V(R) is transmitted immediately if a packet is not being transmitted, or at the completion of the current packet. There are two exceptions, as noted in 3 and 4 below.
3. If a link resetting SABM needs to be transmitted, the SABM is sent ahead of the REJ. However, when the UA is received for the SABM, the REJ condition is cleared.
4. If the receiver is not ready (RNRX = 1), the REJ is not sent until RNRX is cleared.
5. Once the REJ condition is entered, only one REJ will be transmitted. Another REJ is not transmitted unless the REJ condition is cleared and re-entered. The REJ condition is cleared if a packet is received with correct N(S), or if a SABM is received, or if a SABM is transmitted and a UA received.
6. When the REJ is transmitted, error counter #6 is incremented.

RECEIVING A REJ (RESPONSE OR COMMAND)

Suppose a REJ has been received error-free with no I-field, then:

1. If the N(R) is not valid, an interrupt is generated with ER0 = C8, and a FRMR is transmitted.
2. If the N(R) is valid, and greater than NA, at least one transmitted packet is acknowledged. The appropriate ACKED bits in TLOOK are set, and an XBA interrupt is generated.
3. If the N(R) is valid and less than NB, the WD2511 will begin sequential retransmission starting with V(S) = received N(R). If a packet is being transmitted when the REJ was received, that packet is aborted. If the N(R) is valid and equal to NB, and a packet is being transmitted, that packet (which will be #NB) is aborted, and retransmission will begin.
4. If the N(R) is valid and equal to NB, and there is no packet being transmitted, there is no retransmission initiated. In this case the REJ has the same effect as an RR.
5. If in 2, 3, or 4 above, the received REJ is a command, the WD2511 will transmit a RR or RNR response at the next opportunity.

DEFINITION OF VALID RECEIVED N(R)

Reference

CCITT Recommendation X.25 paragraphs 2.4.10 and 2.3.4.10.

Definition

A valid received N(R) is greater than or equal to NA, and less than or equal to NB.

1. The "greater than" and "less than" relationships must be understood in a circular sense. 0 could be greater than 7 depending on the values of NA and NB.
2. If NA = NB, there is only one possible valid received N(R).
3. If NB + 1 = NA, there are seven outstanding packets, and any received N(R) will be valid.
4. Basically, a received N(R) which is not valid is one which acknowledges a packet, or packets, never transmitted.

STATE TABLE IV
SENDING AN RNR (RNRX) (RESPONSE OR COMMAND)

LINK	RECR	REC RDY	ACTION
1	X	X	No S-frame transmitted when link down.
1 → 0	1	1	RLOOK ready. No S-frame sent immediately.
1 → 0	1	0	Interrupt RLNR. RNR response sent immediately after link Up. RNRX set.
1 → 0	0	X	RNR response sent immediately after link Up. RNRX set.
0	1	1	Receiver ready to accept packets.
0	1 → 0	1	Receiver ready to accept packets.
0	1	0	Interrupt RLNR. RNR response sent. RNRX set.
0	0 → 1	1	If RNRX was set, then RNRX will be cleared after the next received packet or S-command. After that, an RR or REJ response is sent.
0	0	0	RNR response sent. RNRX set. There is no RLNR interrupt.

NOTES ON STATE TABLE IV

1. The arrows (→) indicate a change in state from the value on the left to the value on the right.
2. The RNRX status bit is set at the time the receiver-not-ready condition was established. The RNR frame will be sent immediately if no packet is being sent, or after the end of the current packet.
3. When a received packet is brought into memory with RNRX = 0, the packet will be accepted provided the FCS and N(S) are correct, and the I-field is not too long. The N(R) may, or may not, be correct, but is checked separately. If N(R) is not valid, a FRMR is transmitted.
4. Whenever RNRX = 1, the I-field of all received frames is not brought into memory. For received packets, the N(S) and N(R) are checked as usual. If the N(S) is out-of-sequence, the REJ will be transmitted after RNRX is cleared.
5. If a link resetting SABM is transmitted when RNRX = 1, RNRX will be cleared when the UA is received. If the condition which caused receiver-not-ready still exists, an RNR is sent and RNRX is set. However, if the receiver is now ready, I-field data may be brought into memory.

The same also applies is a link resetting SABM is received, and a UA transmitted when RNRX = 1.

STATE TABLE V
RECEIVING AN RNR (RNRR)

SEND	NA AND NB	RECVD ACK?	RECVD RNR	RECVD RR, REJ OR UA	T1 EXPIRES	ACTION
X	not =	Yes	Yes	No	No	Set RNRR. Restart T1 and N2. Update UA.
0	Equal	No	Yes	No	No	Set RNRR.
X	not =	No	No	No	Yes	Send S-command (P = 1). If RNR subsequently received restart T1 and N2.
X	not =	Yes	No	Yes, but not UA	No	Clear RNRR. Restart T1 and N2. Update NA.
X	X	X	No	Yes	No	Clear RNRR.
0 → 1	Equal	No	No	No	No	Send next packet. Increment NB after transmission. (Then, NB does not = NA). Start T1 and N2.

2. NOTES OF TABLE V

1. If SEND = 1, it is assumed for this table that BRDY of the next TLOOK segment is set.
2. If RNRR = 1, an RR or RNR command is transmitted at T1 intervals.

SENDING S-FRAME COMMANDS

When an S-frame command is to be transmitted, an RR command is transmitted if RNRX = 0, or an RNR command is transmitted if RNRX = 1. If RNRX = 0, and a REJ is waiting to be transmitted, a REJ command is transmitted.

For all transmitted S-commands, the P bit is set to 1.

An S-command will be transmitted at T1 intervals if an RNR is received (RNRX = 1), or if T1 has expired due to waiting for an acknowledgement to one, or more, previously transmitted packets.

CONDITIONS FOR SENDING LINK RESET

1. FRMR received.
2. Have sent an S-command N2 times with P = 1, and at T1 intervals, without receiving an S-response with F = 1.

UNSOLICITED UA OR UNSOLICITED F BIT

If an unsolicited UA or an unsolicited F bit is received with the link up, a FRMR will be transmitted with W = 1.

SENDING AN FRMR

An FRMR may be transmitted for any of the reasons indicated in X.25 (W, X, Y, Z). An FRMR is transmitted only if the link is up.

Upon sending a FRMR, the WD2511 will not send a packet until the FRMR condition is cleared. The WD2511 will also discard any received I-field. The FRMR condition is cleared when either a SABM or DISC is received.

If an S or I-frame is received which acknowledges a previously transmitted packet(s), the acknowledgement(s) is accepted, and the appropriate ACKED bit(s) in TLOOK is set, and there is an XBA interrupt.

While in the FRMR condition, the WD2511 will act as shown below:

FRAME RECEIVED	ACTION BY WD2511
SABM	Send UA. Clear FRMR condition. Enter information transfer phase.
DISC	Send UA. Clear FRMR condition. Enter logical disconnect state.
Packet with good N(R)	Retransmit FRMR
S-frame with good N(R) (command or response)	Retransmit FRMR
Packet or S-frame with bad N(R)	Transmit new FRMR (Z = 1)
Any frame with violation W, X, Y	Transmit new FRMR

RECEIVING AN FRMR

After a FRMR has been received:

1. The FRMR I-field will be in the memory referenced by the current NE segment, provided the receiver was ready.
2. The SEND bit is cleared.
3. No more I-field data is allowed to come into memory until the user makes the receiver memory ready.
4. A link resetting SABM is transmitted, and an error interrupt, ER0 = C0, is generated at the beginning of the SABM.
5. After the UA is received for the SABM, then NA, NB, NE, V(R), and V(S) are all cleared to 0.

PROTOCOL SIGNIFICANCE OF TLOOK/RLOOK POINTERS

The NE, NA, and NB pointers have a relationship with the sequence counters used in the LAPB protocol.

The RLOOK pointer NE is equal to V(R) at all times if TRCV = 0. However, when TRCV = 1 when the link is UP, there is no guaranteed relationship between NE and V(R).

TLOOK pointer NB is the Next Block to be transmitted. If not in packet retransmission, NB is equal to the V(S) of the next new packet to be transmitted.

TLOOK pointer NA is the Next packet to be Acknowledged. It represents the V(S) number for the oldest packet in the retransmission buffer.

USE OF THE RECR BIT

The RECR (CR01) bit should be understood as an instruction to the WD2511 to initialize the receiver memory. The WD2511 will test RECR as soon as MDISC is cleared, and will test RECR after each link set-up and each link reset. Once the receiver memory is ready, the WD2511 will not test RECR again unless there is a link set-up, link reset, or a receiver-not-ready condition.

After the link is UP, and at least one packet has been received as indicated by the PKR interrupt, the user should clear RECR. This is an advantage, because if a link reset is either transmitted or received, the WD2511 will enter a receiver-not-ready condition. This will prevent packets received after the link reset from appearing to have arrived before some packets received prior to the link reset.

The receiver-not-ready condition is indicated by RNRX = 1. This condition is cleared after the user makes RECR = 1 with RECRDY = 1 (in RLOOK #0), and after either a packet or an S-frame is received.

If RECRDY of the next RLOOK is 0, but RECR = 0, there will not be an RLNR interrupt, but RNRX will be set. If RECR = 1, but the RECRDY bit of the next RLOOK segment is 0, there will be an RLNR interrupt (error code 10), and RNRX will be set.

ORDERING INFORMATION

Order Number	Maximum Rate
WD2511T-01	100 Kbps
WD2511T-05	500 Kbps
WD2511T-11	1.1 Mbps*

*Higher speeds available on special order.

The following devices have a tighter power supply specification and meet all other parameters in this data sheet.

Power Supply Range +5/−2%:

Order Number	Maximum Rate
WD2511T-91	100 Kbps
WD2511T-92	500 Kbps
WD2511T-93	1.1 Mbps

Printed in U.S.A.

PRELIMINARY

WD2520 Signalling Link Controller

FEATURES

- Performs most of the controls of the Signalling Link Control for Signalling System No. 7.
- Selectable window sizes (8, 32, 128).
- Dual channel DMA for full-duplex operation.
- Unique memory access method for buffer management.
- All formatting of bit-oriented control included: Zero bit insertion and deletion. Automatic appending and testing of flags and FCS fields.
- Automatic control of sequence numbers FSN and BSN, and of control bits FIB and BIB.
- Selectable "Basic" error correction method or the preventive cyclic retransmission error correction method.
- 48-pin dual in-line package. Pin compatible with the WD2501 and WD2511 X.25 link level controllers.
- DC to 1.1M bit/sec transmit-receive rate. (Higher bit rates available.)
- TTL compatible.

APPLICATIONS

TELEPHONE SIGNALLING
SATELLITE LINKS
PBX SWITCHING/ROUTING
HIGH RELIABILITY COMMUNICATIONS

GENERAL DESCRIPTION

Signalling System Number 7 is an international standard for common channel signalling systems, and is used within digital telecommunications networks for call control, remote control, management and maintenance signalling.

The flexible nature of the device allows more general application in systems that have long round-trip delays; notably satellite systems. As with other members of the Western Digital family of Protocol Controllers, all real-time protocol functions are fully handled by the WD2520. Thus, high speed, high efficiency communications are easily integrated into new and existing designs.

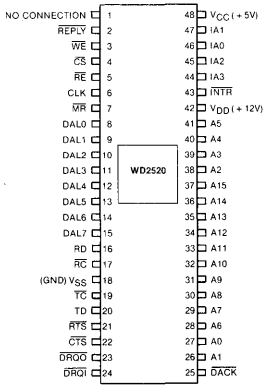

PIN DESIGNATION

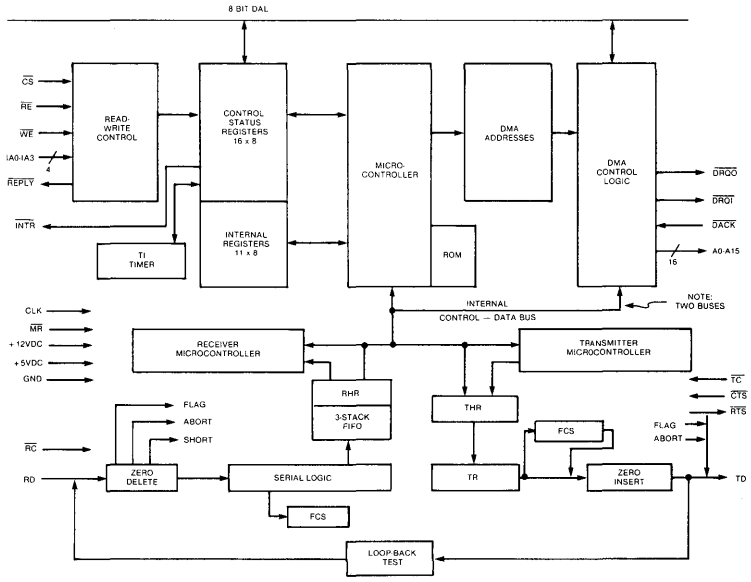

WD2520 BLOCK DIAGRAM

INTERFACE SIGNALS DESCRIPTION (All signals are TTL compatible.)

PIN NUMBER	SYMBOL	PIN NAME	FUNCTION
1		No Connection	Leave pin open
2	$\overline{REPLY}$	Reply	An active low output to indicate the WD2520 has either a $\overline{CS}$•$\overline{RE}$ or a $\overline{CS}$•$\overline{WE}$ input.
3	$\overline{WE}$	Write Enable	The data on the DAL are written into the selected register when $\overline{CS}$ and $\overline{WE}$ are low.
4	$\overline{CS}$	Chip Select	Active low chip select for CPU control of I/O registers.
5	$\overline{RE}$	Read Enable	The contents of the selected register are placed on DAL when $\overline{CS}$ and $\overline{RE}$ are low.
6	CLK	Clock	Clock input used for internal timing. Must be square wave, and greater than 250 KHz.
7	$\overline{MR}$	Master Reset	Initialize on active low. All registers reset to zero, except control bits MDISC and $\overline{LINK}$ are set to 1. $\overline{DACK}$ must be stable high before $\overline{MR}$ goes high.
8-15	DAL0-DAL7	Data Access Lines	An 8-bit bi-directional three-state bus for CPU and DMA controlled transfers.
16	RD	Receive Data	Receive serial data input.
17	$\overline{RC}$	Receive Clock	This is a 1x clock input, and RD is sampled on the rising edge of $\overline{RC}$. RD changes occur on the falling edge of $\overline{RC}$.
18	V_{SS}	Ground	Ground.
19	$\overline{TC}$	Transmit Clock	A 1x clock input. TD changes on the falling edge of TC.
20	TD	Transmit Data	Transmitted serial data output.
21	$\overline{RTS}$	Request-To-Send	An open collector (drain) output which goes low when the WD2520 is ready to transmit either flags or data.
22	$\overline{CTS}$	Clear-To-Send	An active low input which signals the WD2520 that transmission may begin. If high, the TD output is forced high. May be hard-wired to ground.
23	$\overline{DRQO}$	DMA Request Out	An active low output signal to initiate CPU bus request so the WD2520 can output onto the bus.
24	$\overline{DRQI}$	DMA Request In	An active low output signal to initiate CPU bus request so that data may be input to the WD2520. $\overline{DRQO}$ and $\overline{DRQI}$ will not be low at the same time.
25	$\overline{DACK}$	DMA Acknowledge	An active low input from the CPU in response to $\overline{DRQI}$ or $\overline{DRQO}$. $\overline{DACK}$ must not be low if CS and $\overline{RE}$ are low or if $\overline{CS}$ and $\overline{WE}$ are low.
26-41	A0-A15	Address Lines Out (See front page for Pin Assignments)	Sixteen address outputs from the WD2520 for DMA operation. If the control bit ADRV is 1, the outputs are TTL drives at all times. If ADRV is 0, the outputs are three-state, and are HI-Z whenever $\overline{DACK}$ is high. (ADRV is in Control Register #1.)
42	V_{DD}	Power Supply	+ 12VDC power supply input.

INTERFACE SIGNALS DESCRIPTION CONTINUED

PIN NUMBER	SYMBOL	PIN NAME	FUNCTION
43	$\overline{INTR}$	Interrupt Request	An active low interrupt service request output, and returns high when Status Register #1 is read.
44-47	IA0-IA3	Address Lines In (See front page for Pin Assignments)	Four address inputs to the WD2520 for CPU controlled read/write operation with registers in the WD2520. If ADRV = 0, these may be tied to A0-A3. (ADRV is in Control Register #1.)
48	V_{CC}	Power Supply	+ 5VDC power supply input.

TABLE OF CONTENTS

REFERENCE:
CCITT Yellow Book, Volume VI — Fascicle VI.6, Q.703, VIIth Plenary Assembly, Geneva, 10-21 November 1980.

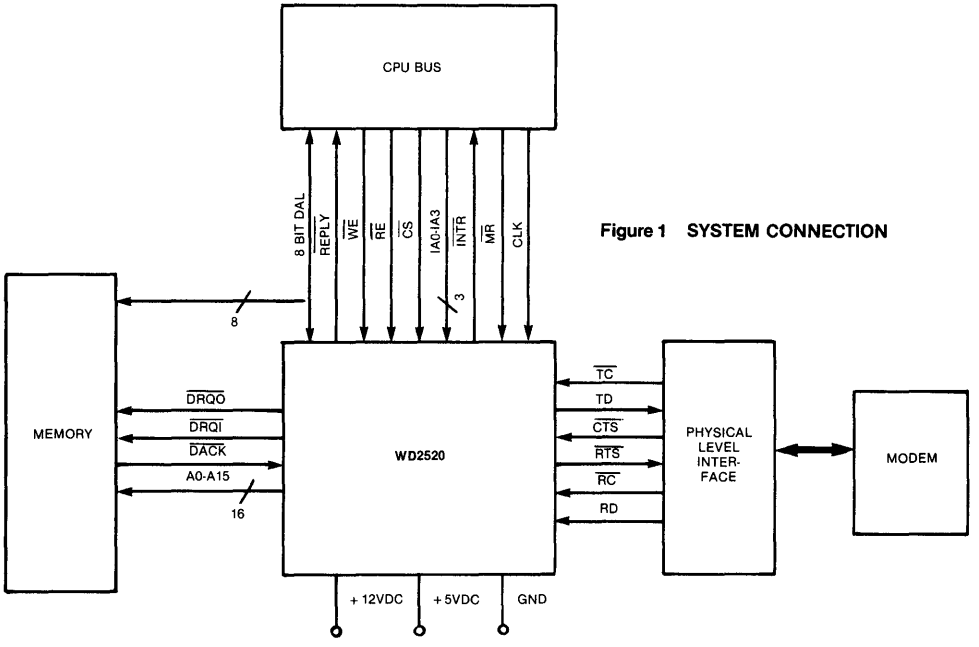

Figure 1 SYSTEM CONNECTION

1.0 ORGANIZATION

A detailed block diagram of the WD2520 is shown on the front page.

Mode control and monitor of status by the user's CPU is performed thru the 16 I/O registers.

SF, SIO, and SIF data are accessed thru DMA control. Serial data is generated and received by the bit-oriented transmitter and receiver. (See Table 1.0).

(See section 8.0 for a GLOSSARY of terms.)

Internal control of the WD2520 is by means of 3 microcontrollers: one for receive, one for transmit, and one for the internal LSI system.

Parallel transmit data is entered into the Transmitter Holding Register (THR), and then presented to the Transmitting Register (TR) which converts the data into a serial bit stream. A five bit serial buffer facilitates zero bit insert under control of the Transmitter Microcontroller. The 16 bit FCS is calculated and transmitted from the FCS register.

Parallel receive data enter the Receiver Holding Register (RHR) from a 3-stack FIFO. The FIFO input is from a 24 bit serial Receiver Register (RR). This 24 bit register prevents received FCS data from entering the RHR. The receiver FCS register is used to test the correctness of the received FCS.

Serial data is defined in blocks called Signal Units (SU). An SU may be a Fill-In Signal Unit (FISU), a Link Status Signal Unit (LSSU), or a Message Signal Unit (MSU). See Table 1.0. Each SU starts and ends with a unique Flag (01111110). In between flags, data transparency is maintained by the insertion of a 0 bit after each sequence of 5 contiguous 1's. Inserted 0 bits are stripped-off by the receiver. The last 16 bits before the closing flag is the FCS.

The transmitter of the WD2520 will always send at least 2 flags between SU's, so that each SU has its own opening and closing flag. However, the WD2520 receiver may receive as few as one flag between SU's.

The FCS calculation includes all data between opening flag and the first bit of the FCS, except for 0's inserted for transparency. The 16 bit FCS has the following characteristics:

Polynomial: $X^{16} + X^{12} + X^5 + 1$
Transmitted Polarity: Inverted
Transmitted Order: High order bit first
Preset Value: All 1's
FCS Register Value: 0001110100001111
if Received Correct (high order bit shown in LSB)

The WD2520 generates and tests Flags, FCS, inserted zeroes, FSN, BSN, FIB, BIB, and LI.

The WD2520 performs signalling link control. All FISU's are automatically generated and tested by the WD2520. The user need only be concerned with the SF field of LSSU's and the SI and SIF fields of MSU's. If the WD2520 has no outstanding MSU's or LSSU's, it will transmit contiguous FISU's.

The user may control the two extra bits in the LI field for MSU's and LSSU's. See paragraphs 3.1 and 3.2.

The WD2520 will discard all received SU's which do not meet all the minimum conditions:

1. FCS Correct.
2. A multiple of 8 bits.
3. There must be at least 5 bytes between flags.
4. The LI (Length Indicator) must be correct. (See paragraph 3.6 for the definition of a "correct" LI.)

TABLE 1.0 The Three Signal Unit (SU) Types

Numbers under SU show number of bits

Fill In Signal Unit (FISU) LI = 0

← FISU's completely controlled by WD2520 →								
FLAG	BSN	BIB	FSN	FIB	LI	00	FCS	FLAG
8	7	1	7	1	6	2	16	8

Link Status Signal Unit (LSSU) LI = 1 or 2

← appended by WD2520 →							←DMA→	←appended→ by WD2520	
FLAG	BSN	BIB	FSN	FIB	LI	XX*	SF	FCS	FLAG
8	7	1	7	1	6	2	8 or 16	16	8

Message Signal Unit (MSU) LI>2

← appended by WD2520 →							← DMA →		←appended→ by WD2520	
FLAG	BSN	BIB	FSN	FIB	LI	XX*	SIO	SIF	FCS	FLAG
8	7	1	7	1	6	2	8	8n, n>2	16	8

*Where the "X" bits are controlled by the user.

When this document refers to an SU as "received," that SU is understood to meet all the above minimum conditions. (In some cases an error counter may be incremented, but the SU is not considered as received unless all minimum conditions are met. See paragraph 3.3 and Figure 5.)

2.0 PROGRAMMING THRU REGISTERS

The WD2520 is programmed by six write/read registers, and is monitored by six status, read-only registers.

REGISTER NUMBER	IA3	IA2	IA1	IA0	REGISTER NAME
0	0	0	0	0	CR0 Control Register 0
1	0	0	0	1	CR1 Control Register 1
2*	0	0	1	0	SR2 Status Register 2
3*	0	0	1	1	SR3 Status Register 3
4*	0	1	0	0	SR4 Status Register 4
5*	0	1	0	1	SR5 Status Register 5
6*	0	1	1	0	SR6 Status Register 6
7*	0	1	1	1	SR7 Status Register 7
8	1	0	0	0	TIMER
9	1	0	0	1	
A	1	0	1	0	BUFFER HI
B	1	0	1	1	BUFFER LO
C	1	1	0	0	(not used)
D	1	1	0	1	
E	1	1	1	0	
F	1	1	1	1	

*NOTE: SR2 thru SR7 may only be read by the CPU.

CONTROL/STATUS REGISTERS SUMMARY

REG	BIT NUMBER 7	6	5	4	3	2	1	0
CR0	PCR/$\overline{\text{BASIC}}$	0	0	0	LOOP TEST	RAMT	RECR	RESET
CR1	0	0	0	ADRV	0	K2	K1	SEND
SR2	X	FSNO						
SR3	MSUR	XMA	ATTN	NAKR	RECEIVER STATE			X
SR4	X	$\overline{\text{IRTS}}$	REC IDLE	IRECR	X	PCRS	NL1	NL0
SR5	ATTENTION REGISTER							
SR6	FIBT	FSNT						
SR7	BIBT	BSNT						

NOTES:
Whenever any of the three bits MSUR, XMA, or ATTN is a 1, $\overline{\text{INTR}}$ will go low. After SR3 is read, all three bits will return to 0, and $\overline{\text{INTR}}$ will return high.

"X" represents a bit used by the WD2520's internal microcontroller. At any time these bits may be "1" or "0", and are to be disregarded by the user's CPU.

Unused control bits must be left at "0".

CONTROL REGISTER 0

REGISTER	CR07	CR06	CR05	CR04	CR03	CR02	CR01	CR00
CR0	PCR/$\overline{\text{BASIC}}$	0	0	0	LOOP TEST	RAMT	RECR	RESET

BIT	NAME	DESCRIPTION
CR00	RESET	This bit, when set, commands the WD2520 to reset SR2 thru SR5 to all 0's, and set SR6 and SR7 to all 1's. The WD2520 will not transmit nor accept serial data until RESET = 0. RESET will go to 1 when $\overline{\text{MR}}$ is enabled.
CR01	RECR	Defines the CPU's receiver memory as initially ready (RECR = 1) or initially not ready (RECR = 0). The CPU should set RECR as part of its initialization procedure, and clear RECR when either IRECR (SR44) is set or when the first MSUR interrupt occurs. See paragraph 3.1.4.
CR02	RAMT	Enables a self test for registers internal to the WD2520. (See section 7.0 on SELF-TESTS.)
CR03	LOOP TEST	Internally ties TD back to RD, and $\overline{\text{TC}}$ to $\overline{\text{RC}}$. The normal inputs of RD and $\overline{\text{RC}}$ are logically disconnected. There will still be an output at TD, and $\overline{\text{TC}}$ needs the transmitter clock. (See 7.0)
CR04-CR06	0	Unused control bits, like CR04, must be left at 0.
CR07	PCR/$\overline{\text{BASIC}}$	This bit selects one of two error correction procedures. If CR07 = 0, the "basic" error correction method is employed which uses retransmission, positive acknowledge, and negative acknowledge. (See section 4.0) If CR07 = 1, the Preventive Cyclic Retransmission (PCR) error control method is used. PCR is essentially a noncompelled forward error correction method. (See section 5.0)

CONTROL REGISTER 1

REGISTER	CR17	CR16	CR15	CR14	CR13	CR12	CR11	CR10
CR1	0	0	0	ADRV	0	K2	K1	SEND

BIT	NAME	DESCRIPTION
CR10	SEND	Used to command the WD2520 to transmit the next MSU, MSU's, or LSSU. If SEND = 1, the WD2520 will read from the next TLOOK segment the value of RTX (See Table 3.1). If RTX = 0, the WD2520 will clear SEND, and no transmission occurs. If RTX = 1, the WD2520 will then read TSADR and LIT from the TLOOK segment, and the buffer (either an MSU or LSSU) is transmitted. After transmission, the WD2520 will clear RTX of the segment just transmitted. If an MSU has been transmitted, the RTX of the next TLOOK segment is now read, and the procedure above is repeated. (See section 3.0) As a matter of good practice, the user's CPU should set SEND each time a new MSU or LSSU is ready to transmit.
CR12 and CR11	K2, K1	Used to define the window size. The window defines the maximum number of outstanding (unacknowledged) MSU's allowed. The window also defines the size of TLOOK and RLOOK.
CR14	ADRV	The ADRV bit is the control for the 16 bit output addresses (A0-A15). If ADRV = 0, the outputs are three-state, and are in HI-Z, except when DACK goes low. If ADRV = 1, the outputs are always low impedance (TTL), and are high level when DACK is high.
CR13, 15, 16, 17	0	Unused control bits.

K2	K1	SIZE OF TLOOK/RLOOK (K + 1)	NUMBER OF OUTSTANDING MSU'S (K)
0	0	8	7
0	1	32	31
1	0	128	127

STATUS REGISTER 2

REGISTER	SR27	SR26 ← SR20
SR2	X	FSNO

BIT	NAME	DESCRIPTION
SR26←SR20	FSNO	Represents the FSN value of the next transmitted MSU to be acknowledged. FSNO relates to the particular TLOOK segment for that MSU. Notice that if the retransmission buffer is empty; FSNO (SR2) = FSNT (SR6) + 1.
SR27	X	Used internally, may be 1 or 0.

STATUS REGISTER 3

REGISTER	SR37	SR36	SR35	SR34	SR33	SR32	SR31	SR30
SR3	MSUR	XMA	ATTN	X	X	X	X	X

BIT	NAME	DESCRIPTION
SR30 to SR34	X	Used internally may be 1 or 0.
SR35	ATTN	The ATTN (Attention) bit indicates that either: 1) An error has occurred which is unrecoverable by the WD2520, or 2) A significant event has occurred, and the CPU must be informed. The ATTN Register (SR5) will contain the coded reason for the ATTN interrupt.
SR36	XMA	Means that a previously transmitted, or retransmitted, MSU, or MSU's, has been acknowledged.
SR37	MSUR	Stands for Message Signal Unit Received. This means that an MSU was received has met the minimum conditions of section 10 and that the received FSN was equal to BSNT + 1. (Where the BSNT referred to is the value of BSNT before the MSU was received.) Also, the receiver memory had to have been ready. This interrupt also means that the SIO and SIF fields are already in RAM, the BSNT has been incremented, and an acknowledge will be transmitted at the next opportunity. (The "next opportunity" to acknowledge an MSU received is with the transmission of the BSN in the next transmitted SU.)

STATUS REGISTER 4

REGISTER	SR47	SR46	SR45	SR44	SR43	SR42	SR41	SR40
SR4	X	$\overline{IRTS}$	REC IDLE	IRECR	X	PCRS	NL1	NL0

BIT	NAME	DESCRIPTION
SR40, SR41	NL0, NL1	Represent a binary number for the segment where the next received LSSU will be placed in the RLSSU (See 3.2).
SR42	PCRS	Used for internal control.
SR43	X	Used internally may be 1 or 0.
SR44	IRECR	A status bit which informs the CPU that the receiver memory through RLOOK has been set-up. The CPU should clear RECR either when IRECR = 1, or after the first MSUR interrupt. After the WD2520 sets IRECR, the WD2520 will not test RECR again until a reset condition, or until a C3 interrupt (see paragraph 2.3.2). Also see paragraph 3.1.4.
SR45	REC IDLE	Indicates that the WD2520 has received at least 15 contiguous 1's. When at least one 0 is received. REC IDLE goes to 0, and remains 0 until 15 contiguous 1's.
SR46	$\overline{IRTS}$	Stands for the Internal Request-To-Send bit, and indicates that the transmitter is attempting (successful or not) to send data. If the $\overline{RTS}$ pin (pin 21) is not tied to ground or WIRED-ORED to another signal, then $\overline{IRTS}$ = $\overline{RTS}$.
SR47	X	Used internally may be 1 or 0.

STATUS REGISTER 5

REGISTER	SR57 ← SR50
SR5	ATTN

BIT	NAME	DESCRIPTION
SR57←SR50	ATTN	The Register shows the coded reason for the ATTN interrupt. See Table XX for description of ATTN codes.

TABLE XX ATTN Register Codes

HEXIDECIMAL CODE	REASON
C1	TUR. Transmitter Under-Run.
C2	ROR. Receiver Over-Run.
C3	RLNR. RLOOK not ready. The RTR bit of the next RLOOK segment is 0.
01	LOST LSSU. An LSSU was received, but temporarily lost, because the RTR bit of the next RLSSU segment was 0. See Table 3.1.
02	At least two out of three consecutively received FIB's indicate start of retransmission when no retransmission was requested. OR, at least two out of three consecutively received BSN's are not valid. A valid received BSN is defined as less than or equal to FSNT and greater than or equal to FSN0-1.
04	The current TLOOK segment for MSU has an illegal byte count. The illegal count is all 0's, and the WD2520 will not attempt to transmit a null MSU.
20	An LSSU has been received and is in the RLSSU.
40	An LSSU has been transmitted.
80	No SU has been received within time T1.

STATUS REGISTERS 6 & (7)

REGISTER	SR67 (SR77)	SR66 ← SR60 (SR76 ← SR70)
SR6 (SR7)	FIBT (BIBT)	FSNT (BSNT)

BIT	NAME	DESCRIPTION
SR66←SR60 (SR76←SR70)	FSNT (BSNT)	The current value of transmitted FSN. SR6 is incremented prior to the transmission of the next new MSU. SR6 applies only to new MSU's, and not MSU's which may be retransmitted. The expected value of the FSN of the next received MSU is SR7 plus one. SR7 (BSNT) also represents the current value of the transmitted BSN. FSNT and BSNT are initialized to all 1's.
SR67 (SR77)	FIBT (BIBT)	The current values of the transmitted FIB and BIB, respectively. The WD2520 will invert FIBT and initiate retransmission of MSU's in response to a received inverted BIB. If the WD2520 receives an incorrect FSN, BIBT will be inverted on the next transmitted SU. (NOTE: A correct received FSN is equal to FSN0 for an MSU, or FSN0-1 for an FISU or LSSU.) FIBT and BIBT are each initialized to 1.

TIMER REGISTER LSB & MSB

REGISTER	TR7 (0)	TR6 (0)	TR5 (0)	TR4 (0)	TR3 (0)	TR2 (0)	TR1 (TR9)	TR0 (TR8)
TIMER LSB	T7	T6	T5	T4	T3	T2	T1	T0
TIMER MSB	(0)	(0)	(0)	(0)	(0)	(0)	(T9)	(T8)

BIT	NAME	DESCRIPTION
TR9←TR0	T9←T0	Registers 8 and 9 define a 10 bit value for timer T1. The LSB is in Register 8, Bit 0, and the MSB is in Register 9, Bit 1. The upper six bits of Register 9 are not used. The time of T1 is given by the expression: T1 = 16,384 N/CLK where CLK = Frequency of CLK input (pin 6) N = Binary value of T1 in Registers 8 and 9. The timer is re-started each time an SU is received. If no SU is received within T1, and ATTN interrupt is generated with ATTN code 80.

BUFFER HI←REGISTER A
BUFFER LO←REGISTER B

REGISTER	7	6	5	4	3	2	1	0
BUFFER HI	B15	B14	B13	B12	B11	B10	B9	B8
BUFFER LO	B7	B6	B5	B4	B3	B2	B1	B0

BIT	NAME	DESCRIPTION
B15-B8 B7-B0	BUFFER HI BUFFER LO	Registers A and B represent the 16 bit starting address of the 24 byte buffer. TLOOK follows that buffer. RLOOK follows TLOOK, RLSSU follows RLOOK, and the Error Counters follow RLSSU. (See Figure 3.) RLOOK and RLOOK will each have either 8, 32, or 128 segments according to the values of K2 and K1.

Regarding the TUR and ROR errors; an ROR means that the receiver Register (RR) had a byte to load into the receiver FIFO, but the FIFO was full. TUR means that the Transmitting Register (TR) needed a byte from the Transmitter Holding Register (THR), but the THR was empty. Either TUR or ROR may be caused by one of two conditions, or both.

1) The bit rate clock, $\overline{TC}$ or $\overline{RC}$, is too fast for the WD2520.
2) The $\overline{DACK}$ response is too slow for the bit rate.

In addition to the above, an ROR could be caused as a result of an RLNR condition. If there is an RLNR condition, and an MSU or LSSU is received before the RLNR is cleared, there will be an ROR interrupt which will override the RLNR.

3.0 TRANSMITTING AND RECEIVING PROCEDURES

Referring to Table 1.0, notice that the SF of LSSU's and the SIO and SIF of MSU's are DMA accessed. The DMA is part of the Memory Access Method (buffer management) shown in Figure 3.

The starting address of a TLOOK (Transmitter Look-Up) table follows the 24 byte buffer which starts at the address given in Registers A and B. TLOOK has K + 1 segments, and each segment consists of four bytes. TLOOK is followed contiguously by RLOOK (Receiver Look-Up) which also has K + 1 segments of four bytes each. The value of K + 1 is programmed by control bits CR12 and CR11.

RLSSU (Received LSSU's) is a four segment table which follows RLOOK. Each segment has four bytes.

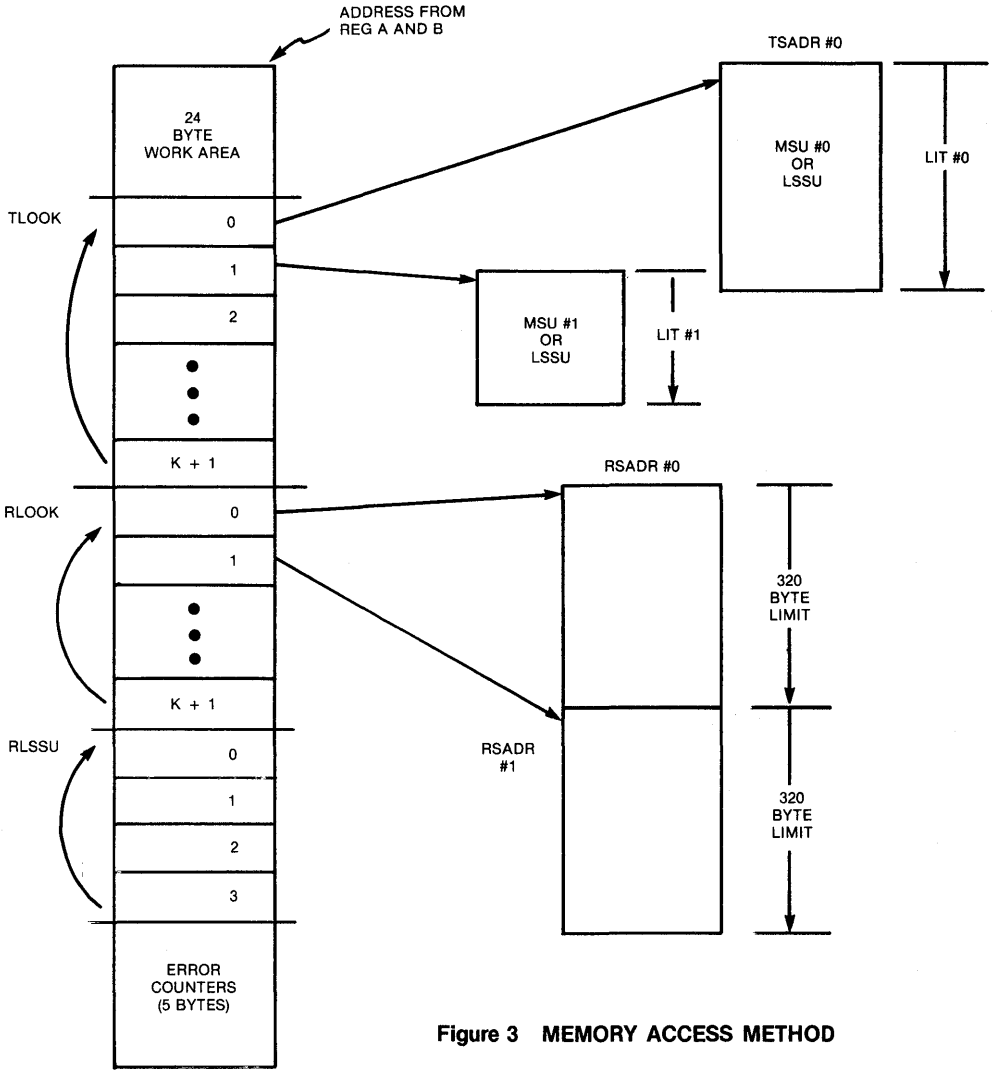

Figure 3 MEMORY ACCESS METHOD

NOTES:
1. ALL RAM IN THE FIGURE IS EXTERNAL TO THE WD2520.
2. K + 1 = 8, 32, OR 128. (SEE CONTROL BITS K2 AND K1).

Lastly, there are five error counters of one byte each. These error counters represent errors recoverable by the WD2520, while those error conditions reported by the ATTN Register are not recoverable by the WD2520.

The 24 byte buffer ahead of TLOOK is to be reserved for the WD2520 for internal control. The beginning of this work area is defined by Registers A and B. The WD2520 will use that work area as follows:

OFFSET BYTE NUMBER	USE
0, 1	Retransmission pointer address
2, 3	FSNO address
4, 5	FSN address
6, 7	BSN address
8, 9	RLOOK starting address
A, B	LSSU pointer address
C, D	LSSU starting address
E thru 17	Reserved

Note that the WD2520 will compute the above addresses, and the user's CPU need only reserve the space. The CPU may read the 24 byte buffer, but must not write into that buffer.

3.1 Sending and Receiving MSU's

The segment number in TLOOK maintains a relationship with the transmitted FSN. The segment number will correspond with the least significant bits of FSNT. For example, if K + 1 = 32, the segment numbers vary from 0 to 1F (hexidecimal) and equal the 5 LSB's of the transmitted FSN. Thus, segment #3 could correspond to transmitted FSN 03, 23, 43, or 63 (hexidecimal).

In like manner, the RLOOK segment has a relationship to the expected value of the received FSN, and the actual expected FSN of the next MSU is maintained is SR7 (BSNT + 1).

To transmit an MSU, the SEND bit (CR10) must first be set. After SEND is set, the WD2520 will go to the next segment (as determined by the value of SR6) and test bit RTX (Ready To Transmit). See Table 4. If RTX = 1, the WD2520 will transmit the MSU with length determined by LIT, ELFT, EX7 and EX6. Prior to the transmission, FSNT will be incremented. The SIO and SIF fields are found in a memory block beginning at address TSADR (Transmitter Starting ADdRess). Of course, LI>2 for an MSU. After sending the MSU, the WD2520 will clear RTX, and go on to the next segment.

When the WD2520 reaches a segment whose RTX = 0, the WD2520 clears the SEND bit. Then, the WD2520 will begin retransmission of the oldest MSU (as determined by FSNO) if PCR is selected (CR07 = 1), or the WD2520 will transmit FISU's if BASIC is selected (CR07 = 0).

After a previously transmitted, or retransmitted, MSU(s) is acknowledged, the XMA bit (SR36) is set and an interrupt is generated. The ACK'ED bit in each TLOOK segment acknowledged will have been set prior to the XMA interrupt.

For all transmitted MSU's and LSSU's, TL6 and TL7 are placed in the SU's byte #3, bits 6 and 7, without alteration.

The length of the MSU is determined as follows: If ELFT (Extended Length Field for Transmitter), EX7, and EX6 are all 0, the MSU will send the length equal to LIT. Also, the LI field in the transmitted MSU will be equal to the TLOOK segment, byte #3 (TL7, TL6, and LIT).

If either ELFT, EX7, or EX6 are set to 1, an MSU with length greater than 62 bytes will be transmitted. The lower 6 bits of LI will be all 1's. The length of the transmitted MSU will be the combination of EX7, EX6, and LIT. If the MSU length is to be greater than 255 bytes, than ELFT is set to 1.

NOTE:
The WD2520 will not send an MSU with length greater than 320 bytes. If the user attempts to put a value greater than 320, the WD2520 will send the 320 bytes, and ABORT. However, this will allow more than 272 bytes as stated in Q.703, paragraph 2.3.8.

To receive MSU's, the WD2520 will test the RTR (Ready To Receive) bit in the next RLOOK segment (as determined by the value of SR7 plus one). If RTR = 1, the WD2520 will prepare to receive beginning at address RSADR (Receiver Starting Address) with maximum buffer size of 320 bytes. If an MSU is received (LI>2), and if the MSU meets all the conditions stated under the description of SR37, then the WD2520 will store the received length in the lower 6 bits of byte #3 of the applicable RLOOK segment. The RTR is cleared to 0 and MRCVD (MSU ReCeiVeD) is set to 1 by the WD2520. After this, an interrupt is generated with MSUR (SR37) set to 1. The WD2520 will go on to the next RLOOK segment.

For all received MSU's and LSSU's, the values of the received SU's byte #3, bits 6 and 7, are placed in RL6 and RL7 without alteration.

The received length is determined as follows: If ER7, ER6, and ELFR are all 0, the length is contained in the six bit number LIR, and the received length is less than, or equal to, 63.

If an MSU of length greater than 63 bytes is received, the received LI field will be all 1's. However, under this condition, at least one of the bits ER7, ER6, or ELFR will be set to 1. The LIR will contain the lower six bits of the received length byte count. ER6 and ER7 are added directly to LIR. If ELFR is set, an MSU greater than 255 bytes has been received.

For example, suppose an MSU of length 70 (hex 46) bytes is received. The LI field in the MSU will be all 1's. In the RLOOK segment, the bits ELFR and ER7 will be 0, and ER6 will be 1. LIR will be 6 (000110).

Like the transmitter, the WD2520 will not receive an MSU of length greater than 320 bytes. The WD2520 will not write into received memory past 320 received bytes per MSU.

If the WD2520 reaches an RLOOK segment with RTR = 0, an ATTN interrupt is generated with ATTN code C3. The WD2520 will not test that RTR again until the user's CPU sets RECR (CR01) to a 1. Therefore, if RECR is set, but the RTR of the next RLOOK is 0, the WD2520 would perform nearly constant DMA's while testing RTR. To prevent this, the CPU should set RECR after a reset, but clear RECR after either IRECR (SR44) is set to 1, or after the first MSUR interrupt. Then, if an ATTN code C3 interrupt happens, the CPU has the opportunity to set-up the next RLOOK before setting RECR, and the WD2520 will not tie-up the DMA channel.

TABLE 4 Segment Descriptions (Note: S = Spare)

	7	6	5	4	3	2	1	0	
Byte 0	ACKED	S	S	S	EX7	EX6	ELFX	RTX	
Byte 1	TSADR HI								TLOOK Segment
Byte 2	TSADR LO								
Byte 3	TL7	TL6	LIT						

	7	6	5	4	3	2	1	0	
Byte 0	MRCVD	S	S	S	ER7	ER6	ELFR	RTR	
Byte 1	RSADR HI								RLOOK Segment
Byte 2	RSADR LO								
Byte 3	RL7	RL6	LIR (length > 2)						

	7	6	5	4	3	2	1	0	
Byte 0	LRCVD	S	S	S	S	S	S	RTR	
Byte 1	RL7	RL6	LIR (1 or 2)						RLSSU Segment
Byte 2	First SF byte received								
Byte 3	Second SF byte received								

3.2 Sending and Receiving LSSU's

If an LSSU is ready to transmit, the same initial procedure will be followed as for an MSU except that the LIT value is 1 or 2, and ENLF, EX7, and EX6 are all 0. After the LSSU is transmitted, the SEND bit (CR10) is cleared, and RTX is cleared. FSNT of SR6 is not incremented, and the WD2520 does not go on to the next segment. After transmission of the LSSU, an ATTN interrupt with code 40 is generated.

When an LSSU is received, the LIR will be either 1 or 2. In addition to the minimum conditions for received SU's (see 1.0), the WD2520 will test the received FSN to be equal to the value of SR7. If it is, the LI and SF fields are placed in the next RLSSU segment as determined by NL1 and NL0 (SR41 and SR40) provided that the RLSSU's segment RTR = 1. (If that RTR = 0, an ATTN interrupt with code 01 is generated.) The WD2520 will then set LRCVD (LSSU ReCeiVeD) and clear RTR, and generate an ATTN interrupt with code 20. The NL1/NL0 pointer is incremented.

Bits TR7 and TR6 are transmitted in byte #3 bits 7 and 6 without alteration. The received bits 7 and 6 in byte #3 are placed in ER7 and ER6 without alteration.

3.3 Error Counters

Following contiguously after RLSSU are five 8-bit error counters. The WD2520 will increment each counter at the occurrence of each defined event. However, the WD2520 will not increment past 255 (all 1's). The CPU has the responsibility of clearing each counter. The first counter past RLSSU is #0, etc.

ERROR COUNTER NUMBER	EVENT DESCRIPTION
0	Received SU's with FCS error (includes ABORTed SU's where an ABORT is 7 contiguous 1's).
1	Number of SU's received which were not a multiple of 8 bits.
2	Received short SU's (less that 5 bytes between flags).
3	Number of SU's received incorrect LI.
4	Number of SU's received with incorrect FSN.

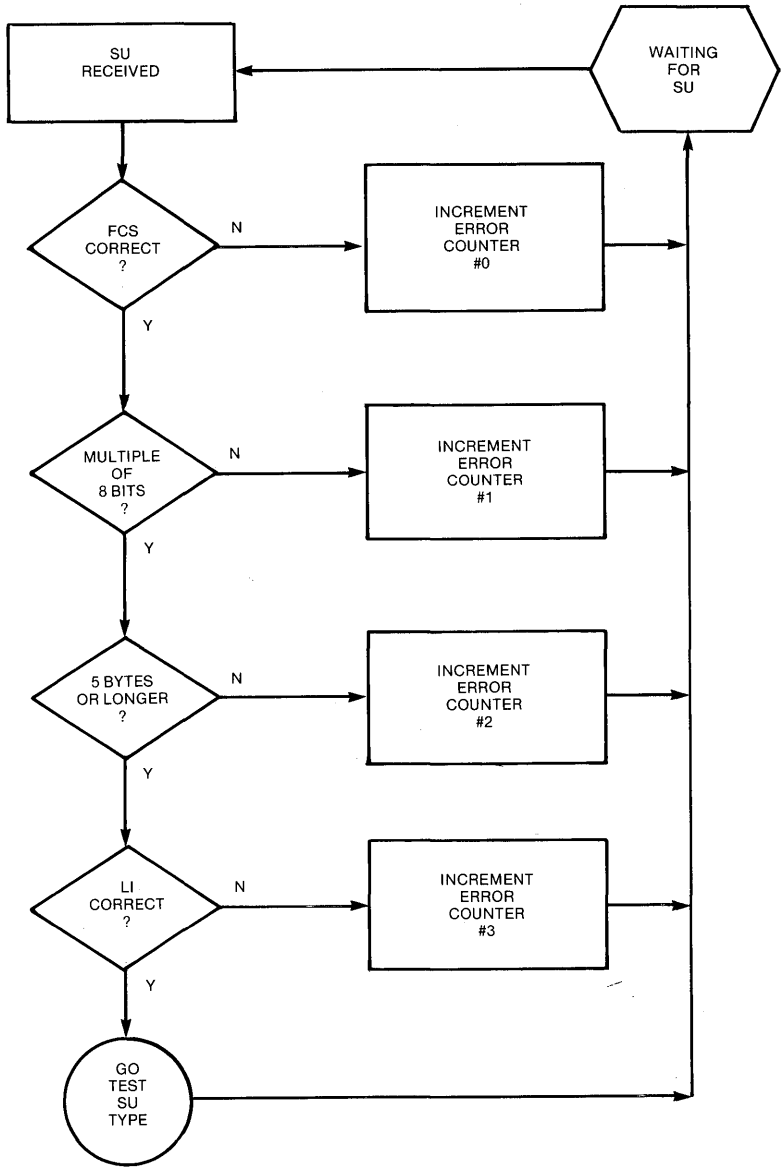

Figure 5 RECEIVER FUNCTIONAL FLOW

The error counters will not be incremented unless the received SU had a correct FCS.

These counters are intended for statistical purposes, only, and are not guaranteed to be incremented at each occurrence of the defined event.

3.4 Deadly Embrace Prevention

A "deadly embrace" is defined here as a state where the user's CPU and the WD2520 are both waiting for an action from the other. This situation could exist unless control bits shared by the WD2520 and the CPU are handled properly. A shared control bit is defined as one which is either set by the CPU and cleared by the WD2520, or is set by WD2520 and

cleared by the CPU. They must both not be allowed to set or clear the same bit. The shared control bits are shown below:

BIT NAME	WHERE FOUND	SET BY	CLEARED BY
SEND	CR10	CPU	WD2520
RTX	each TLOOK segment	CPU	WD2520
ACK'ED	each TLOOK segment	WD2520	CPU
RTR	each RLOOK segment	CPU	WD2520
MRCVD	each RLOOK segment	WD2520	CPU
RTR	each RLSSU segment	CPU	WD2520
LRCVD	each RLSSU segment	WD2520	CPU

Any of these bits may be set to any value by the CPU at initialization. However, after initialization, the deadly-embrace rule must be followed.

3.5 Initialization Procedure

Initialization begins after either the RESET bit is cleared after having been set, or after the CFB bit is cleared after having been set.

At initialization, transmitted FSN's and BSN's are all 1's. Transmitted FIB's and BIB's are 1.

The first transmitted MSU will have FSN = 0, and the expected value of the FSN of the first received MSU will be 0.

3.6 Definition of Correct LI

The LI field is found in each SU as the lower 6 bits in the third byte following the leading flag.

LI is the byte count for the length of the SU. The length of the SU is the number of bytes following the LI field, but preceding the FCS.

For an FISU, LI = 0, and a correct LI means that the FCS field will immediately follow the LI field.

For an LSSU, LI = 1 or 2, and a correct LI means that there are 1 or 2 bytes in the SF field (corresponding to LI).

A correct LI for an MSU will meet either one of two conditions:

For MSU's with lengths from 3 to 63, the LI must equal the number of bytes in the length of the MSU.

For MSU's which are 63 bytes, or longer, the LI will be maintained at a value of 63 (six 1's).

4.0 "BASIC" ERROR CORRECTION METHOD

In the basic method, previously transmitted, but unacknowledged, MSU's are not retransmitted unless a NACK (Negative ACKnowledge) is received. The WD2520 will employ this procedure when CR07 = 0.

A NACK is defined as an SU with the BIB inverted from the previously transmitted SU. When the WD2520 receives a NACK, the WD2520 will begin retransmission of MSU's with the transmitted FSN equal to the received BSN plus one, and the transmitted FIB equal to the received BIB. However, if the received NACK refers to an FSN not in the retransmission buffer, the NACK is discarded, but the next transmitted FIB will still be made equal to the received BIB. If 2 out of 3 consecutively received BSN's reference FSN's of MSU's not in the retransmission buffer, an ATTN interrupt is generated with code 02.

When a valid NACK is received, the WD2520 will begin sequential retransmission starting with the oldest MSU even if there are more new MSU's ready.

A received NACK may also acknowledge one, or more, previously transmitted MSU's.

All previously transmitted MSU's may be acknowledged by a received NACK. In this case, a received NACK will have cleared the retransmission buffer. This is an acceptable condition, and the acknowledgements are accepted. The FIBT is inverted. However, there will be no retransmission of MSU's.

The retransmission buffer is defined as those MSU's awaiting acknowledgement from FSNO to FSNT.

In BASIC, retransmission will not begin unless a NACK is received.

If the WD2520 receives an MSU with an FSN not equal to the BSNT + 1, or if an FISU or LSSU is received with an FSN not equal to BSNT, the WD2520 will transmit a NACK.

5.0 ERROR CORRECTION BY PREVENTIVE CYCLIC RETRANSMISSION (PCR)

In the PCR procedure, previously transmitted, but unacknowledged, MSU's are sequentially retransmitted whenever there are no new MSU's or LSSU's to transmit. This procedure is employed if CR07 = 1.

The absence of new MSU's or LSSU's is determined if the RTX bit of the next TLOOK segment is 0, or if SEND (CR10) is 0, or if there are K outstanding MSU's (K = 7, 31, or 127 depending on control bits K2 and K1). If there are no new MSU's or LSSU's to transmit, and if the retransmission buffer is not empty, the WD2520 will begin retransmission with the oldest MSU. If there are no new MSU's or LSSU's, and if the retransmission buffer is empty, the WD2520 will transmit continuous FISU's.

PCR does not require a NACK to be received before beginning retransmission. However, received NACK's are treated the same in PCR as in BASIC. Also, the WD2520 may send a NACK in PCR just as in BASIC.

The CPU must determine when the maximum number of bytes available for retransmission (NMAX) has been reached. NMAX, measured in bytes, is the total of the lengths of the MSU's in the retransmission buffer. When NMAX is reached, the CPU

must not allow new MSU's to be transmitted by insuring that the RTX bit of those new MSU's is zero. However, new LSSU's may be transmitted. When the CPU has determined that new MSU's may be transmitted (because some of the retransmission buffer has been removed by received acknowledgements), then the CPU may add new MSU's.

6.0 SUMMARY OF CONDITIONS FOR TRANSMISSION

No SU's will be transmitted until RESET (CR00) and CFB (CR17) are both 0, and $\overline{CTS}$ (pin 22) is low. Given these conditions, the WD2520 will always be transmitting SU's. If there are no MSU's nor LSSU's to send, the WD2520 will send FISU's.

The following conditions require that RESET = 0 and that $\overline{CTS}$ is low. In the following paragraphs, the following conventions are used: X means "don't care." (N) refers to the applicable TLOOK segment.

CONDITIONS FOR SENDING AN LSSU

SEND = 1
CR07 = X
RTX(N) = 1
LIR(N) = 1 or 2

After sending the LSSU, RTX(N) and SEND are cleared to 0.

CONDITIONS FOR SENDING AN FISU

Either condition A or B below will cause an FISU to be transmitted.

A. "Basic" procedure selected, and no new MSU's nor LSSU's ready for transmission, and no retransmission in progress.

B. PCR procedure selected, and no new MSU's nor LSSU's ready for transmission, and the retransmission buffer is empty.

SEND = 0
CR07 = 1
FSNT (SR6) + 1 = FSNO (SR2)

CONDITIONS FOR SENDING AN MSU FOR THE FIRST TIME

In either "basic" or PCR, the conditions below will cause an MSU to be transmitted for the first time. After the transmission is complete, RTX(N) is cleared to 0. FSNT (SR6) is incremented prior to the transmission of that MSU.

SEND = 1
CR07 = 0
RTX(N) = 1
LIR(N) > 2
There are less than K outstanding MSU's

CONDITIONS FOR RETRANSMITTING AN MSU

Retransmission of one or more MSU's will occur whenever conditions A, B or C below are met. MSU's are sequentially retransmitted beginning with FSNO.

A. NACK received with received BSN greater than or equal to FSNO, but less than FSNT. Notice that this condition requires that the retransmission buffer is not empty. (The retransmission buffer is empty when FSNT + 1 = FSNO).

B. PCR procedure, and no new MSU's nor LSSU's ready, and retransmission buffer not empty.

C. PCR procedure, and there are K outstanding MSU's.

7.0 SELF TESTS

There are two self tests: 1) An internal RAM test, and 2) A loop-back test. Both tests are suitable for manufacturing testing, user incoming inspection testing, diagnostics, and trouble-shooting.

INTERNAL RAM TEST

There are eleven 8-bit registers in the WD2520 which are not directly accessible by the user's CPU. Four of these cannot be completely tested by using loop-back. Therefore, this test was created as a means to test those four internal registers. The contents of Register A are placed in two even internal registers, and the contents of Register B are placed in two odd internal registers. The four registers are then added together, without carry, and the results are placed in Registers SR2, SR5, SR6, and SR7. This test is initiated when RAMT (CR02) is set to 1 provided RESET (CR00) is also set to 1. Use the following procedure:

1. Insure RESET = 1
2. Set RAMT
3. Set-Up Registers A and B
4. Wait at least 150 times the CLK period
5. Check the results in Registers SR2, SR5, SR6, SR7

To repeat the test for new values in Registers A and B:

6. Clear RAMT
7. Wait at least 120 times the CLK period
8. Go back to step 1

During the execution of this test, the ROM ID code of the WD2520 will be displayed in the lower 5 bits of SR4.

LOOP-BACK TEST

The internal loop-back condition exists whenever CR03 = 1. TD is tied back to RD and $\overline{TC}$ is tied to $\overline{RC}$. The inputs, RD and $\overline{RC}$ are gated-out, but the TD output is still present.

When running this test, it is recommended that MSU's be transmitted and received and the SIO and SIF fields be checked. The transmitted and received SIO and SIF fields should be identical. When this is done, the entire DMA and buffer management is tested as well as the transmitter and receiver.

LSSU's may also be transmitted and received in loop-back.

8.0 GLOSSARY OF TERMS

BASIC	The "basic" error correction method
BIB	Backward Indicator Bit (1 bit)
BIBT	BIB to be transmitted
BSN	Backward Sequence Number (7 bits)
BSNT	BSN to be transmitted
F	Flag
FCS	Frame Check Sequence
FIB	Forward indicator Bit (1 bit)
FIBT	FIB to be transmitted
FISU	Fill In Signal Unit (L = 0)
FSN	Forward Sequence Number (7 bits)
FSNO	FSN of next transmitted MSU to be acknowledged (Also called FSN of oldest MSU in retransmission buffer)
LI	Length Indicator (1 byte field)
LSSU	Link Status Signal Unit (L = 1 or 2)
MSU	Message Signal Unit (L>2)
NACK	Negative ACKnowledgement
NMAX	A procedure which complements PCR, and is used to limit the number of bytes available for retransmission.
PACK	Positive ACKnowledgement
PCR	Preventive Cyclic Retransmission (one of two error correction methods)
SF	Status Field (1 or 2 byte field in an LSSU)
SIF	Signalling information Field (2 or more byte field in an MSU)
SIO	Service Indicator Octet (1 byte field in an MSU)
SU	Signal Unit (may be FISU, LSSU, or MSU)
SU length	The number of bytes in an SU after the LI field, but before the FCS.

9.0 WD2520 ELECTRICAL SPECIFICATIONS:

ABSOLUTE MAXIMUM RATINGS:

Voltages referenced to V_{SS}

High Supply Voltage (V_{DD}) . . . . . −0.3 to 15V
Voltage at any Pin . . . . . −.03 to 15V
Operating Temperature Range . . . . . 0°C to +70°C
Storage Temperature Range . . . . −55°C to +125°C

NOTE:
Maximum limits indicate where permanent device damage occurs. Continuous operation at these limits is not intended and should be limited to those conditions specified in the DC Electrical characteristics.

Operating DC Characteristics: V_{SS} = 0V, V_{CC} = 5.0V ± .25, V_{SS} = 12.0V ± .6V T_A = 0° to 70°C

SYMBOL	PARAMETER	MIN	TYP	MAX	UNIT	CONDITIONS
I_{DD}	V_{DD} Supply Current		20	70	mA	
I_{CC}	V_{CC} Supply Current		200	280	mA	
V_{DD}	High Voltage Supply	11.4	12	12.6	V	
V_{CC}	Low Voltage Supply	4.75	5	5.25	V	
V_{IH}	Input High Voltage	2.4			V	
V_{IL}	Input Low Voltage			0.8	V	
V_{OH}	Output High Voltage	2.8			V	I_O = −0.1mA
V_{OL}	Output Low Voltage			0.4	V	I_O = 1.6mA
I_{LH}	Input Leakage Source or Sink			10	uA	
I_{OZH}	Input Leakage High Impedance			10	uA	Vin = V_{CC}
I_{OZL}	Output Leakage High Impedance			10	uA	Vin = V_{SS}

AC Timing Characteristics: V_{CC} = 5V ± .25; V_{DD} = RV ± .6V, V_{SS} = 0V, T_A = 0° to 70°C

SYMBOL	PARAMETER	MIN	TYP	MAX	UNIT	CONDITIONS
CLK	Clock Frequency	0.5		2.0	MHz	Note 1, 2
RC	Receive Clock Range	0			MHz	Note 3, 4
TC	Transmit Clock Range	0			MHz	Note 4
MR	Master Reset Pulse Width	10			mS	
T_{AR}	Input Address Valid to RE	0			nS	
T_{RD}	Read Strobe (or DACK Read) to Data Valid			375	ns	Note 5
T_{HD}	Data Hold Time From Read to Strobe			80	nS	
T_{HA}	Address Hold Time From Read Strobe	80			ns	
T_{AW}	Input Address Valid to Trailing Edge of WE	200			nS	
T_{WW}	Minimum WE Pulse Width	200			nS	
T_{DW}	Data Valid to Trailing Edge of WE or Trailing Edge of DACK for DMA Write	100			nS	
T_{AHW}	Address Hold Time After WE	80			nS	
T_{DHW}	Data Hold Time After WE or After DACK for DMA Write	100			nS	
T_{DA1}	Time From DRQO (or DRQI) to Output Address Valid if ADRV = 1			80	nS	
T_{DA0}	Time From DACK to Output Address Valid if ADRV = 0			400	nS	Note 5
T_{DD}	Time From Leading Edge of DACK to Trailing Edge of DRQO (or DRQI)			400	nS	Note 5
T_{DAH}	Output Address Hold Time From DACK			100	nS	
T_{DMW}	Data Hold Time From DACK For DMA Read			100	nS	
T_{RP1}	REPLY Response Time (leading edge)			240	nS	Note 5
T_{RP2}	REPLY Response Time (trailing edge)			260	nS	Note 5

NOTES:
1. Clock must have 50% duty cycle.
2. Buffer chaining is not guaranteed when CLK is greater than 1.5 MHz.
3. Residual bit detection logic not guaranteed when RC is greater than 50Kbps.
4. See "Ordering Information" for maximum serial rates.
5. C(load) = 100pf

DMA TIMING

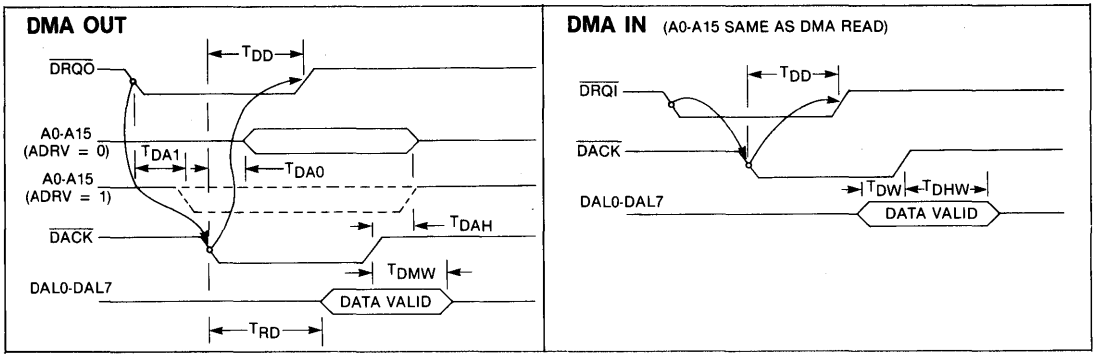

CPU READ/WRITE TIMING

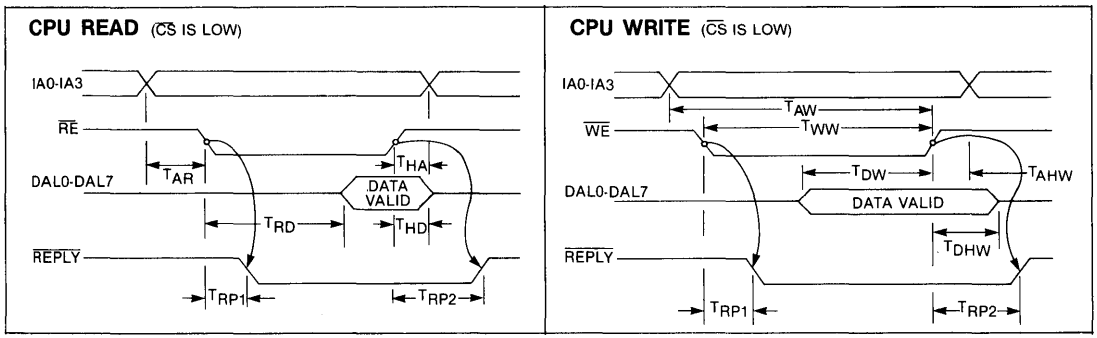

***NOTE:** There must not be a CPU read or write (CS-RE or CS-WE) within 500 nanoseconds after the trailing (rising) edge of $\overline{\text{DACK}}$.

There must not be the leading (falling) edge of $\overline{\text{DACK}}$ allowed within 500 nanoseconds after the completion of a CPU write (CS-WE).

$\overline{\text{RE}}$ and $\overline{\text{WE}}$ must not be low at the same time.

ORDERING INFORMATION

ORDER NUMBER	MAXIMUM RATE
WD2520T-01	100 Kbps
WD2520T-05	500 Kbps
WD2520T-11	1.1 Mbps

Printed in U.S.A.

PRELIMINARY

WD2840 Local Network Token Access Controller

FEATURES

- Broadcast Medium Oriented (Coax, RF, CATV, IR, etc.)
- Up to 254 nodes
- Dual DMA/Highly efficient Memory Block Chaining
- Token based protocol
- Acknowledge option on each datagram
- Adjustable fairness, stations may be prioritized
- Frame format similar to industry standard HDLC
- Supports Global Addressing
- Diagnostic Support: Self-Tests, System and Network
- TTL Compatible

APPLICATIONS

The WD2840 is a general purpose Local Network Token Controller applicable to virtually all types of multi-point communications applications. The token protocol allows the sharing of one bus by up to 254 nodes. WD2840's will be designed into process control equipment, micro-computers, mini-computers, personal computers, proprietary micro-processor based applications, intelligent terminals, front-end processors, and similar equipment.

The great advantage for the design engineer is the ease with which he can implement a local network function. The WD2840 handles autonomously all major communications tasks as they relate to the local network function.

GENERAL DESCRIPTION

The WD2840 is a MOS/LSI device intended for local network applications, where reliable data communications over a shared medium is required. The device uses a buffer chaining scheme to allow efficient memory utilization. This scheme minimizes the host CPU time requirements for handling packets of data. The WD2840 frees the host CPU from extensive overhead by performing network initialization, addressing, coordination, data transmission, acknowledgements and diagnostics.

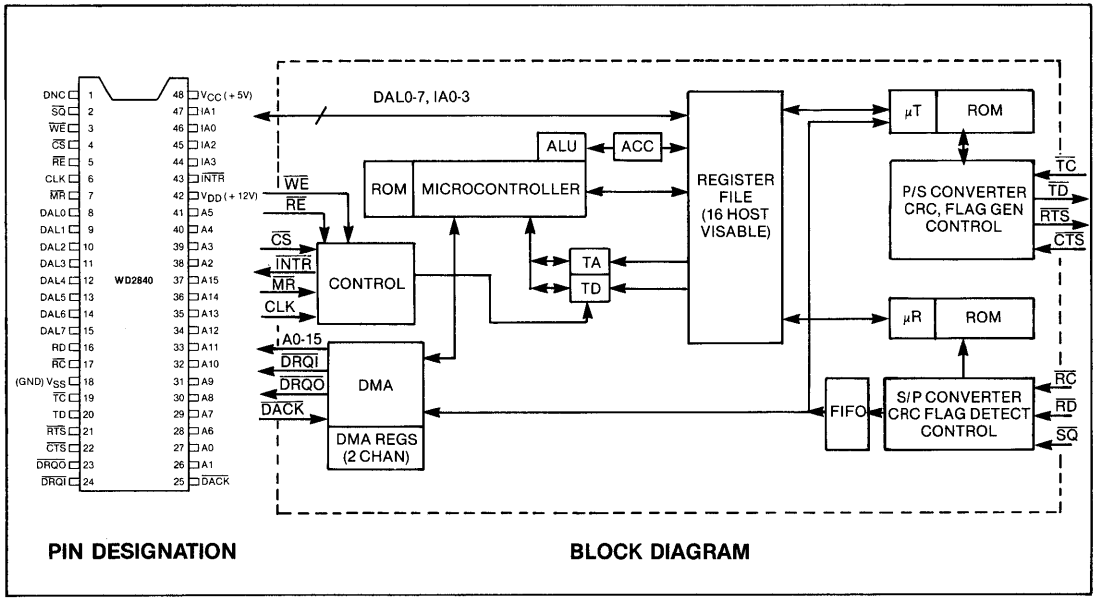

PIN DESIGNATION **BLOCK DIAGRAM**

Figure 1.1

CONTENTS

RELATED DOCUMENTS

Consult Western Digital Corporation for current application notes and related documents.

WD2840 LOCAL NETWORK TOKEN ACCESS CONTROLLER

INTRODUCTION

The WD2840 is a single LSI device which gives systems designers the ability to include networking capabilities into their unique products simply and economically.

A general and fundamental advantage to the use of complex LSI in a given system is the partitioning of required technical expertise. A successful user of the WD2840 need not be a data-communications expert, and further, he need not be at all concerned with low level network details (though these details are documented and available to him if he is interested). The potential user of the WD2840 must simply evaluate the communications facilities provided by the device to determine its suitability for the intended use.

The WD2840 is designed to logically interconnect 2 to 254 user devices over a shared communications medium. Examples of expected mediums include coax cable, twisted pair bus, RF, and CATV. All network control functions, such as data framing and error checking, destination filtering, fair and adjustable transmission scheduling, and network initialization and fault recovery (caused by noise for example) are handled completely by the WD2840.

The protocol implemented allows guaranteed station access intervals allowing applications in factory automation and other critical communications environments where "statistical delays" are not acceptable. The WD2840 token protocol also allows the addition and/or removal of stations to a network at anytime, including while operating.

Serious attention has also been given to the user's interface to the device. The interface is a combination of conventional I/O registers and an elaborate DMA buffer chaining interface. This chaining feature allows the user much more efficient use of his system memory, particularly in situations where the maximum message sent over the network is much longer than the average size. This feature also allows the automatic queueing of messages independently of the user's consumption rate, in effect, speed decoupling the user's CPU and processing requirements from the network.

The WD2840 has several parameters (registers) that allow tailoring to the user's requirements. In this way, network priority and access ordering, to name two, can be manually set if desired.

Using an integrated version of these network algorithms saves not only the development costs already mentioned, but further, the total processing power required for the user's application is not increased. In other words, a CPU upgrade can likely be avoided by "distributing" the network processing task into LSI devices such as the WD2840.

SCOPE

This document differs from traditional LSI data sheets in that it details not only the LSI implementation of a function, but also defines the overall function in detail. Specifically, this document includes details of the communications protocol implemented by the WD2840 Token Access Controller.

The document is organized into three main sections:

SECTION ONE is much like a traditional data sheet including register descriptions, pin definitions, and hardware architecture.

SECTION TWO describes the interfaces to the WD2840. The network side is conventional, the host side consists of an elaborate DMA interface with control blocks and WD2840/host handshaking.

SECTION THREE details the network protocol implemented by the device. Normal operation, initialization, and the handling of error conditions are described.

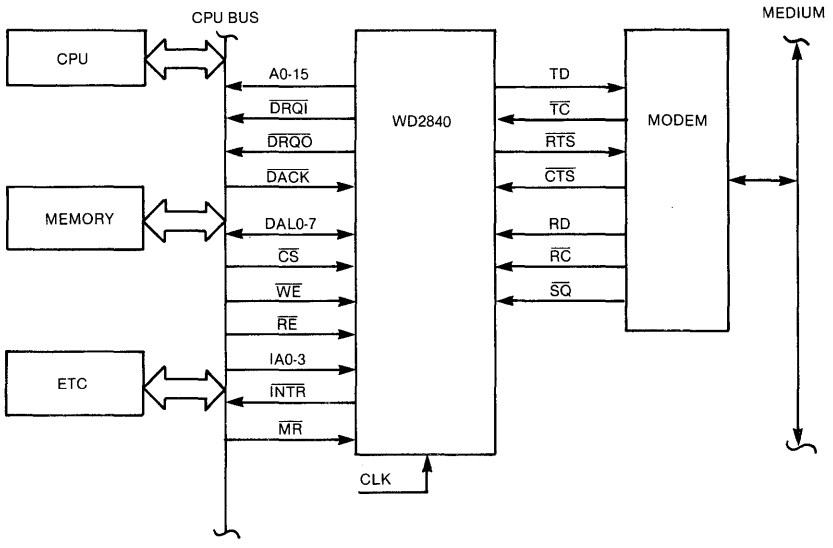

TYPICAL SYSTEM CONNECTION

1.1 PIN DEFINITIONS

PIN NUMBER	SYMBOL	PIN NAME	FUNCTION
1	DNC	DO NOT CONNECT	Leave pin open.
2	$\overline{SQ}$	SIGNAL QUALITY	An active low input which signals the WD2840 that a frame may be received. The modem may negate this signal if its receive signal quality is below a reliability threshold, ensuring that the WD2840 will not accept the frame.
3	$\overline{WE}$	WRITE ENABLE	The data on the DAL are written into the selected register when $\overline{CS}$ and $\overline{WE}$ are low. $\overline{RE}$ and $\overline{WE}$ must not be low at the same time.
4	$\overline{CS}$	CHIP SELECT	Active low chip select for CPU control of I/O registers.
5	$\overline{RE}$	READ ENABLE	The content of the selected register is placed on DAL when $\overline{CS}$ and $\overline{RE}$ are low.
6	CLK	CLOCK	Clock input used for internal timing.
7	$\overline{MR}$	MASTER RESET	Initialize on active low. All registers reset to zero, except control bit ISOL is set to 1. $\overline{DACK}$ must be stable high before $\overline{MR}$ goes high.
8-15	DAL0-7	DATA ACCESS LINES	An 8-bit bi-directional three-state bus for CPU and DMA controlled data transfers.
16	RD	RECEIVE DATA	Receive serial data input.
17	$\overline{RC}$	RECEIVE CLOCK	This is a 1X clock input, and RD is sampled on the rising edge of $\overline{RC}$.
18	V_{SS}	GROUND	Ground.
19	$\overline{TC}$	TRANSMIT CLOCK	A 1X clock input. TD changes on the falling edge of $\overline{TC}$.
20	TD	TRANSMIT DATA	Transmitted serial data output.
21	$\overline{RTS}$	REQUEST-TO-SEND	An open collector output which goes low when the WD2840 is ready to transmit either data or flags.
22	$\overline{CTS}$	CLEAR-TO-SEND	An active low input which signals the WD2840 that transmission may begin.
23	$\overline{DRQO}$	DMA REQUEST OUT	An active low output signal to initiate CPU bus request so that the WD2840 can output onto the bus.
24	$\overline{DRQI}$	DMA REQUEST IN	An active low output signal to initiate CPU bus requests so that data may be input to the WD2840.
25	$\overline{DACK}$	DMA ACKNOWLEDGE	An active low input from the CPU in response to DRQO or DRQI. DACK must not be low if $\overline{CS}$ and $\overline{RE}$ are low or if $\overline{CS}$ and $\overline{WE}$ are low.
26-41	A0-A15	ADDRESS LINES OUT	Sixteen address outputs from the WD2840 for DMA operation.
42	V_{DD}	POWER SUPPLY	+ 12VDC power supply input.
43	$\overline{INTR}$	INTERRUPT REQUEST	An active low interrupt service request output. Returns high when Interrupt Register is read.
44-47	IA0-IA3	ADDRESS LINES IN	Four address inputs to the WD2840 for CPU controlled read/write operations with registers in the WD2840. If ADRV = 0, these may be tied to A0-A3.
48	V_{CC}	POWER SUPPLY	+ 5VDC power supply input.

1.2 DEVICE ARCHITECTURE

A detailed block diagram of the WD2840 is shown in Figure 1.1.

Mode control and monitor of status by the user's CPU is performed through the Read/Write Control circuit, which reads from or writes into registers addressed by IA0-IA3.

Transmit and receive data are accessed through DMA control. Serial data is generated and received by the bit-oriented controllers.

Internal control of the WD2840 is by means of three internal micro-controllers; one for transmit, one for receive, and one for overall control.

Parallel transmit data is entered into the Transmitter Holding Register (THR), and then presented to the Transmitter Register (TR) which converts the data to a serial bit stream. The Frame Check Sequence (FCS) is computed in the sixteen bit CRC register, and the results become the transmitted FCS.

Parallel receive data enters the Receiver Holding Register (RHR) from the 24 bit serial Receive Register (RR). The 24-bit length of RR prevents received FCS data from entering the RHR. The receiver CRC register is used to test the validity of the received FCS. A three level FIFO is included in the receiver.

The WD2840 sends all information, network control and user data, in blocks called frames. Each frame starts and ends with a single flag (binary pattern 01111110). In between flags, data transparency is provided by the insertion of a zero bit after all sequences of five contiguous one bits. The receiver will strip the inserted zero bits. (See section on frame format for location of address, control, and FCS fields.)

1.3 REGISTER DEFINITION

The WD2840 is controlled and monitored by sixteen 8 bit registers. This set of registers consists of two Control Registers, three Status Registers, an Interrupt Event Register, a Counter Register and a variety of Parameter Registers. In general the host is responsible for defining these registers (except certain host read-only registers: SR0-2, IR0, CTR0 and NA) to contain proper and meaningful values prior to entering Network Mode from Isolate State. Furthermore, while the WD2840 is in Network Mode, the CBP (H,L) and MA registers must not be changed by the host. Register NAR may be changed arbitrarily but will only be considered by the WD2840 in response to the NEWNA (CR10) control bit being set. The two Control Registers and the TA, TD, AHOLT, TXLT registers may change dynamically to control the behavior of the WD2840.

REG [1]	NAME	DESCRIPTION
0	CR0	Control Register 0
1	CR1	Control Register 1
2[2]	SR0	Status Register 0
3[2]	IR0	Interrupt Event Register
4[2]	SR1	Status Register 1
5[2]	SR2	Status Register 2
6[2]	CTR0	Counter Register 0
7[2]	NA	Next Address
8	TA	ACK Timer
9	TD	Net Dead Timer
A	CBPH	Control Block Pointer (MSB)
B	CBPL	Control Block Pointer (LSB)
C	NAR	Next Address, Request
D	AHOLT	Access Hold-off Limit
E	TXLT	Transmit Limit
F	MA	My Address

[1] = Hexadecimal representation of IA0-IA3.
[2] = CPU read only, write not possible.

Control, status, and interrupt bits will be referred to as CR, SR, or IR, respectively, along with two digits. For example, SR21 refers to status register #2 and bit 1, which is "STATE."

SUMMARY — CONTROL, STATUS, INTERRUPT REGISTERS

REGISTER	BIT # 7	6	5	4	3	2	1	0
CR0	TXDEN	TXEN	RXEN	TOFF	ILOOP	COPY	NOINT	ISOL [1]
CR1[2] CR1[4]	DIAGC DIAGC	PIGT 0	INIT 0	ADRV ADRV	GIRING DMAT	0 LOOPT	0 RAMT	NEWNA NUDIAG
SR0	LASTF	SENDACK	L2	0	BSZ3	BSZ2	BSZ1	BSZ0
SR1	TAOUT	$\overline{\text{IRTS}}$	RECIDL	1	1	1	1	1
SR2	NXTT0	NXTR0	TR	ACKRQ	RETRY	TSENT	STATE	INRING
IR0 [3]	ITUR	IROR	INS	ITRAN	IREC	ITOK	ITA	ITD/M
NOTE: ZERO BITS (0) SHOWN ABOVE ARE RESERVED AND SHOULD NOT BE USED.								

NOTES:

[1] = Set to 1 on power-up or master reset.

[2] = Non diagnostic mode only (CR17-DIAGC cleared).

[3] = Any bit set causes host interrupt ($\overline{\text{INTR}}$ goes true) when Master Interrupt Suppress (CR01) is clear. All bits are cleared when register is read by the host.

[4] = Diagnostic State only (CR17-DIAGC set). See diagnostic section for register usage in diagnostic mode.

CR0 — CONTROL REGISTER 0 DEFINITION

REGISTER	CR07	CR06	CR05	CR04	CR03	CR02	CR01	CR00
CR0	TXDEN	TXEN	RXEN	TOFF	ILOOP	COPY	NOINT	ISOL

BIT	NAME	DESCRIPTION
CR00	ISOL	Isolate. Set true on power up or master reset. Host clears this bit after the host memory based WD2840 control block and other WD2840 registers have been set up. May be set by the host at any time (will be ignored if WD2840 is in diagnostic state). There is some delay for the WD2840 to respond to any state change request. A state change to network mode is acknowledged by the state confirmation status bit (SR21-STATE) being cleared. Setting ISOL while the WD2840 is in Network State will cause a state change to Isolate State, confirmed by an interrupt event (IR00-ITM) and the STATE status bit (SR21) being set. This transaction will be delayed until the node does not possess the token. Any in-progress frame transmission will be completed normally (at the current frame, regardless of queue length), followed by a normal token pass sequence.
CR01	NOINT	Master Interrupt Suppress. When clear, the WD2840 will generate host interrupt requests ($\overline{\text{INTR}}$ low) if any bit in the WD2840 interrupt request register (IR0) is set. When set, only the interrupt request is suppressed, not the setting of bits in IR0. Note that any interrupt request will be dropped by the WD2840 when IR0 is read since this will clear IR0.
CR02	COPY	Enables COPY mode. When set causes all received data frames to be accepted and DMA'ed into memory regardless of destination address. (See description in Diagnostics Section.)
CR03	ILOOP	Instructs the WD2840 to loop data internally from transmitter to receiver. Used with the LOOP diagnostic. Must NOT be set while in network mode (CR00-ISOL clear).
CR04	TOFF	When set causes WD2840 to ignore timers. This is NOT intended to be used in an operational network but is provided to support network diagnosis. CAUTION: This control bit disables all automatic network error recovery.

BIT	NAME	DESCRIPTION
CR05	RXEN	Receive Data Enable. When clear, the WD2840 still makes normal responses to supervisory frames (scan, token pass), but will not DMA any data frames into memory and ignores the receiver buffer chain. However any data frame which is addressed to this node and for which an ACK is requested, will be NAK'ed with a "receiver not enable" Nak code. When RXEN is set, it allows the receiver to DMA appropriate data frames into memory. RXEN may be arbitrarily set and reset while in Network State but changes will not affect any frames in progress. NOTE: Even when RXEN is clear, the WD2840 is "following" the receiver buffer chain with an internal register pointing either to the next available buffer (NXTR0 set) or, if the chain is exhausted, to a link field of zero (NXTR0 clear). The constraints on host manipulation of the receiver buffer chain are the same regardless of the state of RXEN. See the subsequent section on Receiver Memory Interface for more details.
CR06	TXEN	Master Transmit Enable. When clear no transmissions will occur and the transmit buffer chain will be ignored. When set, transmission activity is further dependent upon TXDEN (CR07). NOTE: Even when TXEN is clear, the WD2840 is "following" the transmitter buffer chain with an internal register pointing either to the next frame to transmit (NXTT0 set) or, if the chain is exhausted, to a link field of zero (NXTT0 clear). The constraints on host manipulation of the transmitter buffer chain are the same regardless of the state of TXEN. See the subsequent section on Transmitter Memory Interface for more details.
CR07	TXDEN	Data Transmit Enable. Has no meaning unless TXEN is set. When set in conjunction with TXEN, normal WD2840 transmission of data and supervisory frames will occur. When clear and with TXEN set, only data frame transmission will be suppressed. That is, token pass and Ack/Nak supervisory frames will still be transmitted when appropriate. NOTE: The note above for TXEN applies.

CR1 — CONTROL REGISTER 1 DEFINITION

REGISTER	CR17	CR16	CR15	CR14	CR13	CR12	CR11	CR10
CR1	DIAGC	PIGT	INIT	ADRV	GIRING	0	0	NEWNA
CR1	DIAGC	0	0	ADRV	DMAT	LOOPT	RAMT	NUDIAG

BIT	NAME	DESCRIPTION (CR17 = 0, Network mode)
CR10	NEWNA	Update NA register. When set causes WD2840 to copy the contents of register NAR into register NA. The WD2840 clears this bit after the function is complete. This mechanism allows the host to define the WD2840's successor in the logical ring. The node's next token pass will be to the new NA node. NOTE: The normal token pass recovery applies. If the token pass to the new NA is not successful, a normal scan sequence will occur where the WD2840 attempts a single token pass to each node address in numerical sequence until a successful pass occurs or the node's address itself is reached.
CR11	—	(Not used, Reserved.)
CR12	—	(Not used, Reserved.)
CR13	GIRING	Get in logical ring. Instructs the WD2840 to gain entry into the logical ring at the next opportunity (i.e. respond to a token pass). The INRING status bit (SR20) is confirmation; when INRING is set, it indicates that the WD2840 is participating in a logical ring of at least two nodes. If the host clears GIRING while INRING is set, the WD2840 will not accept the next token pass to it at which time INRING will be cleared as confirmation.

BIT	NAME	DESCRIPTION (CR17 = 0, Network mode)
CR14	ADRV	Address Driver Enable. Enables the sixteen output address (A0-A15). If ADRV = 0, the outputs are tri-state and are in HI-Z, except when $\overline{DACK}$ goes low. If ADRV = 1, the outputs are always TTL levels.
CR15	INIT	Network Initialization Enable. When clear, the WD2840 will not attempt to (re)initialize the network if the net dead timer (TD) expires. When set, TD timer expiration causes the WD2840 to enter Scan Mode. In this mode it transmits a token pass frame to each node numerically higher in address, one after another, until either network activity occurs (another node responds) or until the node's own address is reached. When Scan Mode begins, the first node address used is the then current NA (Next Address) node address. This value is derived from and is affected by the following actions: 1. At transition into Network State it defaults to MA + 1. 2. It may be set by the host using the NAR register and the NEWNA (CR10) control flag. 3. Upon receipt of a Scan Mode frame, NA is redefined to MA + 1. The successful initialization of the network by Scan Mode causes NA to be defined as the first responding node (hence, this node's successor). All node address computations are ascending and circular within the valid node address range of 1-254. NOTE: Since this network initialization activity comes about because of a timer expiration, TOFF (CR04) must be clear.
CR16	PIGT	If set, instructs WD2840 to piggy back token on last data frame transmitted. This request is honored if the last frame is determined as a result of limit TXLT or the LAST bit set in the TFSB, but not if transmission ends due to the reaching of the end of the chain.
CR17	DIAGC	Enables diagnostic mode. In network mode this bit should be zero.

CR1 — CONTROL REGISTER 1 DEFINITIONS

BIT	NAME	DESCRIPTION (CR17 = 1, Diagnostic mode)
CR10	NUDIAG	Perform a new diagnostic. When set causes WD2840 to perform the selected diagnostics. The host initializes the appropriate registers for the particular diagnostic and by setting this bit can initiate the test. The WD2840 clears this bit after completion of the diagnostic.
CR11	RAMT	Selects internal RAM test if in diagnostic mode.
CR12	LOOPT	Selects Loop Test if in diagnostic mode.
CR13	DMAT	Selects DMA Test if in diagnostic mode.
CR14	ADRV	Address Driver Enable. Enables the sixteen output address (A0-A15). If ADRV = 0, the outputs are tri-state and are in HI-Z, except when $\overline{DACK}$ goes low. If ADRV = 1, the outputs are always TTL levels.
CR15	—	(Not used, Reserved.)
CR16	—	(Not used, Reserved.)
CR17	DIAGC	Enables diagnostic mode. Confirmation of diagnostic mode is via status bit STATE (SR21). When DIAGC and STATE are both set, diagnostic functions of CR1 apply. When DIAGC is cleared, after the selected set of diagnostics in progress complete, the WD2840 will transition to the Isolate state. This transition will cause an interrupt event (ITM).

SR0 — STATUS REGISTER 0 DEFINITION

REGISTER	SR07	SR06	SR05	SR04	SR03	SR02	SR01	SR00
SR0	LASTF	SENDACK	L2	0	BSZ3	BSZ2	BSZ1	BSZ0

BIT	NAME	DESCRIPTION
SR00 :: SR03	BSIZ	Buffer Size, multiples of sixty-four bytes (the multiple ranges from 0 to 15, meaning buffers are 64 to 1024 bytes). This is an indication only (i.e. read only). The value is obtained from the host memory-based WD2840 control block which is read when the WD2840 transitions from Isolate State to Network State (see ISOL-CR00). See the subsequent section on memory interface for more details.
SR04	—	Not used.
SR05	L2	An internal flag set during frame transmission if the length value of the current frame is equal to eight. For normal data frame transmission this means the frame has no data field and for transparent frame transmission this means the frame is an access control frame.
SR06	SENDACK	An internal flag set during data frame reception to indicate that the incoming frame should be acknowledged (send ack/nak frame). This flag is cleared when the acknowledgement has been transmitted.
SR07	LASTF	An internal flag set during data frame transmission to indicate that the current frame will be the last to be transmitted during this access period. Five situations can cause this to occur: 1) ISOL (CR00) becoming set, 2) TXDEN (CR07) becomes clear, 3) current frame flagged (via FSB) to be "last frame," 4) the current token frame count reaching the TXLT limit, 5) transmitter under-run detection. Note in particular that the last frame in the transmit queue will not cause LASTF to set since it's being last is not known until frame end. Also if a piggy-back token is permitted (CR16 set) and no acknowledge is requested (via FSB), the token will be piggybacked on the current (last) data frame. LASTF is not cleared until the next data frame transmission begins.

SR1 — STATUS REGISTER 1 DEFINITION

REGISTER	SR17	SR16	SR15	SR14	SR13	SR12	SR11	SR00
SR1	TAOUT	IRTS	RECIDL	1	1	1	1	1

BIT	NAME	DESCRIPTION
SR10 :: SR14	—	(Not used, reserved.)
SR15	RECIDL	Receiver Idle. Indicates the WD2840 has received at least 15 contiguous ones.
SR16	$\overline{\text{IRTS}}$	Internal Request To Send. Indicates the transmitter is attempting (successful or not) to send either data or flags. If the $\overline{\text{RTS}}$ pin is not tied to ground or WIRE-ORED with another signal, then $\overline{\text{IRTS}} = \overline{\text{RTS}}$.
SR17	TAOUT	Timer TA expired.

SR2 — STATUS REGISTER 2 DEFINITION

REGISTER	SR27	SR26	SR25	SR24	SR23	SR22	SR21	SR20
SR2	NXTTO	NXTRO	TR	ACKRQ	RETRY	TSENT	STATE	INRING

BIT	NAME	DESCRIPTION
SR20	INRING	In logical ring. Indicates the node has had the token and has successfully passed it at least once (therefore it is included in a logical ring of at least two nodes). See GIRING (CR13) for other comments.
SR21	STATE	Mode confirmation. Depending on DIAGC (CR17), the WD2840 is either in Isolate or Diagnostic state. When ISOL (CR00) is set, STATE set confirms the WD2840 is not in Network State. When ISOL is clear, STATE clear confirms Network State. Note any state transition into Isolate State causes an interrupt event to occur (ITM).
SR22	TSENT	An internal flag. TSENT is set when the WD2840 passes the token. It may have been either a piggyback or explicit token pass frame. TSENT is cleared when the next frame is received.
SR23	RETRY	An internal flag which is set when either a data frame or a token pass frame must be retransmitted. Data frames are only retransmitted if they have an acknowledge request and no response at all occurred. Token pass frames (except Scan) are retransmitted if no network activity was detected. Both of these situations are detected as a result of a TA timeout.
SR24	ACKRQ	An internal flag set during data frame transmission if an acknowledgement is requested for the specific frame. If this is the case, the WD2840 pauses to await the ACK/NAK response frame; if the TA timer expires before the response, a single retry will occur (see RETRY-SR23). ACKRQ is not cleared until the beginning of the next data frame transmission.
SR25	TR	An internal flag set when the WD2840 receives a token passed to it. It is cleared when the token is passed (or if it is ignored for any reason. For example, piggyback token on a bad data frame, TXEN clear, or detection of duplicate tokens in the logical ring).
SR26	NXTR0	Internal Receive Buffer Pointer State. Because of the linked list approach used in the buffer chains, the WD2840 internal register used to follow the list is either pointing to the next buffer in the chain or at the address of the next buffer in the chain (prior buffer's link field). The WD2840 will always advance along the chain so that it has the address of the next buffer to be used. However, when a zero link is encountered, the WD2840 retains the link field address expecting eventually that the chain will be extended by the host making the link some non-zero value. When the WD2840 actually needs the next buffer, it looks again at the contents of the link field expecting it to have been changed (chain extended) to the address of an available buffer. The NXTR0 bit differentiates between these two situations. When set it indicates the WD2840 has the address of the next buffer and that all prior frames (denoted by posted FSB's) can be removed from the chain for received frame processing by the host. When NXTR0 is clear it indicates that the WD2840 has advanced to a zero link (end of chain). NOTE: In this situation, the last posted frame CANNOT be removed from the chain for processing since it is the link field of his last buffer that must be set in order to extend the receiver buffer chain.
SR27	NXTT0	Internal Transmit Buffer Pointer State. The comments for NXTR0 (SR26) apply (in an analogous manner) to NXTT0 since the transmit buffer chain is handled by the WD2840 using an identical scheme. When NXTT0 is set it indicates that the WD2840 has the address of the next frame to transmit in its internal register. However when clear, it indicates that the transmit chain internal register points to the link field of the last buffer of the last transmitted frame. This link field contained zero when first read. For the transmit case, this is a normal situation corresponding to no data frames to transmit. NOTE: As in the receive case, when NXTT0 is set, all previously transmitted frames (denoted by posted FSB's) can be removed from the chain for reuse. However, when NXTT0 is clear it indicates that the transmit chain must be extended by the host before removing the very last frame that has been transmitted (posted).

IR0 — INTERRUPT REGISTER DEFINITION

REGISTER	IR07	IR06	IR05	IR04	IR03	IR02	IR01	IR00
IR0	ITUR	IROR	INS	INTRAN	IREC	ITOK	ITA	ITD/M

The setting of any bit in this register by the WD2840 causes an interrupt request ($\overline{INTR}$ = low) if NOINT (CR01) is clear. The reading of this register by the host clears all bits (and any interrupt request).

BIT	NAME	DESCRIPTION (1)
IR00	ITD/M	Network dead or mode change (dual use). When in Network mode, timer TD expiring (with TOFF clear) causes this bit to be set to indicate no network activity has occurred within the timeout period. Also INRING (SR20) is cleared and, if INIT (CR15) is set, the WD2840 will enter Scan Mode (see INIT - CR15 for details). Transition from Network or Dignostic State to the Isolate State will be confirmed by this interrupt. The choice between the ITD and ITM interpretations is easily made based on the ISOL (CR00) bit.
IR01	ITA	Date Frame Transmission Unsuccessful. This interrupt indicates that a transmitted data frame with an acknowledge request was not successfully acknowledged. Either a NAK or no response after two transmissions will cause this. The exact cause can be determined by inspecting the appropriate FSB.
IR02	ITOK	The token has been received.
IR03	IREC	Data Frame Received. This interrupt signifies that a good data frame has been properly received and DMA'ed into the buffer chain. Frames that have been received can be identified by following the buffer chain noting the WD2840 frame status bytes (FSB). A non-zero FSB (host must clear when queuing free buffers) indicates a properly received frame. The host may freely remove all received frames from the chain up to but NOT necessarily including the last one posted. The last one posted may only be removed if the WD2840 NXTR0 (SR26) is set. For more details see the explanation for NXTR0.
IR04	INTRAN	Indicates that at least one data frame has been transmitted. The number of frames transmitted and the status of each (i.e. ACK/NAK, retry count) is determined by following the transmit chain and inspecting frame status bytes (FSB). All transmitted frames up to but NOT including the last posted may be freely removed. The last one posted may only be removed if the WD2840 NXTT0 (SR27) is set. For more details see the explanation for NXTT0.
IR05	INS	New successor. The WD2840 has identified a new successor in the logical ring. This happens when the prior successor either failed to respond to a token pass or as instigated by a network scan frame.
IR06	IROR	Receiver over-run. The WD2840 ran out of buffers or access to the DMA channel was delayed by the host so long as to cause loss of received data.
IR07	ITUR	Transmitter under-run. The WD2840 was delayed access to the DMA channel by the host long enough to cause loss of transmitted data (an abort was sent). The frame is not retried and the token is passed.

(1) = Non diagnostic mode only. See diagnostic section for register usage for diagnostics.

OTHER REGISTER DEFINITIONS

NAME	DESCRIPTION
CTR0	Running Limit Counter. Used by the WD2840 for Access Hold-Off Limit (AHOLT) checking and Transmit Limit (TXLT) checking. When transmitting data frames CTR0 is used for TXLT counting; otherwise it is used for AHOLT counting. The counter runs from zero to the 8-bit limit value.
NA	Next Address. This register shows the current (instantaneous) successor node in the network logical ring. For validity, the WD2840 should be "in the ring" (see GIRING - CR13 and INRING - SR20 for more details). The successor node may be changed for a variety of reasons: 1. Any attempted token pass that fails twice will cause the WD2840 to attempt to locate a new successor by sequentially trying token passes to successively higher node addresses beginning with NA + 1. 2. A received Scan frame will cause NA to be set to MA + 1. If the next token pass fails case 1 applies. 3. The host may arbitrarily redefine NA by using the NAR register and the NEWNA (CR10) control bit. At a convenient point the WD2840 recognizes NEWNA, copies NAR into NA, then clears NEWNA as confirmation. If the next token pass fails case 1 applies.
TA	Acknowledgement Timer. Value of maximum allowed time between frame transmission and ACK/NAK (if requested), or between token sent and network activity. The delay is in increments of 64 times the period of the clock CLK. Thus, if CLK = 2 MHz, then TA may be set in increments of 32 microseconds (range of 32 μs to 8.2 ms).
TD	Network Dead Timer. Value of maximum time interval between received valid frames on the network. Same clock source and range as TA. TD must be >TA.
CBP (H,L)	Control Block Pointer. A sixteen bit pointer to the WD2840 control block in the user's memory. Must not be modified while the WD2840 is in network mode.
NAR	Next Address, Request. Used in conjunction with the NEWNA (CR10) control bit to cause the WD2840 to update the NA register. This redefines the node's successor in the network logical ring. It MUST be an address in the range 1-254. The acceptance of this update is confirmed when the NEWNA control bit is cleared. On the next token pass, if the redefined successor fails to accept the token, this WD2840 enters Scan mode where it sequentially attempts a token pass to successively higher nodes.
AHOLT	This register is set at a value indicating the number of access cycles (tokens received) that must be skipped before a data frame may be transmitted. (A token pass frame will be sent even if a data frame may not be sent at a given access cycle.) Initialized to zero at power up.
TXLT	This register is set at the maximum number of consecutive data frames the WD2840 may transmit during one access cycle. A value of zero allows the WD2840 to transmit all frames queued up to 256. Initialized to zero at power up.
MA	My Address. The WD2840 receives only frames with this destination address (along with the broadcast address) and inserts this address into the SA field of any transmitted frame. Must be set by the host (range is 1 to 254).

1.4 DIAGNOSTIC AIDS

There are three levels of diagnostics supported by the WD2840; those that are associated with the network as a whole, those associated with the individual node, and those that are limited to the WD2840 as a device. These tests are Network Diagnostics, System Diagnostics and Self Diagnostics respectively. The Network Diagnostics can be performed while the WD2840 is in the logical ring, but the System Diagnostics and the Self Diagnostics may be used only while the WD2840 is in the diagnostic mode.

DIAGNOSTIC MODE CONTROL			
CR00 ISOL	CR17 DIAGC	SR21 STATE	DEFINITION
1	0	0	WD2840 "Isolated." Power-up condition or isolate request.
0	0	0	WD2840 active.
1	0	1	Isolate request function confirmed.
1	1	0	Host request to enter diagnostic mode.
1	1	1	Diagnostic mode confirmed. Diagnostic functions of CR1 apply.
0	0	1	Illegal.
0	1	0	Illegal.
0	1	1	Illegal.

Diagnostic mode may be entered after power-up or from the network mode by manipulation of the mode control bits. The mode transition is confirmed by the WD2840 via the STATE status bit.

Once in diagnostic mode, the desired test is selected via CR1. Because most of registers 8 through F are interpreted differently for each test, only one of the diagnostic test bits should be set at a time. In conjunction with setting the diagnostic bits, the NUDIAG (CR10) bit must be set to perform the diagnostic test requested.

At the completion of the selected test NUDIAG is cleared by the WD2840. Therefore the host can initiate a diagnostic by entering the diagnostic mode, initializing the proper registers, setting the desired diagnostic bit, and setting NUDIAG. The host then moniters CR1 for NUDIAG going to zero, indicating the completion of the requested diagnostic.

DIAGNOSTIC STATE FLOW CHART

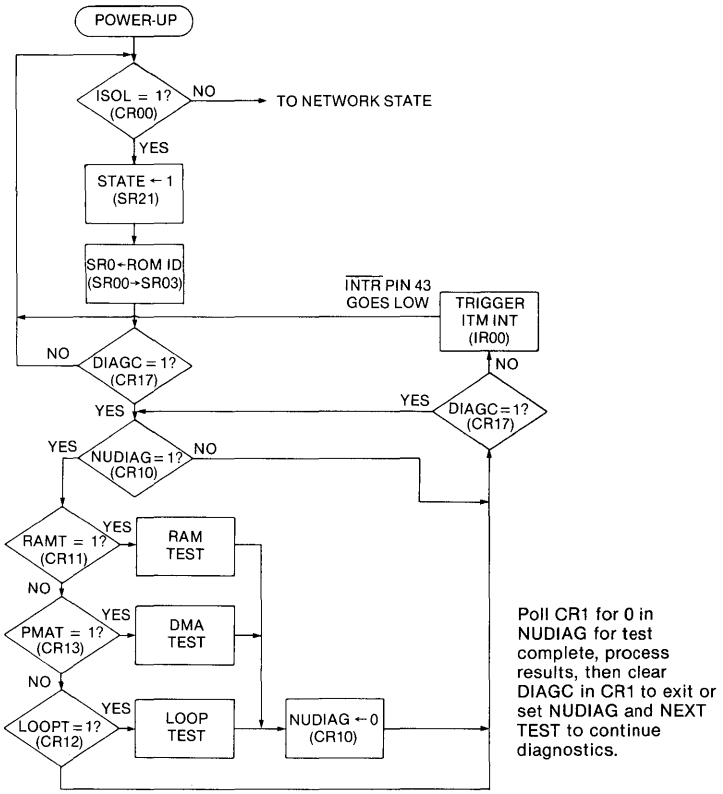

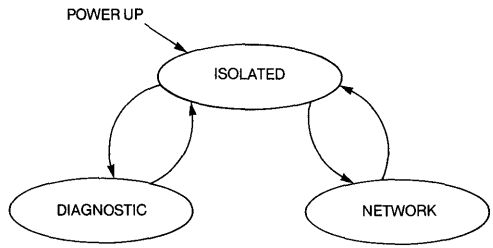

Figure 1.3 FUNCTIONAL STATES

1.4.1 SELF DIAGNOSTICS

Internal Ram and Interrupt Test

There are nine eight bit registers in the WD2840 which are not directly accessable by the users CPU. This test provides a means to check those registers and the interrupt register. The contents of register A are placed into the interrupt register and five even internal registers, and the contents of register B in four odd internal registers. The nine registers are then added together without carry and the result is placed in registers 2, 5, 6, 7.

Use the following procedure to initiate the RAM test:

1. Enter diagnostic mode.
2. Set up registers A and B
3. Set RAMT.
4. Set NUDIAG (can be set with RAMT bit together).
5. Wait for NUDIAG to be cleared.
6. Read registers 2, 5, 6, 7. Clear RAMT.

Note that the setting of any bit in the interrupt register while NOINT is clear will generate a hardware interrupt ($\overline{INTR}$, pin 43 goes true).

1.4.2 SYSTEM DIAGNOSTICS

DMA Test

This test verifies proper operation of the DMA subsystem by reading the value from a register and writing it into the user memory. The test continues by reading the value from the same location in memory and writing it into another register.

The value is read from register C. Using the transmitter DMA sub-system, it is written into memory location addressed by register A and B (location N; register A is the MSB). The receiver DMA sub-system is used and contents of the same address is read and it is stored into the register 7. Next the receiver dma is used and the contents from register D is written into location N + 1. The transmitter dma reads the value from location N + 1 and stores it into register 6.

It is the host's responsibility to check if the contents of registers C and register 7 and memory location N match. The same is true for registers D and 6 and memory location N + 1.

Loop-Back Test

The host can test the WD2840 parallel to serial, serial to parallel converters, CRC, and framing logic by setting LOOPT (CR12) bit of the WD2840, while in the diagnostic mode. The host has the responsibility of initializing a transmit buffer with a known pattern and then verifying its correct reception. The pattern is looped internally to the device if ILOOP (CR03) = 1, or may be looped externally (with outside logic) if ILOOP = 0.

The following procedure should be followed in order to run the loop-back test:

1. Enter diagnostic mode.
2. Set up register A and B to point to a buffer that is initialized with a pattern for transmission.
3. Set up register C and D to point to a buffer to receive the frame. (It is a good practice to initialize this buffer with all '00' or all 'FF' value bytes.)
4. Set up the buffer size in bits 3-0 of register E. (NOTE: In this test the last two bytes of the buffer will not be transmitted.)
5. Set ILOOP bit (CR03). (This is optional, if internal loop-back test is desired.)
6. Set LOOPT bit (CR12).
7. Set NUDIAG (CR10).
8. Wait for NUDIAG (CR10) to be cleared.
9. Compare the two buffers to verify correct reception of the frame.

NOTE:
If this test frame is allowed onto the network, transmission collisions may occur. Further, the first three bytes of the transmit buffers will be interpreted as TC, DA and SA, respectively, by the other stations. Therefore in case this test is initiated while this node is in the logical ring, care should be taken for choosing these three values for external loop-back test.

For proper operation of the internal loop-back test the $\overline{CTS}$ and SQ pins of the WD2840 should be either tied to ground or tied to $\overline{RTS}$ pin of the WD2840.

1.4.3 NETWORK DIAGNOSTICS

Duplicate Station Detection

Duplicate stations (more than one station with the same address) can result from the faulty programming of internal register MA (due to wrong address switch settings on the user's device, for example). This is expected to occur often enough to warrant the addition of a detection algorithm in the users WD2840 initialization procedure.

After initializing all required parameters, the user places the WD2840 in network mode (by setting ISOL false). The WD2840 monitors all frames on the network, and, if one is observed as having been transmitted by its address (source address of the frame equals the value in register MA) an event counter is incremented.

The user should monitor the SA = MA event counter at least long enough for the token to have circulated all the way around the access ring (time is configuration dependent) before enabling the WD2840's transmitter.

It is useful to note that this constraint requiring each node which is participating in the network logical ring to have a unique address does not extend to nodes which are "listening" but not "in the ring." It might be useful to a network designer to have groups of receive only nodes which have the same node address but do not participate in the network token passing (see GIRING - CR13). Data frames transmitted to such clusters must not request acknowledgement since all nodes in the cluster would simultaneously respond.

Copy Mode

The COPY Mode is selected by setting the COPY control bit (CR02). Normally the WD2840 receives (DMA's into the receive buffer chain) data frames only if they contain the general broadcast destination address or if they are specifically addressed to the WD2840. This occurs when the frame's destination address (DA) matches the WD2840 my address (MA, set by the host).

However, when COPY mode is selected data frames which are specifically addressed to other nodes will be treated as broadcast frames by this node. The COPY mode allows a specific node to "evesdrop" on data frame traffic on the network.

Nak Response

The WD2840 sends negative acknowledgements (NAK's) on response to received frames under several circumstances. The NAK prevents the transmitting node from wasting bandwidth retrying indiscriminately, and further, lends visibility to individual network node problems. The NAK includes a reason code which is available to the transmitter's software (via the TFSB).

Each data frame to be transmitted can be specifically marked (via the FCB) by the host to require an ACK/NAK response from the receiving WD2840. In the absence of errors, an acknowledge (ACK) frame will be returned to the transmitter as confirmation. However, several circumstances cause a Negative Acknowledge (NAK) to be returned:

1. Insufficient buffer space
2. Receiver not enabled (RXEN - CR05 cleared)
3. Receiver overrun
4. Frame exceeded 16 buffers in length

This information is placed in the transmitted frames's FSB. See section 2.1.2 for more details on the Transmit Frame Status Byte (TFSB).

2.0 INTERFACES

There are two interfaces to the WD2840: the host computer side, and the network side. The network side is conventional from an electrical point of view, the WD2840 performs all logical functions required to ensure communications capability on broadcast media (such as coax or RF).

The host interface involves two separate functional interfaces: the status/control registers described in section one, and a DMA interface that is described in the following subsection.

2.1 HOST

The WD2840 uses a complex memory buffer architecture allowing it to respond in real time to its network obligations (e.g., to meet network data rate and processing delay requirements). These memory structures are managed cooperatively by the host and the WD2840.

Memory management functions requiring real time response (e.g., traversing chains) are completely handled by the WD2840. Other important, but not time critical operations are the responsibility of the host software (such as removing used buffers from the transmit chain).

All memory references by the WD2840 are pointed to by memory locations (and internal registers) initially defined and set up by the host software. Initial values and memory based registers are grouped together and called the WD2840 Control Block.

The location of this control block is written into the registers CBPH and CBPL anytime the WD2840 is in Isolate State. This control block has the following structure:

CBP→	+0	NXTR(H)	Receive Buffer Chain (MSByte)
	+1	NXTR (L)	Receive Buffer Chain (LSByte)
	+2	NXTT (H)	Transmit Buffer Chain (MSByte)
	+3	NXTT (L)	Transmit Buffer Chain (LSByte)
	+4	BSIZE	Buffer Size / 16 (0-F = 64-1024 bytes)
	+5	EVT0	
	+6	EVT1	Eleven separate Event Counters, see section 2.1.1 for details
	....		
	+F	EVT10	

As the WD2840 transitions to Network State, it reads and uses the first five bytes of the control block. The remaining eleven bytes of event counters are accessed by the WD2840 only when each specific event condition occurs.

Either the Receive (NXTR) or Transmit (NXTT) chain entries in the control block may initially be zero; in such a case the WD2840 expects the chain to be extended by the host's changing the zero link field in the control block. Thereafter any such zero link would be in a buffer.

The WD2840 uses constant size buffers; their length is set by the value in location BSIZE. The buffer size is indicated by a 4-bit count in the least significant 4 bits of the BSIZE byte in the WD2840 control block. The buffer sizes available are multiples of 64; (BSIZE + 1) 64 is the buffer size used by the WD2840. Thus a BSIZE range of 0-15 corresponds to actual buffer sizes of 64 through 1024 bytes. This buffer length is inclusive of control bytes and buffer link pointers.

The WD2840 includes a chained-block feature which allows the user more efficient use of memory, particularly in situations where the maximum packet size is much larger than the average packet size. One or up to 16 buffers may make up a frame but a buffer may not contain more than one frame.

Byte counters are associated with each frame (at the memory interface, not actually transmitted within the frame) so that frames on the network need not be integer multiples of buffers. The byte counters include all buffer management overhead. Therefore, a frame consisting of 100 transmitted data bytes, occupying two 64-byte buffers, would have a byte count of 108 (six bytes per frame + 2 bytes per buffer boundry).

Since the WD2840 receive and transmit buffer chains are linked lists (see section 2.1.2 and 2.1.3) and are "followed" by the WD2840 but managed by the host; it is expected that the host will maintain both a FIRST and a LAST address for each chain. On transition into Network State, the chain origin information in the WD2840 control block is the same as FIRST. In fact, since the WD2840 does not change these control block entries, they can be maintained directly as FIRST by the host. An explicit LAST could be placed in an extended control block section.

The WD2840 "follows" the linked buffer chains by maintaining a NEXT address internally for each chain. This NEXT address can be in one of two states: 1) it can be the address of the next buffer in the chain, or 2) at the chain end (zero link), it can be the address of the buffer containing the zero link. The WD2840 uses a status bit for each chain, NXTR0 (receive) and NXTT0 (transmit), to differentiate the two states. When set they indicate the WD2840 chain NEXT address is in state 1 above; when clear they indicate state 2 above. This is an important distinction since it indicates whether the last buffer posted in a chain can be removed by the host (because the WD2840 has advanced to the buffer beyond) or must be left until the chain can be extended so the WD2840 can advance.

The host software monitors the progress of the NEXT pointer, and updates FIRST and LAST as it adds (and removes) buffers to (from) the chains as required. The WD2840 provides Interrupt Events (see IR0) and NXTR0, NXTT0 status bits to indicate when it advances along the two chains and exactly what state its NEXT address registers are in. The operation of these chains will be explained by example in later sections.

"Deadly Embrace" Prevention

A "Deadly Embrace" can occur when two processors reach a state where each is waiting for the other. In this case, the two processors are the user's CPU and the micro-controller inside the WD2840. Therefore, to prevent the "deadly embrace," the following rule is obeyed by the WD2840 and should also be obeyed by the user's CPU. This rule applies to the WD2840 memory registers and to the I/O registers. The Event Counters are an exception to this rule.

Rule:
If a bit is set by the CPU, it will not be set by the WD2840, and vice versa. If a bit is cleared by the WD2840, it will not be cleared by the CPU, and vice versa.

As an example, the NEWNA (CR10) control bit is only set by the host and is only cleared by the WD2840.

Dual DMA

The WD2840 may, for efficiency, interleave frame data fetch/store operations with fetches and stores of pointers and flags in memory. In all cases, operation sequencing is such as to prevent deadlocks and ambiguities between the WD2840 and software.

2.1.1 EVENT COUNTERS

Several non-fatal logical events are tabulated by the WD2840 and made visible to the host via memory based event counters (see WD2840 control block organization for specific locations). The WD2840 will increment each counter at the occurance of the specified event. Note that the WD2840 will not increment past 255. The host has the responsibility of initializing each counter.

COUNTER	DESCRIPTION
EVT0	"Set scan mode" frame received from the network. The NA register was redefined to MA + 1 at the time.
EVT1	Transmission error first attempt, second try successful. Can only occur for frames requiring an acknowledgement. It indicates no response was received for the first transmission; however, the second transmission was either ACK'ed or NAK'ed.
EVT2	Transmission error. Attempt aborted due to either transmitter underrun or frame length exceeding 16 buffers.
EVT3	Timer TD (network dead) expired.
EVT4	Access Control Frame Reception Error. A one or two byte supervisory frame (ACK/NAK, Token Pass, Scan Mode) has been received in error. This may be due to an FCS error, frame abort, or carrier loss detection.
EVT5	Data Frame Reception Error. An incoming data frame was incorrectly received due to an FCS error, frame abort, carrier loss detection, or receiving a data frame when expecting an ACK/NAK frame.
EVT6	NAK sent. Can occur for any of the following reasons: 1. Insufficient buffers in chain 2. Receiver not enabled (RXEN clear) 3. Receiver overrun 4. Frame length exceeded 16 buffers
EVT7	Invalid frame received. Caused by the detection of certain abnormal network conditions such as receiving an ACK/NAK frame when not expecting one, receiving a Scan mode frame when expecting an ACK/NAK frame, or receiving an invalid supervisory frame.
EVT8	Duplicate token detected. This counter will be incremented when the WD2840 determines that more than one token exists in the logical ring. This happens if a token pass is received when the WD2840 already has the token, or a data frame is received when the WD2840 is waiting for an acknowledgement frame.
EVT9	Not used.
EVT10	Duplicate node address. This counter will be incremented when a data frame being DMA'd into memory has a source address (SA) equal to the WD2840 node address (MA). This counter when used with COPY mode (CR02) is one way for detecting other nodes with the same node number (MA).

2.1.2 TRANSMIT MEMORY INTERFACE

When the token is received, data transmission is enabled (TXEN - CR06 and TXDEN - CR07 both set), and if the access hold-off counter has reached its limit, the WD2840 will determine whether any data frames are pending in the transmit chain. If so, it will transmit the first data frame in the chain. Otherwise the token will be passed. A given data frame will be the last frame transmitted for this token if any of several conditions occur:

1. ISOL (CR00) is set indicating the host has requested a transition to Isolate State.
2. TXDEN (CR07) is clear indicating the host has changed data frame transmission rights.
3. The frame FSB indicates this frame should be the last transmitted for this token.
4. The running frame counter has reached its limit (TXLT).
5. No further frames are pending in the transmit chain.

If any of the first four reasons above are true a token pass will occur. If the last frame does not require an acknowledgement, the WD2840 will piggyback the token pass if that is permitted (CR16). If the token cannot be piggybacked or if the last frame transmitted is the last frame pending (condition #5 above), an explicit token pass will occur. A piggyback token will not occur for the last pending frame because, for the general multiple buffer case, it is not known to be the last pending frame until after the transmission is complete.

The WD2840 will read and evaluate the address of the next frame at two specific points in time:

1. At the end of the prior frame, even if the prior frame is the last to be transmitted for this token.
2. When the token is received and data frame transmission is permitted.

If a non-zero frame address is found at time 1 above, it is kept and used without being re-read at time 2 above. However, if no pending frame is found at time 1, this is noted with the NXTT0 flag clear and the chain re-inspected on each occurrence of time 2 above.

As frame transmission commences, the WD2840 reads the address of the next buffer, the frame control byte, (FCB) and the frame length. It then starts reading bytes from the buffer and sending them until the frame length count or the end of the buffer is reached. The new buffer is read and data transmitted as before. (See Figure 2.1)

The frame length provided in the LENGTH field must include all overhead bytes (LINK, FSB, FCB, LENGTH) in all buffers used for the frame. For example LENGTH = 8 implies DA, SA but no data bytes. If buffer size is 64 then LENGTH = 67 implies DA, SA, and 57 data bytes (one data byte in an overflow buffer. As a result of this convention, certain LENGTH values are not valid (e.g., 65, 66 in the second example).

When the frame length is finally reached, the WD2840 pauses if an acknowledgement has been requested. The frame status byte (FSB) is updated when the frame is completed; its posting indicates frame completion and gives information about the success or failure of the frame transmission. At frame completion, the WD2840 attempts to advance along the transmission chain to identify the next frame regardless of whether it will be transmitted with this token or later.

The host may add frames to the end of the transmit chain at any time by changing the zero link in the last buffer. Also buffers of all posted frames up to but NOT including the last buffer of the most recently posted, may be arbitrarily removed from the chain. The last posted frame (more specifically, the last buffer of the last frame) may only be removed and reused if NXTT0 is set. This indicates that the WD2840 has advanced its NEXT address to the next frame but that its transmission has not been completed (in fact, perhaps not even started).

NOTE:
The WD2840 checks only the most significant byte of the link field for zero link detection. This has the following implications:

1. When writing into a zero link field, the host must write the LSB of the new link field first, followed by the corresponding MSB.
2. All buffers must have a starting address greater than or equal to Hex '0100'.

TRANSMIT FRAME STATUS BYTE (WRITTEN BY WD2840)

BIT #	7	6	5	4	3	2	1	0
Name	DONE	WIRING	X	X	SELF	VAL2	VAL1	VAL0

BIT	NAME	DESCRIPTION
7	DONE	Set to guarantee a non-zero value for the posted FSB.
6	WIRING	Value of the corresponding bit in received ACK frame.
5-4	—	Reserved.
3	SELF	When set, indicates the ACK/NAK code appears in the value field (bit 2-0) of this FSB is assigned by the WD2840 transmitter routine. When clear, indicates value resulted from ACK/NAK code from receiving station.
2-0	VAL	An encoded field whose interpretation depends upon the SELF flag (bit 3) in this FSB. a. SELF clear 0 0 0 — No receive error (= ACK when DONE is set). 0 0 1 — Insufficent buffers for frame. 0 1 0 — Receiver not enabled at frame start. 0 1 1 — Receiver over-run. 1 0 0 — Frame exceeded 16 receive buffers. b. SELF set 0 0 0 — No transmit error. 0 0 1 — Transmitter under-run. 0 1 0 — End of chain reached. 0 1 1 — Frame exceeded 16 transmit buffers. 1 0 0 — Transmission failed. Two attempts occurred.

Transmit Frame Status and Control Bytes

Each frame has two bytes reserved, one for host control information needed by the WD2840, the other for status information posted by the WD2840 at frame transmission completion. The frame control byte (FCB) is only read by the WD2840, never changed; the frame status byte (FSB), is written (posted) by the WD2840 with no regard for its prior contents. On completion, the FSB value will always be non-zero; it is important that the host zero the FSB byte in order to be able to recognize a posted frame.

NOTE:
Specifically note in Figure 2.1 that the first buffer of each frame has a different structure than any overflow buffers for that frame. In particular, each frame has only one set of FSB, FCB, and LENGTH fields regardless of the number of buffers required by the frame.

TRANSMIT FRAME CONTROL BYTE (WRITTEN BY HOST)

BIT #	7	6	5	4	3	2	1	0
Name	WACK	FCBLF	TRANSP	X	X	X	X	X

BIT	NAME	DESCRIPTION
7	WACK	Wait for Acknowledgement. Instructs the WD2840 to wait for an ACK/NAK response from the receiver for this particular frame only. The token control (TC) byte in the frame is automatically set to cause the destination node to respond. This bit must NOT be set if the frame uses the broadcast destination address. Inadvertently doing so will cause the frame to be posted "Transmission failed, due to max retries."
6	FCBLF	Last Frame. This bit will cause the WD2840 to pass the token either piggybacked with this frame (if possible) or explicitly after the frame transmission completes.
5	TRANSP	Transparent Frame. This bit will cause the WD2840 to interpret the buffer contents to be the exact sequence of bytes to be transmitted. The normal token control (TC) byte and source address (SA) byte generation is suppressed. Note that for a non-transparent data frame the TC byte must NOT appear in the buffer.
4-0	—	Reserved.

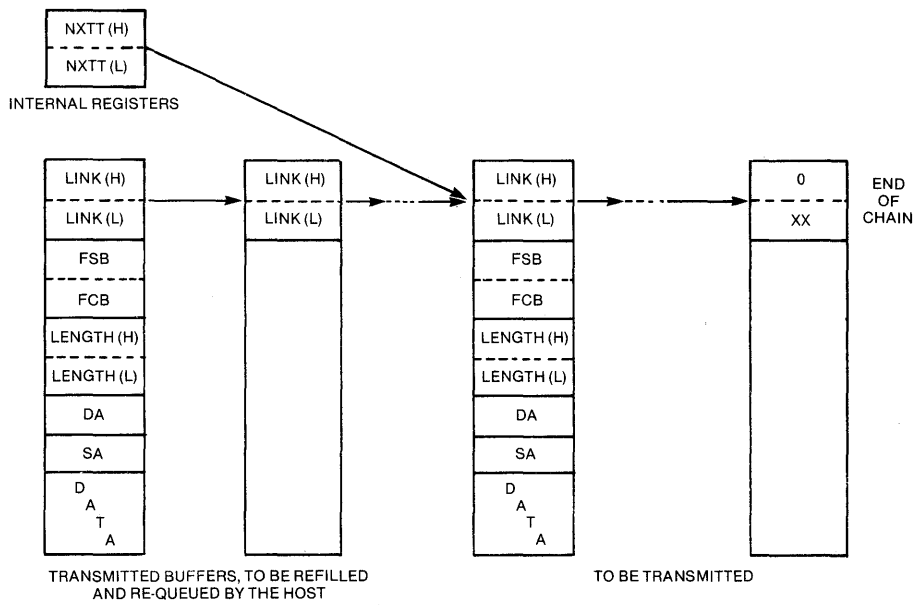

Figure 2.1 TRANSMIT BUFFER CHAIN

2.1.3 Receive Memory Interface

After the third byte of an incoming data frame is detected, the WD2840 will begin to place frame data into memory if several conditions are satisfied:

1. Receiver Enabled (RXEN-CR05 set).
2. There is an available buffer in the receive buffer chain.
3. The frame is addressed to this node specifically, it is a broadcast frame, or COPY mode has been selected by the host.

As the frame continues, it may completely fill its buffer. If this happens the WD2840 reads and inspects the link field of the current buffer. If this link is zero, an error occurs and the receive chain is reset to reuse from the first buffer used by the dropped frame. However, if another buffer is available, the incoming frame is continued beginning in the third byte of that buffer. This continues until one of several things happen:

1. Receiver overrun. The WD2840 has a four byte FIFO to buffer incoming frame data; however, if the host DMA responds too slowly a receiver overrun will occur. If this happens an event counter is incremented, the frame is dropped, and the receiver buffer chain is reset to reuse buffers of the dropped frame.
2. Current buffer capacity exhausted. If 16 buffers have been used for the current frame, an event occurs with the frame being dropped and the chain reset. Otherwise the WD2840 attempts to advance to the next buffer in the receiver buffer chain. The frame data will be continued in this subsequent buffer. If the end of the receiver buffer chain is reached an event counter is incremented, the frame is dropped, and the chain reset.
3. Frame ends. If the FCS is not corrected an event counter is incremented, the frame is dropped, and the chain is reset. If corrected however, the frame length is placed in the LENGTH field and the Frame Status Byte (FSB) is posted "done, no error."

If the frame is addressed to this node and indicates an acknowledgement is required (TC = 255), whether or not an error occurs, the WD2840 responds with an ACK/NAK supervisory frame indicating either success or failure. In case of receiver over-run, bad FCS, and SA = MA acknowledgement request will be ignored. (See section 1.4.3 for details)

It is the host's responsibility to ensure that buffers are available, initialized (FSB zero'ed), and attached to the end of the receive buffer chain.

RECEIVE FRAME STATUS BYTE (WRITTEN BY WD2840)

BIT #	7	6	5	4	3	2	1	0
Name	DONE	X	X	X	X	X	X	X

BIT	NAME	DESCRIPTION
7	DONE	Set to indicate the frame reception is complete.
6-0	—	Reserved.

RECEIVE FRAME CONTROL BYTE (WRITTEN BY HOST)

BIT #	7	6	5	4	3	2	1	0
Name	X	X	X	X	X	X	X	X

BIT	NAME	DESCRIPTION
7-0	—	Reserved.

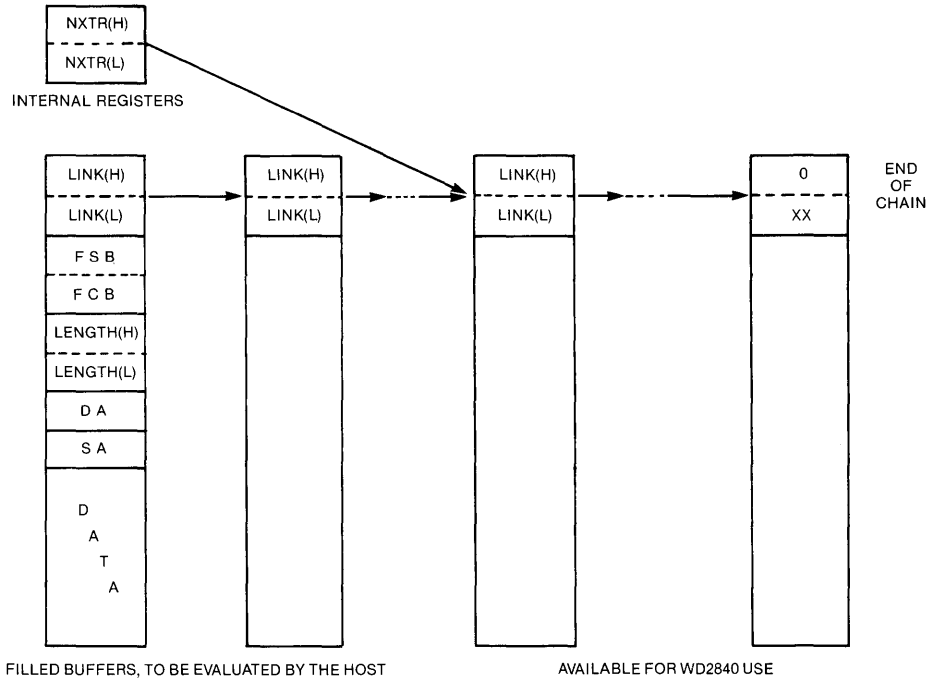

Figure 2.2 RECEIVE BUFFER CHAIN

2.2 MODEM INTERFACE

The modem interface is the conventional half duplex NRZ type with separate data and clock (Figure 2.3). When the WD2840 desires to transmit, it asserts $\overline{RTS}$ and awaits $\overline{CTS}$. $\overline{RTS}$ is generally used to enable the modem transmitter. After a system dependent preamble is generated, the modem asserts $\overline{CTS}$ which allows the WD2840 to begin the actual transmission of the frame. (Note: $\overline{CTS}$ may be asserted permanently if the transmission system does not need to generate a preamble).

The $\overline{SQ}$ input is used on receive to indicate a valid carrier. If this term is negated anytime during a receive message, the WD2840 will presume the message is in error and treat it as an abort. This signal is used to augment message integrity beyond that of the CRC by allowing a modem to detect and report low level faults (such as out-of-frequency carrier or missing clock).

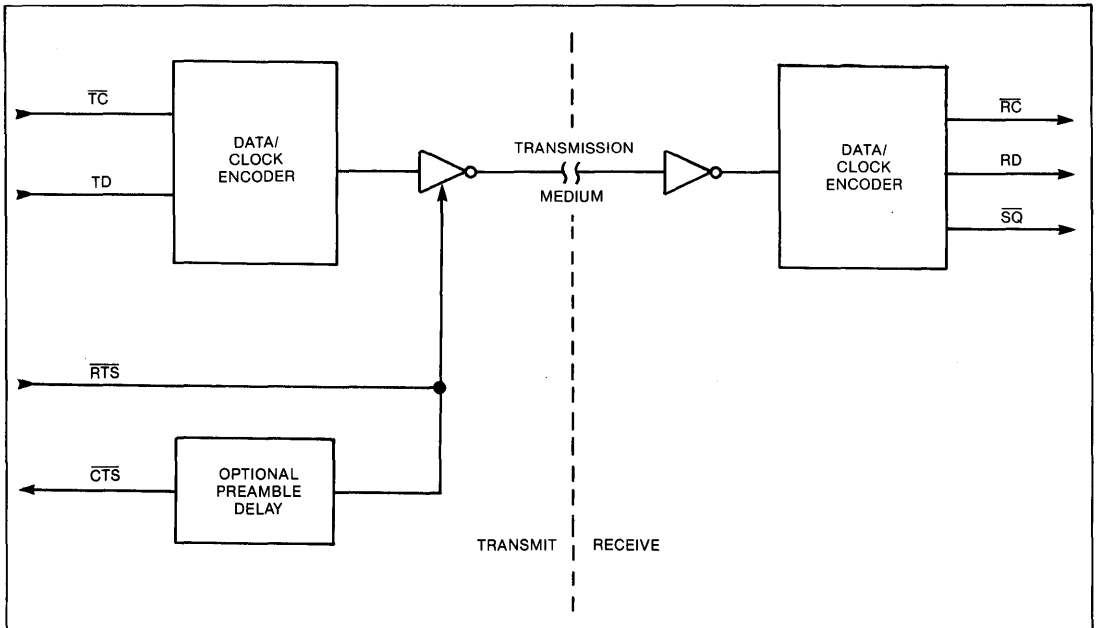

Figure 2.3 CONCEPTUAL MODEM

3.0 NETWORK PROTOCOL

To enable operation on a broadcast medium without the need for a central controller performing device polling, the WD2840 implements a media access protocol. The particular access protocol designed into the WD2840 prevents self-induced transmission collisions and ensures a fair and guaranteed distribution of transmission time among attached controllers.

This design-out of collisions allows the WD2840 a greatly expanded selection of transmission media, since no physical characteristics of a particular medium are relied upon for proper network operation. Another benefit of this lack of collisions is the visibility of network faults. If a collision is detected, it is treated consistently in a error recovery mode by the WD2840 and is also unambiguously visible to service personnel as a fault.

Secondly, the WD2840 can ensure that a transmitted message was correctly received and buffered by requiring acknowledgement of its receipt. This is sometimes called "acknowledging datagrams" where the sender awaits a predefined period after a frame is sent for a reply from its destination. With this method, no sequence counters nor multi-frame retransmission buffering is required. The scheme is efficient since local network applications such as the WD2840 address do not encounter extremely long transmission delays (such as satellite links) as in conventional data networks (such as X.25).

Both functions are parameterized, allowing tuning and optimization by the user to his unique application. These parameters may be adjusted in real time by the user's software, allowing a dynamic network, responsive to constantly changing requirements.

The two functions, access control and data transmission, function simultaneously though independently. Thus they are described separately as subprotocols for clarity.

3.1 Data Transmission

The data transmission cycle is entered after the token has been received and data transmission rights validated (see section 3.2 "access method"). The WD2840 determines if there is a frame to be sent and, if not, simply sends the token to the next station.

If something is queued for transmit, the WD2840 DMA's it from memory and sends it. After the complete frame has been sent, the WACK (Wait for ACK) bit is tested in the TFSB (Transmit Frame Status Byte). If set, the WD2840 waits for, and expects, an acknowledgement from the frames recipient. A timer (TA) is started. In the normal case, the ACK is received before TA expires which causes the WD2840 to send the next frame queued, repeating this procedure. Thus, the WD2840 sends multiple frames to various destinations until the transmit queue is emptied or a programmed limit (register TXLT) is exceeded.

In the event TA expires, the frame is re-transmitted once. (Note: it is the responsibility of higher level protocol operating in the host to protect against the possibility of duplicate frame reception.) If TA expires again, usually indicating the destination node is off-line, the FSB is updated to reflect the unsuccessful transmission, interrupt bit ITA is set, and the frame is skipped.

A frame is also skipped and tagged if the destination station sends a NAK, indicating it cannot presently process the frame.

TRANSMISSION OF ABORT

An ABORT is transmitted by the WD2840 to terminate a frame in such a manner that the receiving station will ignore the frame. An ABORT is sent when there is a Transmitter Under-Run. The abort sequence is a zero, followed by seven ones, after which RTS is set false.

3.2 ACCESS METHOD

The WD2840 network access method is based on the use of tokens, the specific granting of transmission rights passed from station to station. At any given time, exactly one station has the right to transmit (this right is called the token) and is obligated to pass it on when finished with it.

This can be clarified by referring to Figure 3.1. We assume in this figure that the network has already been initialized (meaning that the linkages in the access ring have already been established) and the token is held at this instant by station 4 (the station whose MA register = 4).

When station 4 is ready to pass his access right on, he sends a message to the station number called out in his internal register NA, in this case 11. The message, and thus the token, are received by station 11 who can now transmit its message(s). When station 11 is ready to pass the token, it sends a message to station 19, as directed by its internal register NA and the cycle continues, in a circular fashion, from station 4 to 11 to 19 to 54 to 4...

Notice that the station numbers need not be contiguous. This relatively arbitrary station numbering (in the example) poses no inefficiency to the access method. The value of this is the ability to add and remove stations (re-configure) to the network without re-arranging everyone elses addresses. (See section 3.2.2 for an example.)

In this way, the token is passed from one station to the next in a logical ring.

3.2.1 ACCESS INITIALIZATION/ ERROR RECOVERY

When the WD2840 is commanded into Network State, the Next Address Request (NAR) register and the NEWNA (CR10) flag must be used to define the Next Address (NA) register. When it is necessary to pass the token, it is passed to the current node number in register NA. If station NA is not on-line, determined by its lack of response, station NA + 1 is tried. This process continues until a station is found which does respond. The responding station number is written into register NA so that this scanning procedure need not be repeated on subsequent access cycles.

NOTE: 1. Node numbers 0 and 255 are reserved and cannot be used. Consequently scanning occurs circularly in the range 1-254.

2. During Scan mode token passing each node is only tried once.

Anytime a station cannot successfully pass a token within two attempts, register NA is updated to NA + 1, and a new "next" station is searched for. The result is the removal of non-responding station(s) from the access ring. An interrupt (INS) is generated indicating a network exception caused a change to NA.

The above description covers network recovery from station failure and purposeful removal of stations during on-line network operation. Setting stations in the scan mode can also be accomplished by sending control frames (a Scan frame redefines NA = MA + 1) over the network. The control frame may be directed to a single station, or all stations simultaneously (using the broadcast address). It is this scanning for new stations that permits on line addition to the access ring.

NOTE:
The policy of the SCAN frame is redefined by the user software as required by the application. For example: in a process control environment where stations are not often added while the network is in use, this procedure would be initiated rarely if at all.

3.2.2 REMOVING A STATION

There are two ways a station can be removed from the access ring: non-response due to station failure and non-response due to host commanded transition to the Isolate State. Both are treated identically from a network point of view.

Referring to Figure 3.1, assume that station 19 is removed from the network (either physically or logically). In this example, station 11 would detect a network fault when trying to pass the token to 19 (time TA would expire since station 19 will not respond). Station 11 detects this and finds the next station in the access ring by using the "scan" function (similar to initialization). The next attempt at passing the token would be to station 20, register NA + 1.

By starting the token ring recovery procedure at the intended station plus one (station 20) rather than MA + 1 (station 12) as is done in initialization, recovery delays are minimized (since fewer stations are tested for presence, 8 less in this example).

The next station found would be number 54 in the example which station 11 writes into his register NA (now "patching out" dead station 19). The next time station 11 is finished with the token, it directly sends it to 54, making the sequence now 11 to 54 to 4 to 11 to 54...

3.2.3 INTERACTION OF THE SUB-PROTOCOLS

After a station is given the token, it will send an information frame, a token frame, or a combination of both. It is this combination frame, referred to as a "piggy back" token, that causes the sub-protocols to interact slightly.

In the normal case (no time-out), the SOURCE may transmit a combination frame to the DATASINK when his access period is over. All stations on the network observe this; after the reception of the current frame is complete, the one whose MA register matches the token address in the frame (TC) knows it has the token.

In the case of a combination frame, the SENDer resets his timer TA on transmission complete and waits for the NA station to transmit something valid, to verify his reception of the piggy back token. If the timer expires, the sender sends an explicit token (the data from the combination frame is assumed to have been accepted) and enters the normal token sub-protocol.

The user is prevented from sending a combination frame and requesting an acknowledgement at the same time to prevent possible network state confusions under time-out conditions.

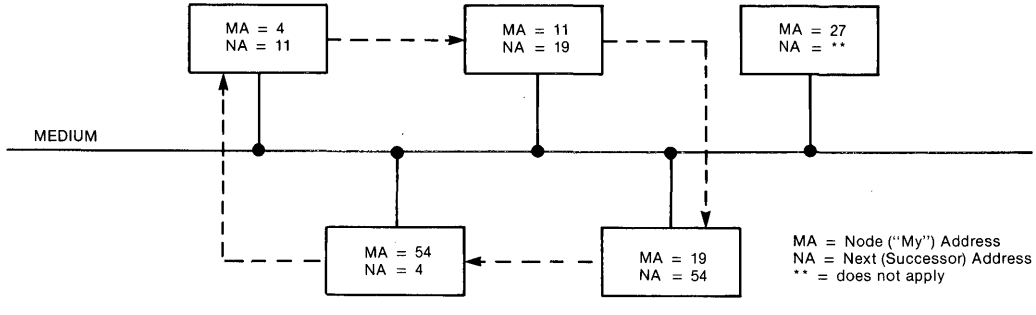

Figure 3.1 TOKEN PASSING ON A LOGICAL RING

3.3 FRAME FORMAT

The frame format the WD2840 uses to transmit all data and control frames is similar to the industry standard HDLC. A 16 bit CRC is implemented and standard zero insertion (CRC16-CCITT) is used for framing. This framing method allows the use of standard network monitoring and diagnostic equipment such as data scopes and logic analyzers.

Additional address fields and control points are defined as required to support the protocol.

Normal Frame Format:

F - TC - DA - SA - I - FCS - F

F = Flag, binary pattern 01111110
TC = Token Control (8 bit)
DA = Destination Address (8 bit)
SA = Source Address (8 bit)
I = Information Field (0 to 2048 bytes or 16 buffers, whichever is less).
FCS = Frame Check Sequence (16 bit)

Access Control Format:

F - DA - AC - FCS - F

F = Flag, binary pattern 01111110
DA = Destination Address (8 bit)
AC = Access Control Field (8 bit)
FCS = Frame Check Sequence (16 bit)

Token Pass Format:

F - TC - FCS - F

F = Flag, binary pattern 01111110
TC = Token Control (8 bit)
FCS = Frame Check Sequence (16 bit)

FIELD DESCRIPTIONS AND ENCODING

TC

The token control byte has the dual purpose of transferring access control between stations and conveying a request for immediate acknowledgement of the frame by its intended receiver.

There is no interaction between the TC field and the DA or SA fields. Thus the token may be transferred to one station and data sent to the same or a different station, with one single frame. The value entered into the TC field is determined by the WD2840 and does not appear in the buffer (except for transparent frames). See WD2840 state document for details.

TC Value	Meaning
0	Token not affected at this time.
1-254	After current frame, the token belongs to station TC. (The sending station has recovery responsibility).
255	Immediate ACK requested. Token not affected.

NOTE:
The sharing of this field prevents the passing of the token with data (piggy-back) and acknowledgement requests on the same frame. This combination is specifically disallowed because of its undesirable characteristics in network error situations.

DA

Destination address. Value of zero is reserved, 1 to 254 indicates the destination address of the frame. The value 255 is the global (or broadcast) address.

SA

Source address. The values of 0 and 255 are reserved. A value of 1 thru 254 is the address of the sender of the frame.

I

Information Field. User defines format and content.

FCS

Frame Check Sequence. The FCS calculation includes all data between the opening flag and the first bit of the FCS, except for 0's inserted for transparency. The sixteen bit FCS is compatible with the standard HDLC FCS.

AC

Access Control. Conveys supervisory information. May be sent as a command using transparent mode or received in response to an ACK/NAK request. Its format is shown below:

ACCESS CONTROL FIELD

BIT #	7	6	5	4	3	2 1 0
Name	SCANF	WIRING	0	0	0	NVAL2 — NVAL1 — NVAL0

BIT	NAME	DESCRIPTION
7	SCANF	Scan Mode (Command). Indicates that the addressed node(s) must redefine NA = MA + 1 for use on its next token pass.
6	WIRING	Wants in ring (Response). This bit indicates the node that transmitted the frame is not in the logical ring but would like to be. It is the logical function of the transmitting node's GIRING .AND. INRING. (see CR13 and SR20) The WD2840 does not act on this information but merely passes it to the host via the ACK'ed frame's FSB.
5-3	—	Reserved.
2-0	NVAL	An encoded NAK/ACK value (Response). The receiving node will set one of the following codes depending upon the state of the last received frame: 0 0 0 — No error 0 0 1 — Insufficient buffers for frame 0 1 0 — Receiver not enabled at frame start 0 1 1 — Receiver overrun 1 0 0 — Frame exceeded 16 receive buffers

4.0 ELECTRICAL SPECIFICATIONS

ABSOLUTE MAXIMUM RATINGS:

Voltages referenced to V_{SS}
High Supply Voltage (V_{DD}) . . . . . −0.3 to 15V
Voltage at any Pin . . . . . −.03 to 15V
Storage Temperature Range . . . . −55°C to +125°C
Electro-static voltage at any pin . . . . . 400V (Note 6)

NOTE:
Absolute maximum ratings indicate limits beyond which permanent damage may occur. Continuous operation at these limits is not intended and should be limited to those conditions specified under DC Electrical Characteristics.

OPERATING CHARACTERISTICS (DC):

Operating Temperature Range . . . . . 0°C to +70°C

SYMBOL	PARAMETER	MIN.	TYP.	MAX.	UNIT	CONDITIONS
I_{DD}	V_{DD} Supply Current		18	30	mA	
I_{CC}	V_{CC} Supply Current		160	220	mA	
V_{DD}	High Voltage Supply	11.4	12	12.6	V	
V_{CC}	Los Voltage Supply	4.75	5	5.25	V	
V_{IH}	Input High Voltage	2.4			V	
V_{IL}	Input Low Voltage			0.8	V	
V_{OH}	Output High Voltage	2.8			V	$I_O = -0.1mA$
V_{OL}	Output Low Voltage			0.4	V	$I_O = 1.6mA$
I_{OZH}	Three-State Leakage			50	μA	$V_{IN} = V_{CC}$
I_{OZL}	Three-State Leakage			50	μA	$V_{IN} = 0.4V$
I_{IH}	Input Current			10	μA	$V_{IN} = V_{CC}$
I_{IL}	Input Current			1.6	mA	$V_{IN} = 0.4V$

5.0 TIMING CHARACTERISTICS (AC):

SYMBOL	PARAMETER	MIN.	TYP.	MAX.	UNIT	CONDITIONS
CLK	Clock Frequency	0.5		2.05	MHz	Note 1
RC	Receive Clock Range	0			MHz	Note 4
TC	Transmit Clock Range	0			MHz	Note 4
MR	Master Reset Pulse Width	10			mS	
T_{AR}	Input Address Valid to $\overline{RE}$	0			nS	
T_{RD}	Read Strobe (or $\overline{DACK}$ Read) to Data Valid			375	nS	Note 5, 2
T_{HD}	Data Hold Time From Read to Strobe			80	nS	
T_{HA}	Address Hold Time From Read Strobe	80			nS	
T_{AW}	Input Address Valid to Trailing Edge of $\overline{WE}$	200			nS	
T_{WW}	Minimum $\overline{WE}$ Pulse Width	200			nS	
T_{DW}	Data Valid to Trailing Edge of $\overline{WE}$ or Trailing Edge of $\overline{DACK}$ for DMA Write	100			nS	Note 2, 3

TIMING CHARACTERISTICS (AC):

SYMBOL	PARAMETER	MIN.	TYP.	MAX.	UNIT	CONDITIONS
T_{AHW}	Address Hold Time After $\overline{WE}$	80			nS	
T_{DHW}	Data Hold Time After $\overline{WE}$ or After $\overline{DACK}$ for DMA Write	100			nS	
T_{DA1}	Time From $\overline{DRQO}$ (or $\overline{DRQI}$) to Output Address Valid if ADRV = 1			80	nS	
T_{DA0}	Time From $\overline{DACK}$ to Output Address Valid if ADRV = 0			400	nS	Note 5
T_{DD}	Time From Leading Edge of $\overline{DACK}$ to Trailing Edge of $\overline{DRQO}$ (or $\overline{DRQI}$)			400	nS	Note 5
T_{DAH}	Output Address Hold Time From $\overline{DACK}$			100	nS	
T_{DMW}	Data Hold Time From $\overline{DACK}$ For DMA Read			100	nS	Note 2

NOTES:

1. Clock must have 50% duty cycle.
2. There must not be a CPU read or write ($\overline{CS}$-$\overline{RE}$ or $\overline{CS}$-$\overline{WE}$) within 500 nanoseconds after the trailing (rising) edge of $\overline{DACK}$.
3. There must not be the leading (falling) edge of $\overline{DACK}$ allowed within 500 nanoseconds after the completion of a CPU write ($\overline{CS}$-$\overline{WE}$).
4. See "Ordering Information" for maximum serial rates.
5. C(load) = 100pf
6. Measured by discharging a 100pf capacitor to each pin through a 1K ohm resistor.

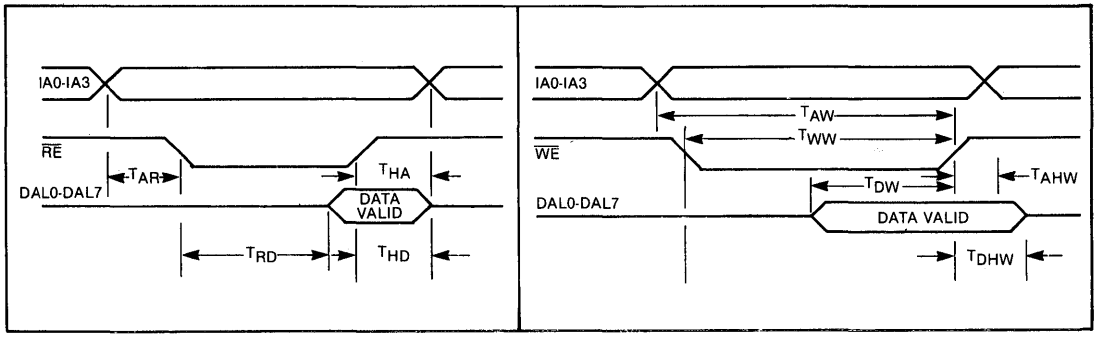

CPU READ ($\overline{CS}$ IS LOW) **CPU WRITE ($\overline{CS}$ IS LOW)**

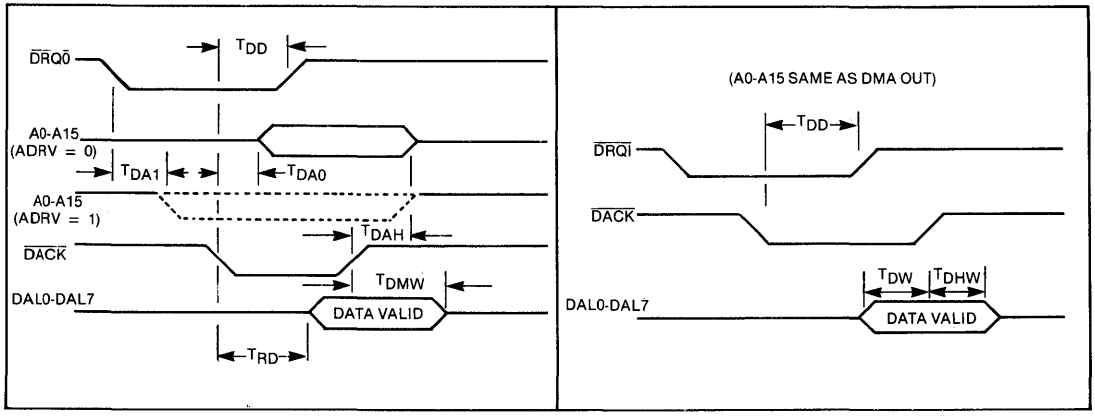

DMA OUT **DMA IN**

6.0 ORDERING INFORMATION

ORDER NUMBER	MAXIMUM RATE
WD2840T-01	100 Kbps
WD2840T-05	500 Kbps
WD2840T-11	1.1 Mbps

Printed in U.S.A.

PRELIMINARY

WDK25001 PACKIT

FEATURES

- X.25 (WD2501 OR WD2511) NETWORK CONTROLLER
- BAUD RATE GENERATOR
- RS-423 INTERFACE
- 8 SOCKETS FOR OPTIONAL USER RAM, ROM, EPROM
- WIRE-WRAP HOST INTERFACE AREA
- INTERNAL DATA RATE TO 64 KBPS
- +5°C TO +50°C OPERATION

GENERAL DESCRIPTION

The Western Digital WDK25001 PACKIT is a preconfigured breadboard designed to be an instant development tool for the implementation of bit-oriented, full duplex, serial X.25 data communication systems. Most of the hardware interfacing (Level 1/2 interfacing) is already done on the board. Level 2/3 interfacing is described in the long form specification. To perform Link Control (Level 2) functions, the PACKIT utilizes the Western Digital LSI Packet Network Interface Device. Two versions of the LSI Packet Network Interface Device are available for use with the PACKIT: the WD2501 and the WD2511. The WD2501 LSI device handles the LAP (Link Access Procedure). The WD2511 LSI device handles LAPB (Link Access Procedure Balanced).

The PACKIT is composed of two major physical areas. One side of the printed circuit board contains the factory installed circuitry required to perform X.25 Link Level operations: the other side of the circuit board contains the wire-wrapping area for user implementation of the individual system design.

The wire-wrapping area of the circuit board is designed to accommodate any industrial standard IC package. IC width is defined in 0.15 inch increments, with IC pin separation of 0.1 inch. The wire-wrap area is large enough to permit the installation of approximately one hundred 16-pin IC's.

ORGANIZATION

The WDK25001 "PACKIT", is composed of two major physical areas. The printed circuit area contains factory installed circuitry which handles the Packet Network Interface device, programmable bit rate generator, memory address decoder, and the EIA RS-423 interface. The wire-wrap area is set up to allow the user to design and implement their own host interface and or level 3 and higher modules.

INTERFACES

The PACKIT interfacing (Level 1/2) is divided into two classifications: the user interface and the serial communications interface.

USER INTERFACE

Voltage Levels:

TTL

Architecture:

Microprocessor interface oriented. 16 Address lines, 8 data lines, 10 control lines.

SERIAL COMMUNICATIONS INTERFACE

Electrical Characteristics:

RS-423

Mode:

Synchronous, Bit stuffing oriented, full-duplex, X.25 Asynchronous Response Mode (WD2501) or X.25 Asynchronous Balanced Mode (WD2511).

Clock Rates:

Independent transmit and receive clocks. Either internal or external clock source, selected by jumper strapping. Internal clock rates selectable by DIP switches or under user program control via the User Interface.

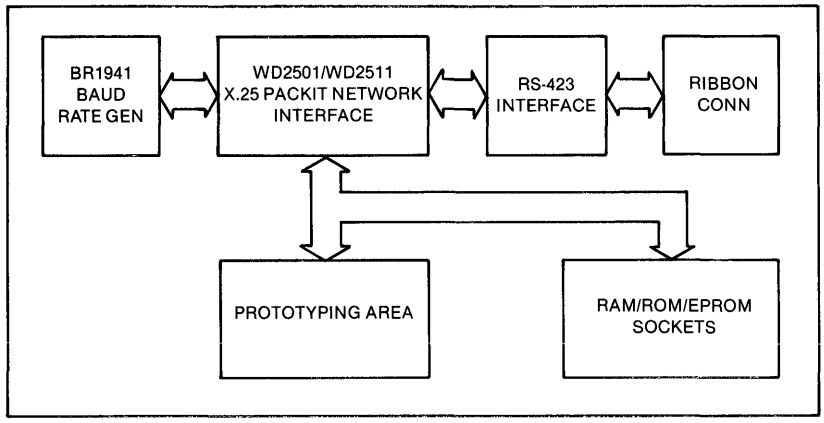

PACKIT BLOCK DIAGRAM

ELECTRICAL SPECIFICATIONS

Voltage Requirements:

+5v, +/− 5%
+12v, +/− 5%
−5v, +/− 5%

Maximum Voltage Ripple:

100mv, PTP

Factory-Installed Circuitry Power Requirements:

+5v @ .76A (max.)
(No memory device) +12v @ .20A (max.)
−5v @ .04A (max.)

MEMORY SPECIFICATIONS

Maximum Size Allowable:

64K x 8.

Memory Type:

8-bit wide RAM, ROM, EPROM

Package Type:

24-pin or 28-pin

ENVIRONMENTAL SPECIFICATIONS

Operating Temperatures:

+5 to +50 °C (Maximum temperature is reduced by 1.8 degree centigrade per 1000 meters altitude above sea level.)

Relative Humidity:

10 to 95%, maximum wet bulb of 32 degrees Centigrade and a minimum dew-point of 2 degrees Centigrade (90 degrees and 36 degrees Farenheit, respectively.)

PHYSICAL DIMENSIONS

Size:

Single printed circuit board, 8 inches by 14 inches (203.2 by 355.6mm)

Weight:

2 Pounds (0.907Kg).

THE LSI PACKIT NETWORK INTERFACE DEVICE

The Packit Interface Device is a 48-pin, n-channel silicon-gate, MOS chip designed to perform CCITT X.25 Level 2 (link control), with selected enhancements. In addition to the link control functions required by X.25, the Device eliminates the need for separate DMA circuits, timing chips, and the system software previously required to perform the link control between a data terminal equipment (DTE) and a data circuit-terminating equipment (DCE).

The WD2501/WD2511 interfaces directly with the on-board memory (when installed) or with the user's memory data and address lines. The WD2501/WD2511 is controlled through nine user interface lines: Master Reset ($\overline{MR}$), $\overline{REPLY}$, DMA Request Out ($\overline{DRQO}$), DMA Request In ($\overline{DRQI}$), DMA Acknowledge ($\overline{DACK}$), Interrupt Request ($\overline{INTR}$), Write Enable ($\overline{WE}$), Read Enable ($\overline{RE}$), and Chip Select ($\overline{CS}$).

Additionally, the WD2501/WD2511 is driven by the 1 MHz system clock.

THE PROGRAMMABLE BIT RATE GENERATOR

Circuit board IC location U3 contains a Western Digital BR1941 Programmable Bit Rate Clock. This n-channel MOS silicon-gate device is capable of generating sixteen externally selected clock rates, where-of six are standard rates.

The bit rate for both transmit and receive frequencies may be selected independently through the on-board

DIP switch, S1, or through the user's program via the TA, TB, TC, TD, RA, RB, RC, and RD user interface control lines.

To further increase PACKIT flexibility, the bit rate generator is installed in an IC socket to enable easy removal for specialized user bit rates not conforming to those available from the BR1941.

PACKIT MEMORY

Located on each PACKIT printed circuit board are eight 28-pin IC sockets designed for optional user-installed RAM, ROM, or EPROM memory chips. These sockets present the user with the option of either using the system memory or off-loading the PACKIT memory functions to on-board memory of the user's choice. Once the optional memory is installed it may be accessed by both the user's computer and by the WD2501/WD2511 Network Interface Device via DMA.

NOTE:
All memory devices installed on the PACKIT board must be 8-bit wide devices. Total memory installed must not exceed 64K-bytes.

The design of the memory sockets permits the user a wide range of RAM, ROM, and EPROM selectability. Memories such as the 4802 (2K x 8 RAM), 4118 (1K x 8 RAM), and the 2716 (2K x 8 EPROM), allow the user to select the type, capacity, and expense of memory desired for each individual application.

Also, the design of the sockets provides for the use of either 28-pin or 24-pin memory circuits. If a 24-pin circuit is to be installed, the chip is simply inserted in a manner that places pin 1 of the memory chip in IC socket position 3. Polarization of the sockets is indicated by a white dot, on the PC board, adjacent to both Pin 1 and Pin 3 of the socket.

When calculating the size of on-board memory, for a particular application, it is imperative that the user remember that the on-board memory must be large enough to contain five major arrays.

1. The data to be transmitted,
2. The data received,
3. The transmit-data look-up table (TLOOK),
4. The receive-data look-up table (RLOOK),
5. The PACKIT error counters.

For a detailed description of the use of memory by the Packit Network Interface Device, refer to the WD2501/WD2511 Long Form Specification.

Industrial standard memory types like 4118 or 4801 (1K x 8), 4802 or 2716 (2K x 8), and 37000 or 2764 (8K x 8) can be used in the PACKIT.

THE EIA RS-423 INTERFACE

The communications link between the Western Digital PACKIT and the Data Communications Equipment (DCE) is provided through the EIA RS-423 interface. This interface contains a subset of the RS449 signals. The interface signals are supplied to a 40-pin ribbon cable connector. Each of the RS-449 signals are buffered by 26LS29-type RS-423 line drivers and 26LS32 RS-423 line receivers.

If RS-232-C interfacing is desired, instead of RS-449 interfacing, the ribbon connector may be jumpered to provide RS-232-C signal compatability. However, due to differences in signal voltage levels, certain precautions must be taken when converting to RS-232-C interfacing.

BIT RATE GENERATOR — PROGRAMMING AND STRAPPING

The Western Digital BR1941 Bit Rate Generator, on board the PACKIT, may be programmed for receive and transmit bit rate, by on-board switch setting or by program command from the user circuitry. The selection between these two options is controlled by the S1 switchpack. To use program control all switches of S1 must be in the "OFF" position otherwise the bit rate is the "wired-and" of the program control and the setting of the switches. Transmit and receive bit rates are independently selected.

Program control of the bit rate, when enabled, is received by the PACKIT via the TA, TB, TC, and TD jumper pads, for the transmit bit rate, and via the RA, RB, RC, and RD jumper pads, for the receive bit rate. (Each of these control lines is tied to + 5v via a 2.2K resistor pull-up.)

THE PACKIT USER INTERFACE

The Western Digital PACKIT user interface defines all of the signals available to the user. These signals include: the microprocessor oriented address, data, and control lines, the RS-423 communications interface signals, the programmable bit rate generator control lines and on-board strapping.

USER INTERFACE SIGNALS

The user designed circuitry interfaces with the PACKIT through the interface signal jumper pads located at the center of the PACKIT PC board. The signals are listed in the table on following page. Signals are defined as "IN", "OUT", or "BI", to denote input from or output to the user interface, or bidirectionality, respectively. RS-449 mnemonics are presented in parenthesis.

USER INTERFACE SIGNALS

SIGNAL NAME	DIR	SIGNAL DEFINITION
A0-A15	BI	The sixteen memory address lines. These lines may be connected to the user's memory address bus or to the PACKIT on-board optional memory chips, when installed. These lines also carry the memory address outputs of the Packit Network Interface Device for DMA operations.
D0-D7	BI	The eight bit data lines used to transmit and receive byte-oriented data between the interface chip and the user circuitry. These lines also carry data between the WD2501/WD2511 and the optional on-board memory (when installed).
TMI (TM)	OUT	Test mode Indication. Signal returned from the RS-423 communications interface indicating that the local DCE is in the test mode. This signal is propagated to the user interface.
RDY REC (RR)	OUT	Ready to receive. Signal returned from the RS-423 communications interface indicating that the communications link is ready to receive data. This signal is propagated to the user interface.
$\overline{\text{DACK}}$	IN	DMA Acknowledge. The CPU signal generated in response to the WD2501/WD2511 transmitted DRQO or DRQI DMA request signals. An active low, sent to the PACKIT, on this line informs the PACKIT that the DMA request is acknowledged and the CPU has relinquished control of the system bus.
$\overline{\text{DRQI}}$	OUT	DMA Request IN. The WD2501/WD2511 requests a DMA bus access. A $\overline{\text{DRQI}}$ is a request for a transmission of data FROM the memory TO the WD2501/WD2511. This signal is active low.
$\overline{\text{DRQO}}$	OUT	DMA Request OUT. The WD2501/WD2511 initiated signal requesting access for a DMA data transfer. The $\overline{\text{DRQO}}$ signal requests a DMA cycle to enable transfer of data FROM the WD2501/WD2511 TO the memory. This signal is active low.
$\overline{\text{MR}}$	OUT	Master Reset. The master reset, generated by the on-board RESET momentary closure switch, clears all of the WD2501/WD2511 control and status registers, with the exception of two internal control bits (refer to the WD2501/WD2511 Long Form Data Sheet, sheet 2). This signal is active low.
$\overline{\text{INTR}}$	OUT	Interrupt Request. The WD2501/WD2511 issues $\overline{\text{INTR}}$ to request an interrupt. This signal is active low.
$\overline{\text{CS}}$	IN	Chip Select. $\overline{\text{CS}}$ is driven low, by the user's circuitry to enable the WD2501/WD2511 for programmed I/O read or write operations. $\overline{\text{CS}}$ may be permanently activated by jumpering pad location S11 to ground.
DATA MODE (DM)	OUT	The DATA MODE signal is returned from the communications link, over the RS-423 interface, to inform the PACKIT that the data link is in the data mode.
$\overline{\text{REPLY}}$	OUT	Reply. An active low signal, generated by the WD2501/WD2511 to indicate that it is selected ($\overline{\text{CS}}$ is low) and it is either read enabled ($\overline{\text{RE}}$ is low) or it is write enabled ($\overline{\text{WE}}$ is low).
$\overline{\text{MWE}}$	IN	An active low signal generated by the user's system to enable the PACKIT on-board memory chips for a memory-write operation. Not applicable on 24-pin memories. All 24-pin memories are enabled through the J2-0 through J2-7 interface signal lines. This signal is connected to pin 27 of all 28-pin memories.
J2-0 to J2-7	IN	User activated signal lines to individually enable each of the PACKIT on-board memory sockets. Each active low signal enables the Write Enable ($\overline{\text{WE}}$) input of the respective memory socket where: J2-0 write-enables memory socket U13 J2-1 write-enables memory socket U14 J2-2 write-enables memory socket U15 J2-3 write-enables memory socket U16 J2-4 write-enables memory socket U17 J2-5 write-enables memory socket U18 J2-6 write-enables memory socket U19 J2-7 write-enables memory socket U20 Each signal is connected to Pin 23 of the corresponding memory socket.

CONTINUED USER INTERFACE SIGNALS

SIGNAL NAME	DIR	SIGNAL DEFINITION
TA, TB, TC, TD	IN	These four inputs combined select the Transmit bit rate to be generated by the Western Digital BR1941 Programmable Bit Rate Clock. The values presented to the bit rate clock generator, over these lines, may be determined either by the user program or by on-board switch setting, as determined by switch-pack S1.
REMOTE LOOP TEST (RL)	IN	When activated, the REMOTE LOOP TEST line forces the communications link into a diagnostic test. Data is transferred from the memory, to the communications link, to the remote DCE, and back to the PACKIT for verification.
TERM RDY (TR)	IN	Terminal Ready. Input line to the PACKIT, from the user circuitry, informing the DCE that the PACKIT is ready to set-up the communications link. This signal is optionally generated, on a permanent basis, by the PACKIT when jumper S12 is connected to ground.
LOCAL LOOP TEST (LL)	IN	The LOCAL LOOP TEST performs a diagnostic operation similar to that of the REMOTE LOOP TEST with the exception that the data being tested is transmitted to the local DCE and then returned to the PACKIT for verification.
$\overline{\text{MOE}}$	IN	Memory Output Enable. A user supplied low active signal used to enable the 3-state output of the optional on-board PACKIT memory, when installed. This signal is connected to pin 22 (28-pin) and pin 20 (24-pin) of the memory devices.

WESTERN DIGITAL
CORPORATION

TR1402/TR1602
Universal Asynchronous Receiver/Transmitter (UART)

FEATURES

- FULL DUPLEX OR HALF DUPLEX OPERATION
- AUTOMATIC INTERNAL SYNCHRONIZATION OF DATA AND CLOCK
- AUTOMATIC START BIT GENERATION
- EXTERNALLY SELECTABLE
  Word Length
  Baud Rate
  Even/Odd Parity (Receiver/Verification — Transmitter/Generation)
  Parity Inhibit
  One, One and One-Half, or Two Stop Bit Generation (1½ at 5 Bit Level for TR1602)
- AUTOMATIC DATA RECEIVED/TRANSMITTED STATUS GENERATION
  Transmission Complete
  Buffer Register Transfer Complete
  Received Data Available
  Parity Error
  Framing Error
  Overrun Error
- BUFFERED RECEIVER AND TRANSMITTER REGISTERS
- THREE-STATE OUTPUTS
  Receiver Register Outputs
  Status Flags
- TTL COMPATIBLE
- PULL-UP RESISTORS ON ALL INPUTS

APPLICATIONS

- PERIPHERALS
- TERMINALS
- MINI COMPUTERS
- FACSIMILE TRANSMISSION
- MODEMS
- CONCENTRATORS
- ASYNCHRONOUS DATA MULTIPLEXERS
- CARD AND TAPE READERS
- PRINTERS
- DATA SETS
- CONTROLLERS
- KEYBOARD ENCODERS
- REMOTE DATA ACQUISITION SYSTEMS
- ASYNCHRONOUS DATA CASSETTES

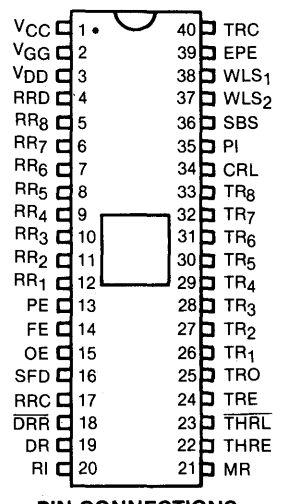

PIN CONNECTIONS

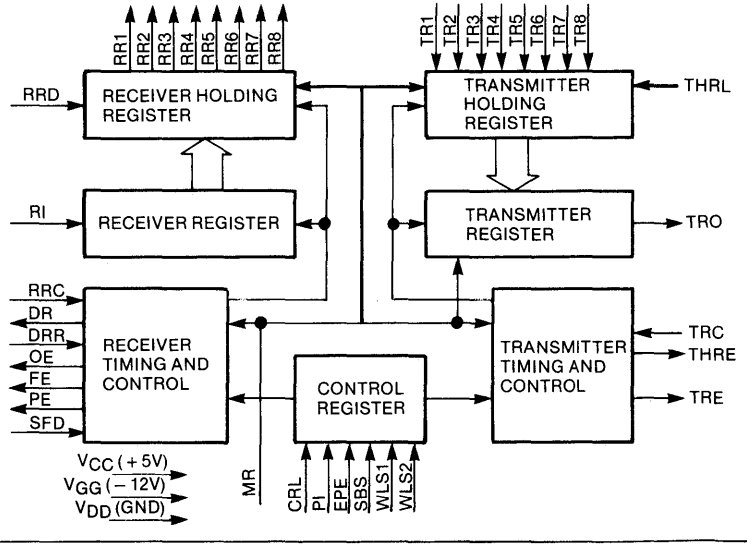

TR1602/TR1402 BLOCK DIAGRAM

WESTERN DIGITAL
CORPORATION

TR1863/TR1865
Universal Asynchronous Receiver/Transmitter (UART)

FEATURES

- SINGLE POWER SUPPLY — +5VDC
- D.C. TO 1 MHZ (64 KB) (STANDARD PART) TR1863/5
- FULL DUPLEX OR HALF DUPLEX OPERATION
- AUTOMATIC INTERNAL SYNCHRONIZATION OF DATA AND CLOCK
- AUTOMATIC START BIT GENERATION
- EXTERNALLY SELECTABLE
  Word Length
  Baud Rate
  Even/Odd Parity (Receiver/Verification — Transmitter/Generation)
  Parity Inhibit
  One, One and One-Half, or Two Stop Bit Generation (1½ at 5 Bit Level)
- AUTOMATIC DATA RECEIVED/TRANSMITTED STATUS GENERATION
  Transmission Complete
  Buffer Register Transfer Complete
  Received Data Available
  Parity Error
  Framing Error
  Overrun Error
- BUFFERED RECEIVER AND TRANSMITTER REGISTERS
- THREE-STATE OUTPUTS
  Receiver Register Outputs
  Status Flags
- TTL COMPATIBLE
- TR1865 HAS PULL-UP RESISTORS ON ALL INPUTS

APPLICATIONS

- PERIPHERALS
- TERMINALS
- MINI COMPUTERS
- FACSIMILE TRANSMISSION
- MODEMS
- CONCENTRATORS
- ASYNCHRONOUS DATA MULTIPLEXERS
- CARD AND TAPE READERS
- PRINTERS
- DATA SETS
- CONTROLLERS
- KEYBOARD ENCODERS
- REMOTE DATA ACQUISITION SYSTEMS
- ASYNCHRONOUS DATA CASSETTES

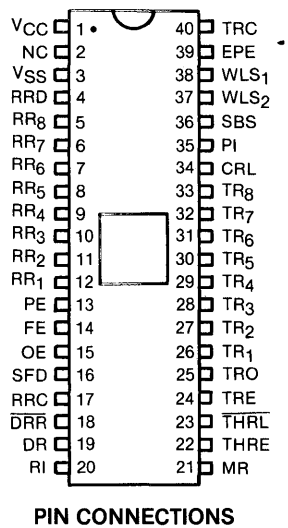

PIN CONNECTIONS

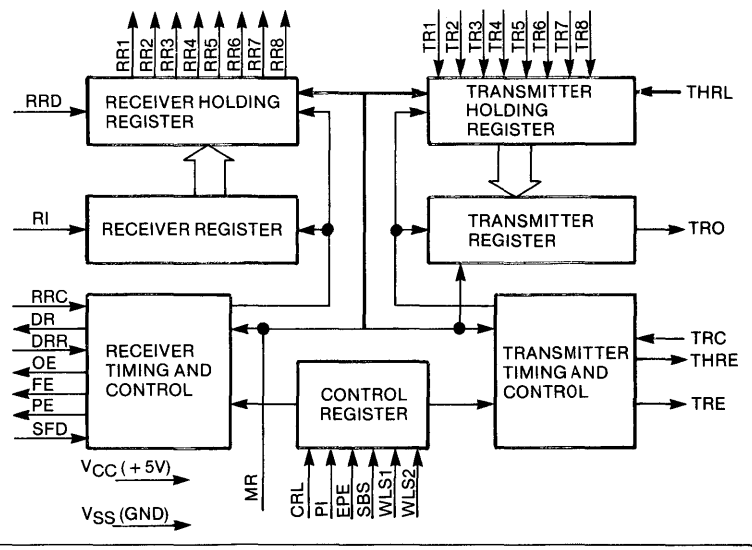

TR1863/TR1865 BLOCK DIAGRAM

WESTERN DIGITAL
CORPORATION

WD1983 (BOART)
Bus Oriented Asynchronous Receiver/Transmitter

FEATURES

ASYNCHRONOUS MODE

- FULL DUPLEX OPERATION
- SELECTABLE 5,6,7, & 8 BIT CHARACTERS
- LINE BREAK DETECTION AND GENERATION
- 1, 1½, or 2 STOP BIT SELECTION
- FALSE START BIT DETECTION
- OVERRUN AND FRAMING ERROR DETECTION
- DC TO 36K BITS/SEC (16X)
- DC TO 600K BITS/SEC (1X)
- 8251/8251A ASYNCHRONOUS ONLY REPLACEMENT
- REQUIRES NO ASYNCHRONOUS SYSTEM CLOCK
- 28 PIN PLASTIC OR CERAMIC
- +5 VOLT ONLY

SYSTEM COMPATIBILITY

- DOUBLE BUFFERING OF DATA
- 8 BIT BI-DIRECTIONAL BUS FOR DATA, STATUS, AND CONTROL WORDS
- ALL INPUTS AND OUTPUTS TTL COMPATIBLE
- CHIP SELECT, $\overline{RE}$, $\overline{WE}$, C/$\overline{D}$ INTERFACE TO CPU
- ON-LINE DIAGNOSTIC CAPABILITY
- THREE STATE DATA BUS

BAUD RATE-DC TO 36K BITS/SEC (16X) SELECTABLE CLOCK RATES

- 1X, 16X, 64X, BAUD RATE CLOCK INPUTS
- UP TO 47% DISTORTION ALLOWANCE WITH 64X CLOCK

APPLICATIONS

ASYNCHRONOUS COMMUNICATIONS
SERIAL/PARALLEL INTERFACE

GENERAL DESCRIPTION

The WD1983 is an N channel silicon gate MOS/LSI device that interfaces a digital asynchronous channel with a parallel channel. It is available in a ceramic or plastic standard 28 pin dual in line package.

The WD1983 is a fully programmable microprocessor I/O peripheral with two control registers and a status register. It is capable of full duplex operations.

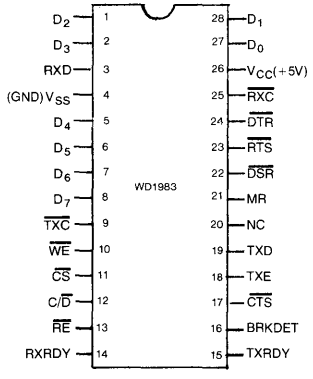

FIGURE 1 WD1983 PIN-OUT

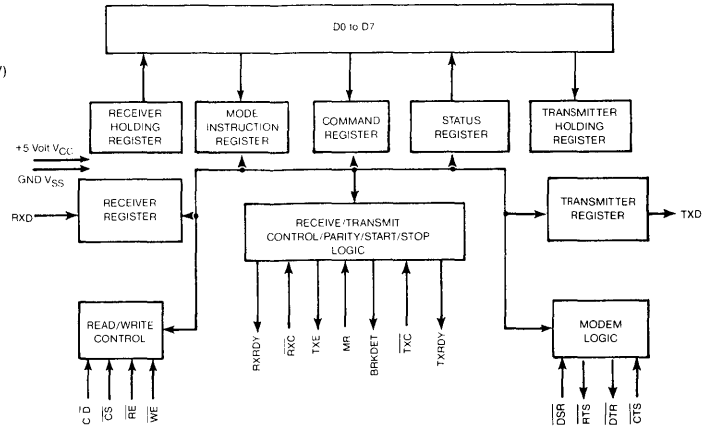

FIGURE 2 WD1983 BLOCK DIAGRAM

WD8250 Asynchronous Communications Element

FEATURES

- Designed to be Easily Interfaced to Most Popular Microprocessors (Z-80, 8080A, 6800, etc.)
- Full Double Buffering
- Independently Controlled Transmit, Receive, Line Status, and Data Set Interrupts
- Programmable Baud Rate Generator Allows Division of Any Input Clock by 1 to (2^{16} - 1) and Generates the Internal 16x Clock
- Independent Receiver Clock Input
- Fully Programmable Serial-Interface Characteristics
  - —5-, 6-, 7-, or 8-Bit Characters
  - —Even, Odd, or No-Parity Bit Generation and Detection
  - —1-, 1½-, or 2-Stop Bit Generation
  - —Baud Rate Generation (DC to 56K Baud)
- False Start Bit Detector
- Complete Status Reporting Capabilities
- THREE-STATE TTL Drive Capabilities for Bidirectional Data Bus and Control Bus
- Line Break Generation and Detection
- Internal Diagnostic Capabilities
  - —Loopback Controls for Communications Link Fault Isolation
  - —Break, Parity, Overrun, Framing Error Simulation
- Full Prioritized Interrupt System Controls
- Single +5-Volt Power Supply

GENERAL DESCRIPTION

The WD8250 is a programmable Asynchronous Communication Element (ACE) in a 40-pin package. The device is fabricated in N/MOS silicon gate technology.

The ACE is a software-oriented device using a three-state 8-bit bi-directional data bus.

The ACE is used to convert parallel data to a serial format on the transmit side, and convert serial data to parallel on the receiver side. The serial format, in order of transmission and reception, is a start bit, followed by five to eight data bits, a parity bit (if programmed) and one, one and one half (five bit format only) or two stop bits. The maximum recommended data rate is 56K baud.

Internal registers enable the user to program various types of interrupts, modem controls, and character formats. The user can read the status of the ACE at any time monitoring word conditions, interrupts and modem status.

An additional feature of the ACE is a programmable baud rate generator that is capable of dividing an internal XTAL or TTL signal clock by a division of 1 to $2^{16} - 1$.

The ACE is designed to work in either a polling or interrupt driven system, which is programmable by users software controlling an internal register.

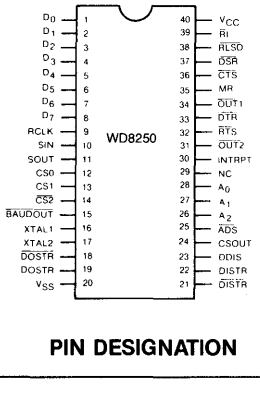

PIN DESIGNATION

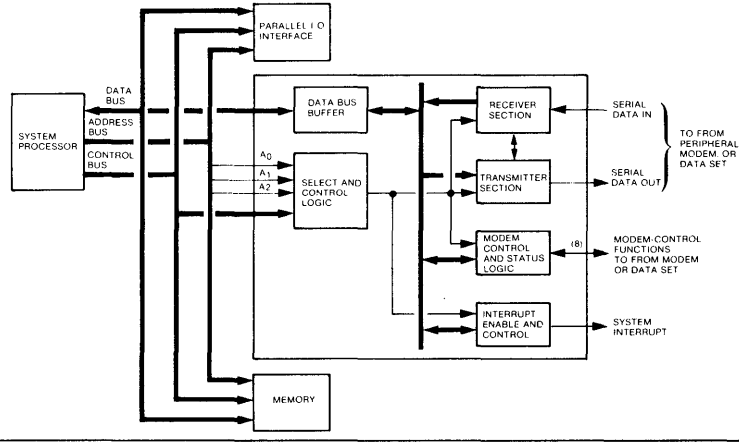

WD8250 GENERAL SYSTEM CONFIGURATION

WESTERN DIGITAL
CORPORATION

WD2123 DEUCE
Dual Enhanced Universal Communications Element

FEATURES

- TWO INDEPENDENT ASYNCHRONOUS FULL DUPLEX DATA COMMUNICATION CHANNELS (2 BOARTS)
- TWO INDEPENDENT BAUD RATE GENERATORS (ONE PER CHANNEL)
- EACH CHANNEL WITH FOLLOWING FEATURES:
- SELECTABLE 5 TO 8 BIT CHARACTERS
- 1X, 16X, 64X CLOCK RATES
- 16 SELECTABLE BAUD RATE CLOCK FREQUENCIES (INTERNAL)
- LINE BREAK DETECTION AND GENERATION
- 1, 1½, OR 2 STOP BIT SELECTION
- FALSE START BIT DETECTION
- ODD OR EVEN PARITY GENERATE AND DETECTION
- OVERRUN AND FRAMING DETECTION
- DOUBLE BUFFERING OF DATA
- TTL COMPATIBLE INPUTS AND OUTPUTS
- COMPATIBLE WITH 8251A (ASYNC ONLY) AND WD1983 DEVICES
- DIAGNOSTIC LOCAL LOOP-BACK MODE
- RXD INITIALIZATION UPON MASTER RESET
- ON-BOARD OSCILLATOR FOR EASE OF USE WITH A CRYSTAL
- VERSATILE CLOCK SELECT OPTIONS FOR INDEPENDENT TRANSMIT AND RECEIVE RATES

INTRODUCTION

The Western Digital WD2123 Dual Enhanced Universal Communications Element (DEUCE) is a single chip MOS/LSI Data Communications Controller Circuit that contains two independent full-duplex asynchronous RECEIVER/TRANSMITTER CHANNELS and two independent BAUD RATE GENERATORS. The WD2123 is fabricated in N-Channel silicon gate technology and is packaged in a 40 pin plastic or ceramic package. All inputs and outputs are TTL compatible.

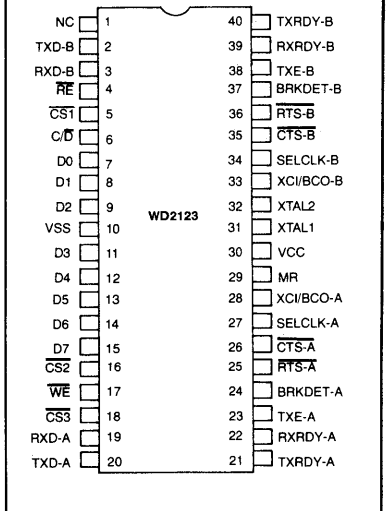

Figure 1. WD2123 PINOUT DIAGRAM

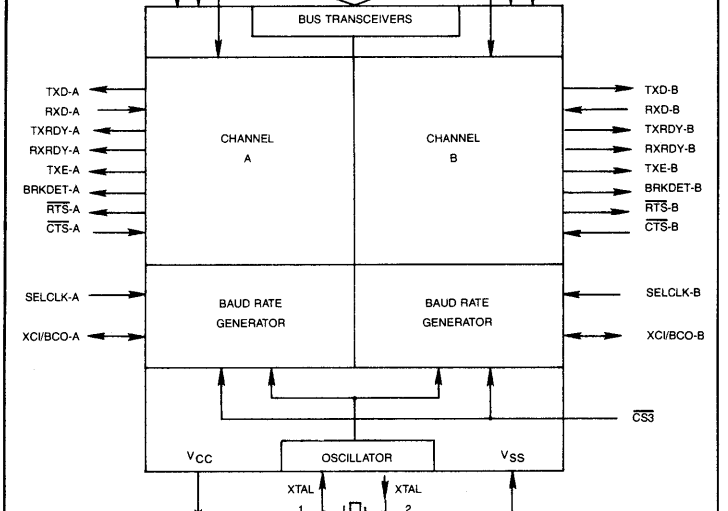

Figure 2. WD2123 BLOCK DIAGRAM

BR1941(5016) Dual Baud Rate Clock

FEATURES

- 16 SELECTABLE BAUD RATE CLOCK FREQUENCIES
- SELECTABLE 1X, 16X OR 32X CLOCK OUTPUTS FOR FULL DUPLEX OPERATIONS
- OPERATES WITH CRYSTAL OSCILLATOR OR EXTERNALLY GENERATED FREQUENCY INPUT
- ROM MASKABLE FOR NON-STANDARD FREQUENCY SELECTIONS
- INTERFACES EASILY WITH MICROCOMPUTERS
- OUTPUTS A 50% DUTY CYCLE CLOCK WITH 0.01% ACCURACY
- 6 DIFFERENT FREQUENCY/DIVISOR PAIRS AVAILABLE
- TTL, MOS COMPATIBILITY
- PIN COMPATIBLE WITH COM5016

GENERAL DESCRIPTION

The BR1941 is a combination Baud Rate Clock Generator and Programmable Divider. It is manufactured in N-channel MOS using silicon gate technology. This device is capable of generating 16 externally selected clock rates whose frequency is determined by either a single crystal or an externally generated input clock. The BR1941 is a programmable counter capable of generating a division from 2 to $(2^{15} - 1)$.

The BR1941 is available programmed with the most used frequencies in data communication. Each frequency is selectable by strobing or hard wiring each of the two sets of four Rate Select inputs. Other frequencies/division rates can be generated by reprogramming the internal ROM coding through a MOS mask change. Additionally, further clock division may be accomplished through cascading of devices. The frequency output is fed into the XTAL/EXT input on a subsequent device.

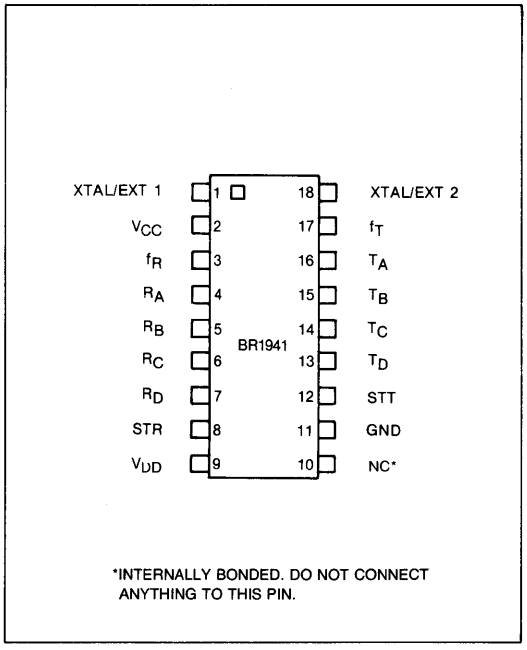

PIN CONNECTIONS

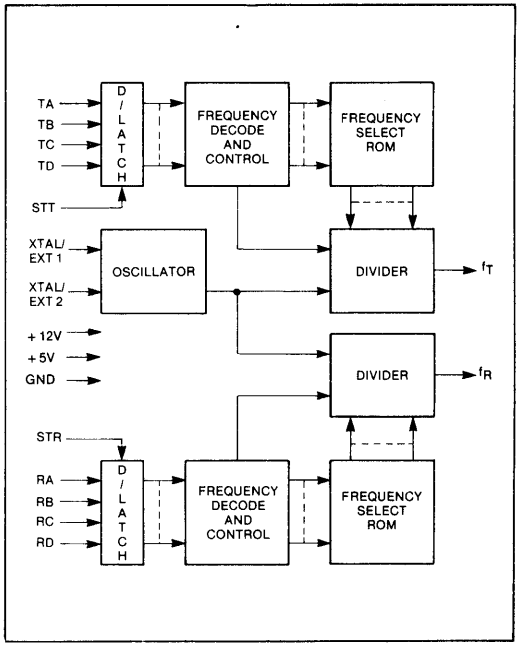

BR1941 BLOCK DIAGRAM

WD1943(8116)/WD1945(8136) Dual Baud Rate Clock

FEATURES

- 16 SELECTABLE BAUD RATE CLOCK FREQUENCIES
- OPERATES WITH CRYSTAL OSCILLATOR OR EXTERNALLY GENERATED FREQUENCY INPUT
- ROM MASKABLE FOR NON-STANDARD FREQUENCY SELECTIONS
- INTERFACES EASILY WITH MICROCOMPUTERS
- OUTPUTS A 50% DUTY CYCLE CLOCK WITH 0.01% ACCURACY
- 6 DIFFERENT FREQUENCY/DIVISOR PAIRS AVAILABLE
- SINGLE +5V POWER SUPPLY
- COMPATIBLE WITH BR1941
- TTL, MOS COMPATIBILITY

- WD1943 IS PIN COMPATIBLE TO THE COM8116
- WD1945 IS PIN COMPATIBLE TO THE COM8136 AND COM5036 (PIN 9 ON WD1945 IS A NO CONNECT)

GENERAL DESCRIPTION

The WD1943/45 is an enhanced version of the BR1941 Dual Baud Rate Clock. The WD1943/45 is a combination Baud Rate Clock Generator and Programmable Divider. It is manufactured in N-channel MOS using silicon gate technology. This device is capable of generating 16 externally selected clock rates whose frequency is determined by either a single crystal or an externally generated input clock. The WD1943/45 is a programmable counter capable of generating a division by any integer from 4 to $2^{15} - 1$, inclusive.

The WD1943/45 is available programmed with the most used frequencies in data communication. Each frequency is selectable by strobing or hard wiring each of the two sets of four Rate Select inputs. Other frequencies/division rates can be generated by reprogramming the internal ROM coding through a MOS mask change. Additionally, further clock division may be accomplished through cascading of devices. The frequency output is fed into the XTAL/EXT input on a subsequent device.

The WD1943/45 can be driven by an external crystal or by TTL logic.

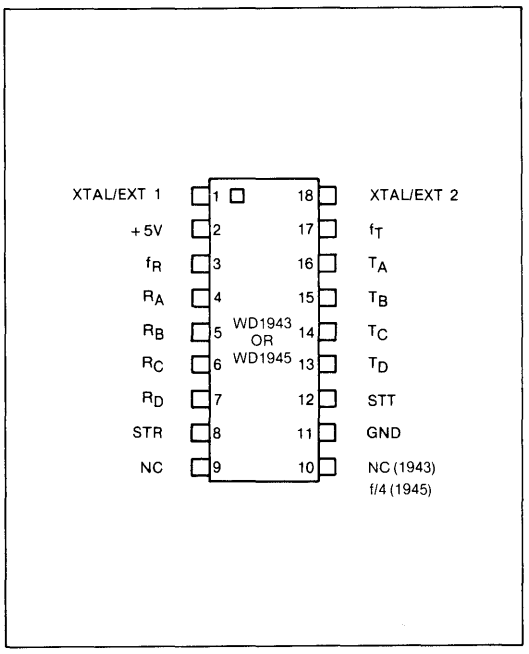

PIN CONNECTIONS

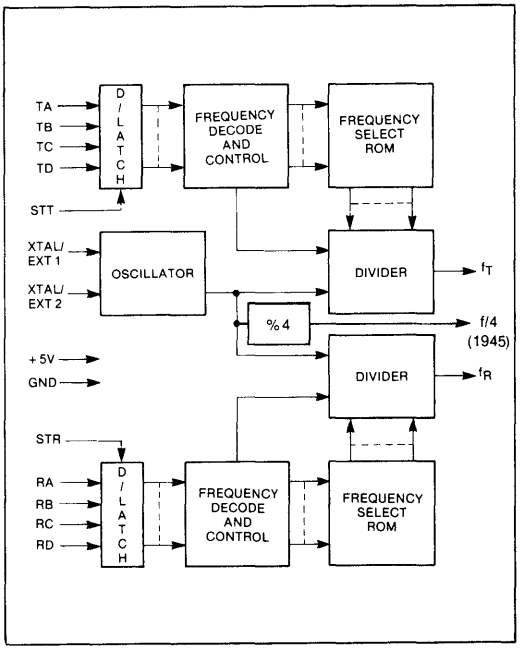

BLOCK DIAGRAM

UC1671 ASTRO

FEATURES

SYNCHRONOUS AND ASYNCHRONOUS
- Full Duplex Operations

SYNCHRONOUS MODE
- Selectable 5-8 Bit Characters
- Two Successive SYN Characters Sets Synchronization
- Programmable SYN and DLE Character Stripping
- Programmable SYN and DLE-SYN Fill

ASYNCHRONOUS MODE
- Selectable 5-8 Bit Characters
- Line Break Detection and Generation
- 1-, 1½-, or 2-Stop Bit Selection
- False Start Bit Detection Automatic Serial Echo Mode

SYSTEM COMPATIBILITY
- Double Buffering of Data
- 8-Bit Bi-Directional Bus For Data, Status, and Control Words
- All Inputs and Outputs TTL Compatible
- Up to 32 ASTROS Can Be Addressed On Bus
- On-Line Diagnostic Capability

TRANSMISSION ERROR DETECTION-PARITY
- Overrun and Framing

BAUD RATE — DC TO 1M BIT/SEC

8 SELECTABLE CLOCK RATES
- Accepts 1X Clock and Up to 4 Different 32X Baud Rate Clock Inputs
- Up to 47% Distortion Allowance with 32X Clock

APPLICATIONS

SYNCHRONOUS COMMUNICATIONS
ASYNCHRONOUS COMMUNICATIONS
SERIAL/PARALLEL COMMUNICATIONS

GENERAL DESCRIPTION

The UC1671 (ASTRO) is a MOS/LSI device which performs the functions of interfacing a serial data communication channel to a parallel digital system. The device is capable of full duplex communications (receiving and transmitting) with synchronous or asynchronous systems. The ASTRO is designed to operate on a multiplexed bus with other bus-oriented devices. Its operation is programmed by a processor or controller via the bus and all parallel data transfers with these machines are accomplished over the bus lines.

The ASTRO is fabricated in n-channel silicon gate MOS technology and is TTL compatible on all inputs and outputs.

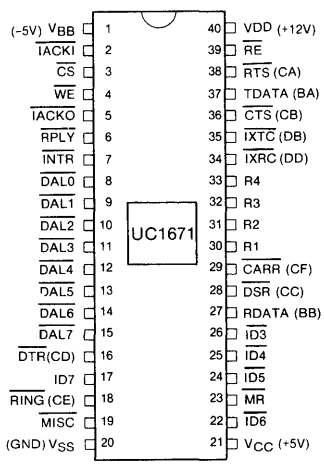

PIN CONNECTIONS

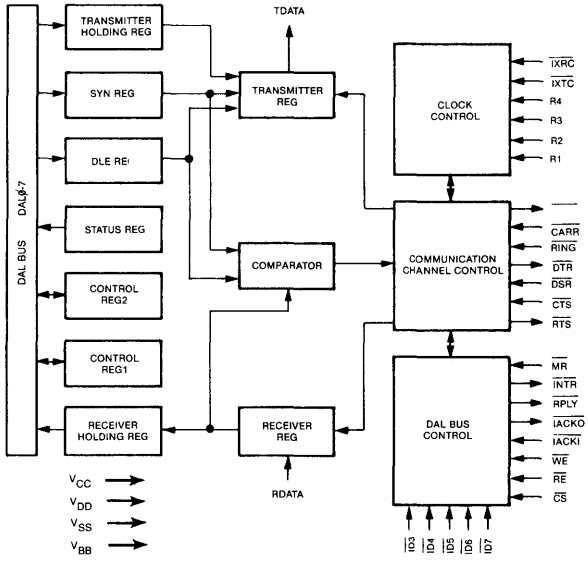

UC1671 BLOCK DIAGRAM

WD1931 Asynchronous/Synchronous Receiver/Transmitter

FEATURES

SYNCHRONOUS AND ASYNCHRONOUS
- Full Duplex Operations
- Selectable Character Length (5, 6, 7 or 8 Bits)

SYNCHRONOUS MODE
- Two Successive SYN Characters Sets Synchronization
- Programmable SYN and DLE Character Detection and Stripping
- Programmable SYN and DLE-SYN Fill
- Transparent BI-SYNC Operation
- DDCMP Compatible

ASYNCHRONOUS MODE
- Line Break Detection and Generation
- 1-, 1½-, or 2-Stop Bit Selection
- False Start Bit Detection
- Automatic Serial Echo Mode
- Overrun and Framing Error Detection

SYSTEM COMPATIBILITY
- Double Buffering of Data
- 8-Bit Bi-Directional Bus for Data, Status, and Control Words
- All Inputs and Outputs TTL Compatible
- Chip Select, RE, WE, A0, A1 Interface to CPU
- On-Line Diagnostic Capability
- Data Set, Carrier Detect, and Ring Interrupts

BAUD RATE — DC TO 1M BIT/SEC

8 SELECTABLE CLOCK RATES
- Accepts 1X Clock and Up to Four Different 32X Baud Rate Clock Inputs
- Up to 47% Distortion Allowance with 32X Clock

PINOUT COMPATIBLE TO WD193X FOR MULTIPROTOCOL BOARD APPLICATIONS

APPLICATIONS

SYNCHRONOUS COMMUNICATIONS
ASYNCHRONOUS COMMUNICATIONS
SERIAL/PARALLEL COMMUNICATIONS

GENERAL DESCRIPTIONS

The WD1931 is a MOS/LSI device which performs the functions of interfacing a serial data communications channel to a parallel digital system. This device is capable of full duplex communications with asynchronous and/or synchronous systems. Western Digital has made device pin assignments for the WD1931 to make it compatible with the WD193X (Synchronous Data Link Controller). This pin out allows the user to implement a one-board multiprotocol design. For character-oriented asynchronous and/or synchronous (bi-sync) protocols, the WD1931 is used, and for bit-oriented SDLC, HDLC and ADCCP protocols the WD193X is used (see WD193X data sheets and WD1931/WD193X compatibility application notes).

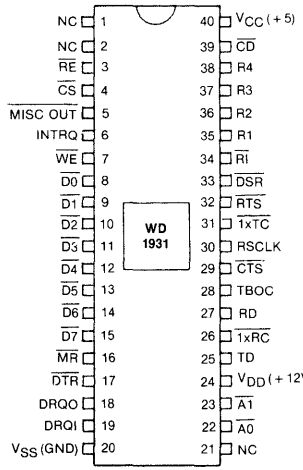

PIN CONNECTIONS

Figure 1. WD1931 PIN CONNECTIONS

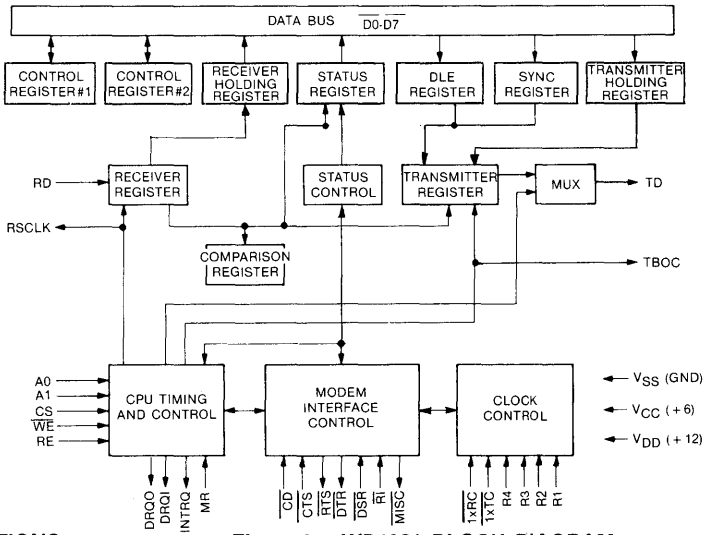

Figure 2. WD1931 BLOCK DIAGRAM

WESTERN DIGITAL
CORPORATION

WD1933/WD1935
Synchronous Data Link Controller

FEATURES

- HDLC, SDLC, ADCCP AND CCITT X.25 COMPATIBLE
- SDLC LOOP DATA LINK CAPABILITY
- FULL OR HALF DUPLEX OPERATION
- DC TO 2.0 MBITS/SEC DATA RATE
- PROGRAMMABLE/AUTOMATIC FCS (CRC) GENERATION AND CHECKING
- PROGRAMMABLE NRZI ENCODE/DECODE
- FULL SET OF MODEM CONTROL SIGNALS
- DIGITAL PHASE LOCKED LOOP
- FULLY COMPATIBLE WITH MOST CPU'S
- MINIMUM CPU OVERHEAD
- ASYNCHRONOUS/SYNCHRONOUS MULTI-PROTOCOL BOARD CAPABILITY (PIN COMPATIBLE WITH WD 1931)
- FULLY TTL COMPATIBLE
- SINGLE +5V SUPPLY
- ERROR DETECTION: CRC, UNDERRUN, OVERRUN, ABORTED OR INVALID FRAME ERRORS
- STRAIGHT FORWARD CPU INTERRUPTS
- PROGRAMMABLE MODEM CONTROL INTERRUPTS
- DOUBLE BUFFERING OF DATA
- DMA COMPATABILITY
- END OF BLOCK OPTION
- VARIABLE CHARACTER LENGTH (5, 6, 7 OR 8 BITS)
- RESIDUAL CHARACTER CAPABILITY
- ADDRESS COMPARE
- GLOBAL ADDRESS RECOGNITION
- EXTENDABLE ADDRESS FIELD
- EXTENDABLE CONTROL FIELD
- AUTOMATIC ZERO INSERTION AND DELETION
- MAINTENANCE MODE FOR SELF-TESTING

APPLICATIONS

COMPUTER COMMUNICATIONS

TERMINAL COMMUNICATIONS

COMPUTER TO MODEM INTERFACING

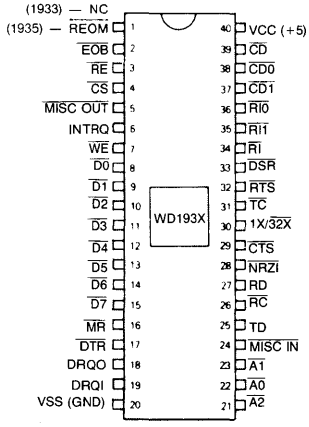

Figure 1. WD193X

PIN DESIGNATION

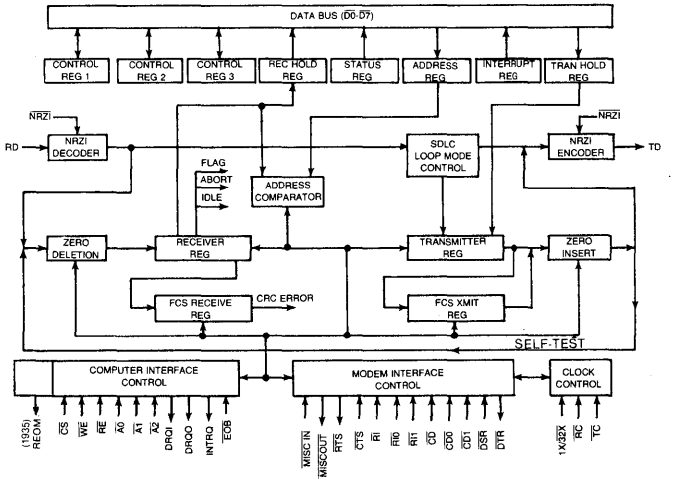

Figure 2. WD193X BLOCK DIAGRAM

WD1993 Arinc 429 Receiver/Transmitter and Multi-Character Receiver/Transmitter

FEATURES

- PRESENT UPON MASTER RESET FOR ARINC 429 PROTOCOL
- PROGRAMMABLE WORD LENGTH FROM 1 CHARACTER TO 8 CHARACTERS
- PROGRAMMABLE CHARACTER LENGTH, 5, 6, 7, OR 8 BITS
- RETURN TO ZERO (RZ) OUTPUT
- AUTO SPACE GENERATION
- DOUBLE BUFFERED RECEIVER AND TRANSMITTER
- UNDERRUN ERROR DETECTION FOR TRANSMISSION
- OVERRUN, FRAMING AND PARITY ERROR DETECTION ON RECEIVER
- WORD ERROR FLAG FOR COMPREHENSIVE ERROR REPORTING
- FIRST CHARACTER OF WORD FLAG FOR SINGLE INTERRUPT APPLICATIONS
- DIAGNOSTIC LOCAL LOOP-BACK TEST MODE
- DC TO 200 KILOBITS PER SECOND OPERATION
- TTL COMPATIBLE INPUTS AND OUTPUTS
- SINGLE +5 VOLT SUPPLY
- TEMPERATURE RANGES 0°C to 70°C, — 1993-03, −40°C to +85°C — 1993-02, −55°C to +125°C — 1993-01

INTRODUCTION

The Western Digital WD1993 Avionic Receiver/Transmitter is designed to handle digital data transmission, according to the Avionic Arinc 429 protocol. Also, the word length is programmable from one to eight characters of 5, 6, 7, or 8 bits. Parallel data is converted into a serial data stream during transmission and serial to parallel during reception. The WD1993 is packaged in a 28 pin plastic or ceramic package and is available in three temperature ranges: Commercial, Industrial and Military.

GENERAL DESCRIPTION

The WD1993 is a bus-orientated MOS/LSI device designed to provide the Avionics Arinc 429 Data Communication Protocol, along with programmable character length capabilities.

Also, the WD1993 contains a local loop-back test mode of operation, which is controlled by the Loop Test Enable (LTE) bit in the command register. In this diagnostic mode, the transmitter output is "looped-back" into the receiver input. The REN and TEN control bits must also be active ("1") and the $\overline{CTS}$ input must be low. The status and output flags operate normally.

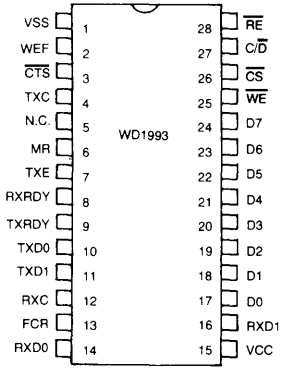

Figure 1 PIN DIAGRAM

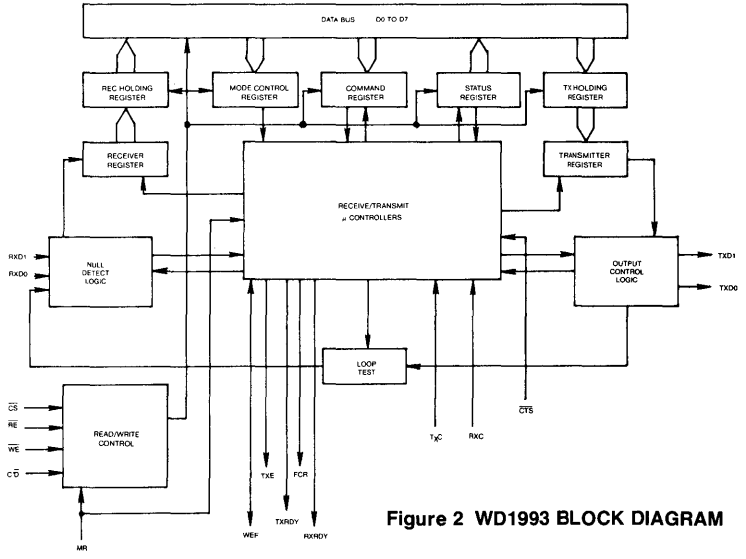

Figure 2 WD1993 BLOCK DIAGRAM

WD2001/WD2002 Data Encryption Devices

FEATURES

- CERTIFIED BY NATIONAL BUREAU OF STANDARDS.
- TRANSFER RATE:
  WD2001/2-05 300Kbs with 500KHz clock
  WD2001/2-20 1.3 Mbs with 2MHz clock
  WD2001/2-30 1.8 Mbs with 3MHz clock
- ENCRYPTS/DECRYPTS 64 BIT DATA WORDS USING 56 BIT KEY WORD
- SINGLE PORT 28 PIN PACKAGE WD2001 OR DUAL PORT 40 PIN PACKAGE WD2002
- COMMAND BIT PROGRAMMING VIA DAL BUS OR INPUT PINS
- DMA COMPATIBLE (SEE WESTERN DIGITAL DM1883)
- PARITY CHECK ON KEY WORD LOADING
- STANDARD 8 BIT MICROPROCESSOR INTERFACE
- INPUTS AND OUTPUTS TTL COMPATIBLE
- KEY STORED ON CHIP IS NOT EXTERNALLY ACCESSIBLE

  SEPARATE CLEAR AND CIPHER BUS STRUCTURE ON WD2002

APPLICATIONS

- SECURE BROKERAGE TRANSACTIONS
- ELECTRONIC FUNDS TRANSFERS
- SECURE BANKING/BUSINESS ACCOUNTING
- MAINFRAME COMMUNICATIONS
- REMOTE AND HOST COMPUTER COMMUNICATIONS
- SECURE A/D
- SECURE DISK OR MAG TAPE DATA STORAGE
- SECURE PACKET SWITCHING TRANSMISSION

GENERAL DESCRIPTION

The Western Digital WD2001 and WD2002 Data Encryption/Decryption devices are designed to encrypt and decrypt 64-bit blocks of data using the algorithm specified in the Federal Information Processing Data Encryption Standard (#46). These devices encrypt a 64-Bit clear text word using a 56-Bit user-specified key to produce a 64-Bit cipher text word. When reversed, the cipher text word is decrypted to produce the original clear text word.

The DE2001/2 are fabricated in N-channel silicon gate MOS technology and are TTL compatible on all inputs and outputs.

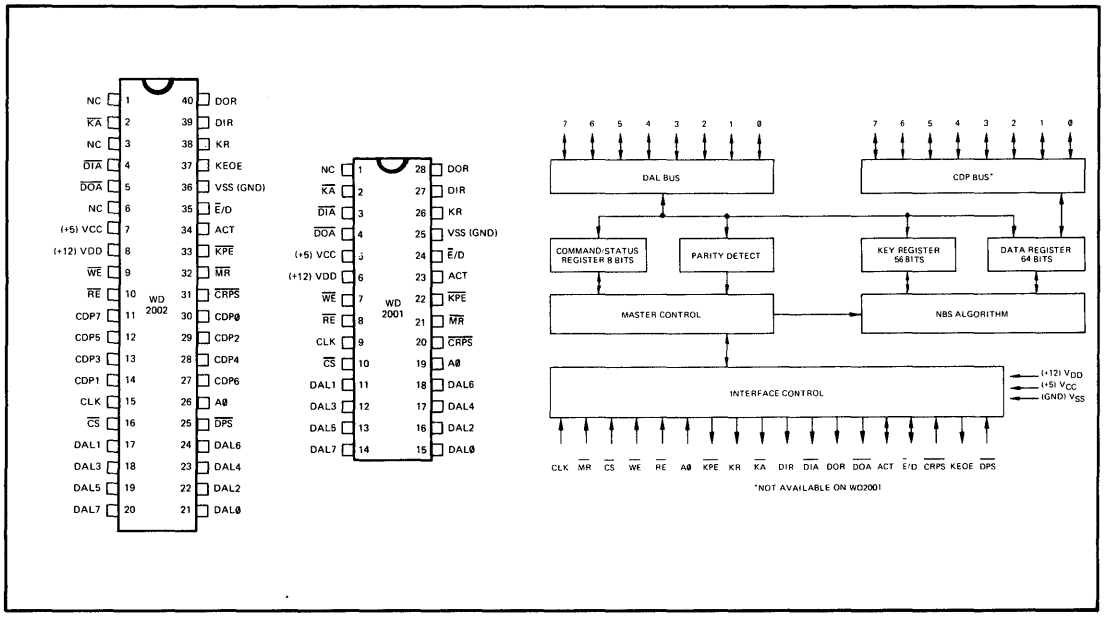

WD2001/WD2002 BLOCK DIAGRAM

DM1883A/B Direct Memory Access Controller

FEATURES

- AUTOMATIC DAISY CHAINING OF BUS AND INTERRUPT ACKNOWLEDGE SIGNALS
- AUTO LOAD OPTION
- SINGLE +5 VDC POWER SUPPLY
- 8 BIT BI-DIRECTIONAL DATA BUS
- TRUE OR COMPLEMENT DATA BUS
- 8 CPU ADDRESSABLE DMAC REGISTERS
- 8 CPU ADDRESSABLE DEVICE REGISTERS
- AUTOMATIC GENERATION OF DEVICE CS DURING DMA AND CPU DEVICE ACCESSES
- 256K MEMORY ADDRESSING
- 64K PROGRAMMABLE PAGE PROTECTION
- BYTE OR WORD DMA TRANSFERS
- INTERRUPT AND BUS REQUEST CAPABILITIES
- END-OF-BLOCK SHUT OFF BY DMAC
- TIME-OUT INTERRUPT CAPABILITY
- SINGLE CLOCK INPUT
- CS, RE, WE, A0-A3 ADDRESSING
- STOP REQUEST INPUT TO DELAY INTERRUPT OR BUS REQUESTS
- COMPATIBLE WITH OUR FLOPPY DISC CONTROLLERS
- 8 BIT PROGRAMMABLE INTERRUPT ID CODE

GENERAL DESCRIPTION

The DM1883 Direct Memory Access Controller (DMAC) is packaged in a 40 pin standard dual in-line package. The chip requires a single +5 power supply input and a single clock input. The device contains 8 CPU addressable registers, and allows for up to 8 CPU addressable device registers if the automatic device chip select feature is used. Byte or word transfers can be programmed, and all memory DMA operations are handshaked for compatibility with a variety of bus structures. Up to 256K bytes of memory can be accessed directly with 64K page protection and nonexistent memory interrupt as options. Bus and Interrupt Acknowledge signals are internally daisy chained, and a STOP REQUEST input prevents new requests while a current request is active. Device accesses are not handshaked, and a BUS HOLD feature is present for high speed devices. Device interrupt input, end-of-block output, and I/O read/write output pins simplify hardware interfacing to the device and the CPU bus. The AUTO LOAD feature allows automatic boot-loading of up to 64K bytes or words into memory starting at location zero. An 8 bit interrupt ID code is also provided.

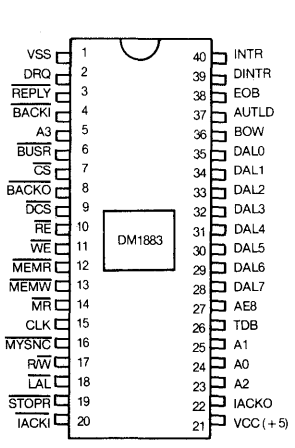

PIN CONNECTIONS

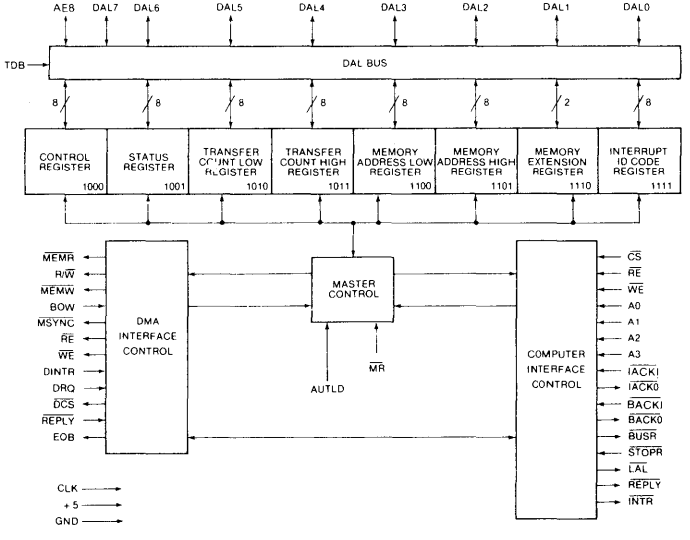

DM1883 BLOCK DIAGRAM

Application Note
Using The WD2501/2511

This application note provides an introduction to the X.25 communication protocol and introduces the ISO reference model. The link layer of X.25 is highlighted as it can today be implemented with a single LSI device, the WD2501/WD2511.

The bulk of this document provides details of the hardware and software interfaces that a user typically encounters when using the WD2501. Schematic and timing diagrams for a typical Z80 interface along with high level flowcharts for initialization and operation are given.

Three appendices cover a glossary, a discussion of X.25 LAP vs LAPB, and using the WD2501/WD2511 in multi-point configuration.

CONTENTS

RELATED DOCUMENTS

WD2501 Specification

WD2511 Specification

"LSI Ready to Make a Mark on Packet-Switching Networks," Geary Leger, Electronics, December 20, 1979.

"LSI Circuit Simplifies Packet-Network Connection," Geary Leger, Electronics, December 20, 1979.

1.0 THE WD2501/2511 General Description

The WD2501 and WD2511 are LSI devices that fully handle the link level (level 2) of the CCITT X.25 communications protocol. They are pin-compatible with each other and with other members of the Western Digital Network Controller Family.

The distinction between the WD2501 and the WD2511 lies in their internal firmware: the WD2501 implements LAP (Link Access Procedure) whereas the WD2511 implements LAPB (Link Access Procedure Balanced). The protocol differences between LAP and LAPB are described in Appendix B.

NOTE:
Unless otherwise noted in this document, all references to the WD2501 should be assumed to pertain to the WD2511 as well.

In addition to the traditional parallel/serial converters and FCS logic, the WD2501 incorporates a highly efficient micro-programmed processor that fully handles the required link set-up and frame sequencing operations conventionally delegated to a "user defined" processor. The WD2501 also contains an intelligent two-channel DMA controller to further simplify its integration into a user's system.

2.0 THE WD2501 AND THE ISO MODEL

The CCITT X.25 recommendation comprises three levels of protocols (Level 1 to 3). See Figure 1.

Level 1 is the physical level, which concerns the actual means of bit transmission across a physical medium.

Level 2 is the link level which includes frame formatting, error control and link control.

Level 3 is the packet (network) level which controls the traffic of the different virtual calls and multiplexes these for passage over the physical line.

These three levels are completely independent of each other, which allows changes to be made to one level without disrupting the operation of any other level. An adjacent level is affected only if the changes affect the interface to that level.

Each level performs one well defined set of functions, using only a well defined set of services provided by the level below. These functions implement a set of services that can be accessed only from the level above. Each level is strictly controlled by the systems engineer according to formal functional and interface specifications.

The WD2501 implements level 2. Without additional logic, it generates the frame, performs error checking, performs link management (set up/disconnect) and ensures reliable data transmission by evaluating the sequence numbers associated with each I-frame. The device automatically acknowledges received I-frames and fully supports up to 7 outstanding (unacknowledged) frames, including retransmission if required.

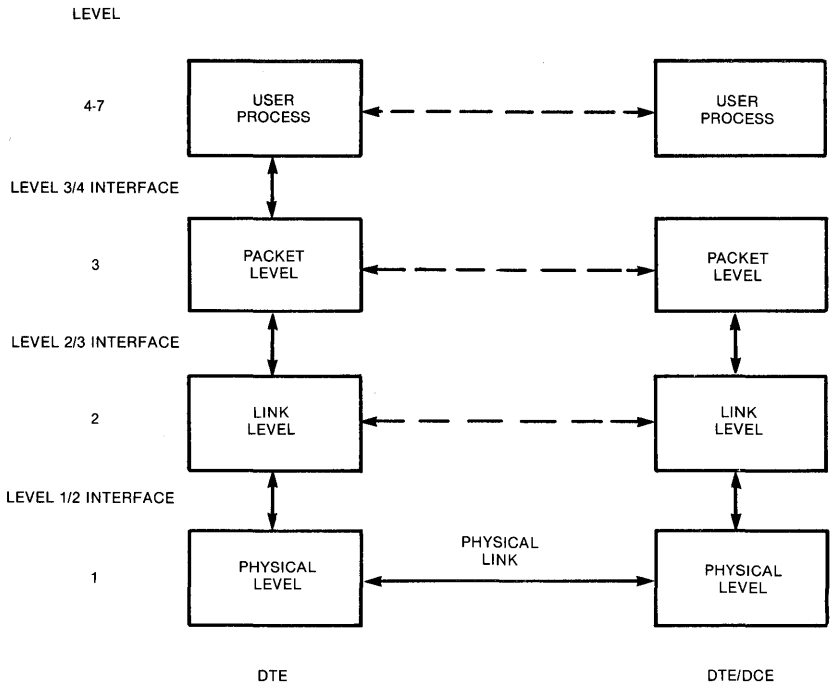

Figure 1. LAYERED ARCHITECTURE FOR COMPUTER NETWORKS

The only real physical connection between the two stations (DTE to DTE/DCE) is the Physical Link between the two physical layers. The other connections shown between two of the same layers (peer to peer interface) is not a physical but rather a logical connection made up by the respective protocols for that particular level.

Each level n "interfaces" to the corresponding level n on the other side of the Data Communication Link through the level n-1, then n-2 etc., via the physical link and up through the levels to n-2, n-1 and to level n.

3.0 HARDWARE

The WD2501 must be connected to the Physical Level (Level 1). This generally amounts to simple line drivers/receivers.

A typical X.25 DTE/DCE station block diagram is shown in Figure 2. Figure 3 shows a circuit diagram of the actual X.25 hardware interface of this same station. Table 1 is a description of signal functions for this circuit diagram. This is to be connected directly to a Z80 microprocessor on one side and an EIA RS-422 interface on the other side.

Figures 4 and 5 are DMA cycle timing diagrams for this particular station.

General notes to this interface:

- A modem would be needed for long-distance communication lines.
- The hardware interface in Figure 3 includes all hardware options. Simpler interfacing is possible.
- The function of the CPU Bus Driver Control Circuit (CBDCC) is to control the direction and/or timing of the data-line transceivers and the two address latches.
- If the CPU clock frequency is not higher than the WD2501 CLK maximum frequency, the High Speed Control Circuit (HSCC) is not needed. The function of the HSCC circuit is to divide an high speed CPU clock signal (0) down to half the frequency (01A). It also delays the reset of $\overline{\text{BUSRQ}}$ with one additional 01 clock cycle when a high speed CPU clock is used. These functions are needed to establish a time window of at least 500ns between $\overline{\text{DACK}}$ being active and a CPU Write/Read function.

- When a high speed CPU clock is used, connect 01A signal to 01 signal and $\overline{\text{BUSRQA}}$ signal to $\overline{\text{BUSRQ}}$ signal. When a low speed CPU clock is used, connect CPU clock (0) direct to 01 signal and $\overline{\text{BUSRQ2}}$ signal to $\overline{\text{BUSRQ}}$ signal.
- The CLK circuit is needed in all applications. Its purpose is to stop the CLK signal from going LO during a CPU read cycle. This logic will avoid any possible interference between an internal updating of a status register while the CPU reads this same register.
- The DMA I/O circuit matches the timing between the Z80 and the WD2501.
- The $\overline{\text{RTS}}$ open collector output needs a pull-up resistor.
- In this particular example, line drivers/receivers are of type EIA RS-422. However, RS-232C or RS-423 can also easily be used.
- Port A of a PIO in this example is programmed to be an output. In this case, the CPU controls the DTR output to the modem. Port B of the PIO is programmed to be interrupt controlled inputs; the CPU can be interrupted by $\overline{\text{DSR}}$ and/or $\overline{\text{INTR}}$ as programmed.
- The WD2501 CLK input is derived from the Z80 CPU clock. The CLK may be supplied from an independent clock generator. In that case, instead of connecting the 01 clock to the CLK circuit, connect the independent clock to where the 01 clock is shown in the CLK circuit in Figure 3.
- $\overline{\text{MRW}}$ (Memory Read/Write) signal enables the output of the memory address decoder for the computer system memory chips. As an example, if a PROM type 28S42 is used as the memory address decoder, $\overline{\text{MRW}}$ is connected direct to $\overline{\text{E}}$ (Pin 15) input.
- The WD2501 $\overline{\text{CS}}$ input is to be connected to a port address decoder (or memory address decoder). $\overline{\text{MWE}}$ is connected to all $\overline{\text{WE}}$ inputs, and $\overline{\text{MOE}}$ is connected to all $\overline{\text{OE}}$ inputs of the system memory chips.
- $\overline{\text{REPLY}}$ output is not used in this application.

3.1 READ/WRITE CONTROL OF I/O REGISTERS

The sixteen I/O registers are directly accessible from the CPU data bus (DAL0-DAL7) by a read and/or write operation by the CPU. The CPU must activate the WD2501 register address (IA0-IA3), Chip Select ($\overline{\text{CS}}$), Write Enable ($\overline{\text{WE}}$) or Read Enable ($\overline{\text{RE}}$) before each data bus transfer operation. The read/write operation is completed when $\overline{\text{CS}}$ or $\overline{\text{RE}}$ /$\overline{\text{WE}}$ is brought high. During a write operation, the falling edge of $\overline{\text{WE}}$ will initiate a WD2501 write cycle. The addressed register will then be loaded with the content of the Data Bus. The rising edge of $\overline{\text{WE}}$ will latch that data into the addressed register.

During a read operation, the falling edge of $\overline{\text{RE}}$ will initiate a WD2501 read cycle. The addressed register will then place its content onto the Data Bus.

To insure that all six status registers (registers 2 thru 7, inclusive) are read correctly, either rule A **or** rule B must be followed (or both):

A) Enable $\overline{\text{CS}} \bullet \overline{\text{RE}}$ only while the CLK input is high.

B) Read the status registers only after an interrupt, and read within one frame time of the interrupt.

If rule A) is used, the high CLK may (at the user's option) be extended by skipping one, and only one, negative going CLK whenever a $\overline{\text{CS}} \bullet \overline{\text{RE}}$ is executed. This is shown below:

One negative CLK
is skipped

As a reminder, the long form specification already states that CLK must be a square wave.

Non-conformance to the Rule above could mean that the CPU will at rare times read incorrect values of NA, NB, NE, ERO, or any other status.

The CPU must set-up all transmit data, TSADR HI and LO, TCNT HI and LO, and residual bits before setting BRDY in the applicable TLOOK segment.

The CPU must set aside receiver memory (at least one chain segment with transfer address), and set-up RSADR HI and LO before setting REC RDY in the applicable RLOOK segment.

3.2 DMA IN/OUT OPERATION

The Direct Memory Access (DMA) operation is completely controlled by the WD2501. During a DMA cycle, the CPU sets its address bus, data bus and three-state control signals to their high impedance states.

(See DMA In/Out timing diagrams, Figures 4 and 5.)

In this application example, the data bus transceivers are permanently enabled (low impedance state). When the CPU has control, the direction of these transceivers is pointing from the CPU bus towards the WD2501. During a DMA In cycle, this is not changed. During a DMA Out cycle however, the direction is reversed (WD2501 towards the CPU bus).

The address bus latches are in high impedance state while the CPU has control of the bus. When the WD2501 has control of the CPU bus, the address latches are in the low impedance state. During the DMA Out cycle, these latches function as regular bus drivers. During the DMA In cycle however, the address gets latched to assure enough data hold time for the WD2501.

3.2.1 DMA IN

During a DMA In cycle, the task of transferring one byte of I-field data from memory into the WD2501 is performed. The CPU time (in the example described in this paragraph to execute this task) is five T-states for a low speed CPU clock system and ten T-states for a high speed CPU clock system.

The DMA In function starts when the WD2501 is ready to receive a byte from memory to be transmitted out to the remote station. This condition causes the $\overline{DRQI}$ signal to go LO, which in turn activates the $\overline{BUSRQ}$ (Bus Request) signal. Also at this time (ADRV bit = 1), the WD2501 presents the address (on A0-A15) of the data byte to be retrieved from memory.

The $\overline{BUSRQ}$ signal is sampled by the CPU with the rising edge of the last CPU clock (0) period of any machine cycle. In this case, because the $\overline{BUSRQ}$ signal is active, the CPU goes into high impedance state with the rising edge of the next CPU clock pulse. At this time, the CPU also switches the control over to the WD2501 by activating the $\overline{BUSAK}$ signal. This causes $\overline{DACK}$ to go LO at the following rising edge of 01 clock. This is the actual indication for the WD2501 to start the DMA In cycle. $\overline{DACK}$ also causes $\overline{DRQI}$ to return to the HI state.

At the next rising edge of the 01 clock, $\overline{MOE}$ (Memory Output Enable) is activated. This causes the memory to output the addressed data byte onto the Data Bus. Also, the address is now latched into the address bus latches (74LS373) at this time.

At the next falling edge of the 01 clock, $\overline{DACK}$ gets deactivated, causing the WD2501 to latch the data byte (DAL0-DAL7) and to set its address lines (A0-A15) to logical HI state (ADRV bit set). The address bus latches hold the address active until $\overline{DMOE}$ signal is deactivated.

At the next rising edge of the 01 clock (low speed CPU clock), $\overline{BUSRQ}$ gets deactivated. When high speed CPU clock is used, $\overline{BUSRQ}$ is deactivated after an additional 01 clock cycle.

At the next following rising edge of 01 clock, $\overline{BUSRQ}$ is sampled by the CPU. This causes $\overline{BUSAK}$ and $\overline{MOE}$ to become deactivated, but not until the next falling edge of the CPU (0) clock. This is the end of the DMA In cycle. At the next rising edge of the CPU clock, the CPU again controls the CPU bus.

3.2.2 DMA OUT

This operation is very similar to the DMA In function. During this cycle, one byte of I-field data is transferred from the WD2501 to the memory. The CPU-time in this example described to perform this task is the same as for the DMA In cycle.

The DMA Out function starts when the WD2501 is holding a received I-field byte and is ready to transfer this to the memory. This condition activates the $\overline{DRQO}$ signal, which in turn sets the $\overline{BUSRQ}$ to LO. Also at this time (ADRV bit = 1), the WD2501 presents the address to the memory location to where the respective data byte is to be loaded.

The $\overline{BUSRQ}$ signal is sampled by the CPU with the rising edge of the last CPU clock period of any machine cycle. Since the $\overline{BUSRQ}$ signal is active, the CPU goes into high impedance state with the rising edge of the following CPU clock pulse. Now the CPU also switches the control over to the WD2501 by activating the $\overline{BUSAK}$ signal. This causes $\overline{DACK}$ to go LO at the next rising edge of 01 clock, which indicates to the WD2501 to start the DMA Out cycle. This causes $\overline{DRQO}$ to reset back to HI state and to load the data byte to be transferred onto the data-bus.

At the next rising edge of the 01 clock, $\overline{MWE}$ (Memory Write Enable) is activated. This causes the memory to input the addressed data byte.

At the following rising edge of the 01 clock, $\overline{MWE}$ goes HI, latching the data into the memory. Also at this time (low speed CPU clock), $\overline{BUSRQ}$ signal gets deactivated. When high speed CPU clock is used, $\overline{BUSRQ}$ is deactivated after an additional 01 clock cycle.

At the next rising edge of 01 clock, the $\overline{BUSRQ}$ signal is sampled by the CPU.

$\overline{DACK}$ goes HI half a 01 clock cycle after $\overline{MWE}$ goes HI. This ends the DMA Out cycle by the WD2501 setting its data-lines in high impedance state and the address-lines (A0-A15) to logical HI state (ADRV bit set).

After CPU has sampled and detected $\overline{BUSRQ}$ being deactivated, it resets $\overline{BUSAK}$ to HI at the next falling edge of the CPU clock.

At the next rising edge of the CPU clock, the CPU again controls the bus.

TABLE 1
SIGNAL NAMES FOR THE HARDWARE INTERFACE (See Note)

NAME	SYMBOL	FUNCTION
RECEIVE	$\overline{RCV}$	When activated (LO), sets the direction of the data bus transceivers from WD2501 towards the CPU bus. This is done only during CPU Read or DMA Out cycle.
DRM TRANSFER	$\overline{DTFR}$	When activated (LO), enables the output of the address bus latches. This is done during a DMA In/Out cycle.
DMA MEMORY OUTPUT ENABLE	$\overline{DMOE}$	Is activated during a DMA In cycle. Generates the $\overline{MOE}$ signal and latches the DMA In addresses.
MEMORY OUTPUT ENABLE	$\overline{MOE}$	Is activated during a DMA In or a CPU Read cycle. Enables the memory outputs. Is to be connected to the $\overline{OE}$ pin of the memory circuits.
MEMORY WRITE ENABLE	$\overline{MWE}$	Is activated during a DMA Out or a CPU Write Cycle. Enables the memory write function. Is to be connected to the $\overline{WE}$ input of the memory circuits.
MEMORY READ/WRITE	$\overline{MRW}$	Is activated during a DMA In/Out function or a Memory Read/Write cycle by the CPU. Enables the output of the Memory Address decoder.
DMA OUT	DMA OUT	Is activated during a DMA Out function.
DMA IN	DMA IN	Is activated during a DMA In function.
INTERNAL LOOP	ILOOP	Is activated during an internal loop-back test. Keeps the $\overline{RTS}$ signal to the modem in off condition and logically connects $\overline{RTS}$ to $\overline{CTS}$.
BUS ACKNOWLEDGE 1	$\overline{BUSAK1}$	When active, indicates that the CPU has switched bus control over to the WD2501. Compared to $\overline{BUSAK}$ signal, this is delayed one 01 clock cycle when going LO to allow a time window of at least 500 ns before $\overline{DACK}$ becomes activated.
BUS REQUEST 1	$\overline{BUSRQ1}$	When active, requests the CPU via $\overline{BUSRQ2}$ and $\overline{BUSRQ}$ (low speed CPU clock) to switch control over to the WD2501.
BUS REQUEST 2	$\overline{BUSRQ2}$	Same function as $\overline{BUSRQ1}$, except that $\overline{BUSRQ2}$ is delayed one 01 cycle when going HI. The delay allows a time window of at least 500 ns between $\overline{DACK}$ being active and a CPU Read/Write function.
BUS REQUEST A	$\overline{BUSRQA}$	Same function as $\overline{BUSRQ2}$, except is delayed an additional 01 cycle when going HI. This delay allows the necessary 500 ns time window between $\overline{DACK}$ being active and a CPU Read/Write function when an high speed CPU clock is used. This is then directly driving the $\overline{BUSRQ}$ signal.

TABLE 1 (Continued)
SIGNAL NAMES FOR THE HARDWARE INTERFACE (See Note)

NAME	SYMBOL	FUNCTION
CLOCK	CLK	Clock used for internal timing of WD2501. Is synchronous with 01 clock with the exception that when the CPU is performing a read function, CLK signal will not go to LO state. This function avoids an internal up-date of a status register while the CPU is reading this same register.
01	01	Clock used for timing of this hardware interface. The 01 frequency is not allowed to be higher than the CLK max. frequency of the WD2501. When a low speed CPU clock is used, 01 signal is connected directly to the CPU clock (0). When an high speed CPU clock is used, this is connected to the 01A signal.
01A	01A	Clock signal with half the CPU clock frequency. This is driving the 01 clock when an high speed CPU clock is used. Toggles only during DMA In/Out cycles.
+5 RESISTIVE	+5R	+5V through a resistor.

NOTE: Signals described in this paragraph are signals generated by this circuitry only. Other signals are described in either the WD2501 device specification or in the Z80 CPU data sheets.

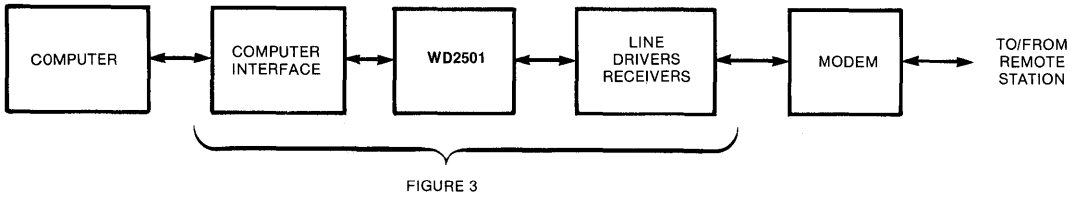

FIGURE 3

Figure 2. DTE/DCE STATION BLOCK DIAGRAM

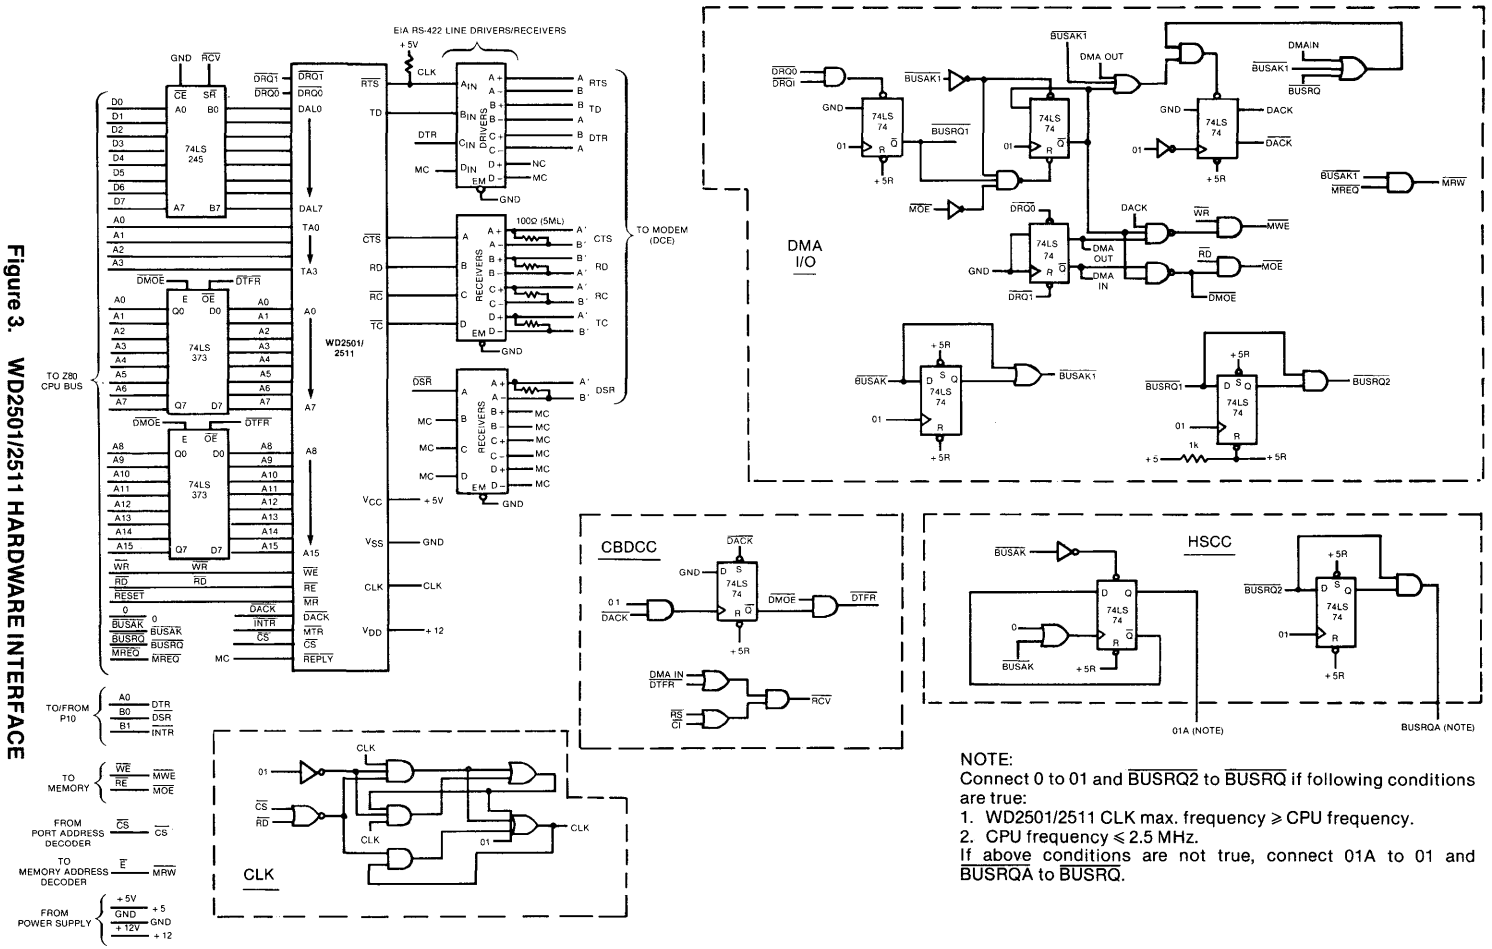

Figure 3. WD2501/2511 HARDWARE INTERFACE

NOTE:
Connect 0 to 01 and $\overline{BUSRQ2}$ to $\overline{BUSRQ}$ if following conditions are true:

1. WD2501/2511 CLK max. frequency ≥ CPU frequency.
2. CPU frequency ≤ 2.5 MHz.

If above conditions are not true, connect 01A to 01 and $\overline{BUSRQA}$ to $\overline{BUSRQ}$.

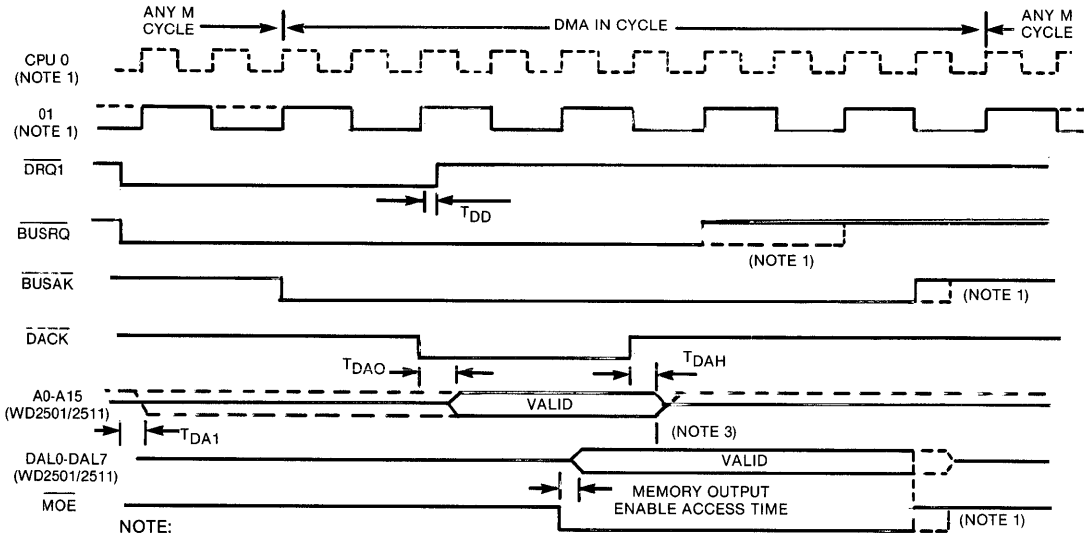

NOTE:
1. These dashed lines show the respective signals when a fast CPU 0 clock is used. Solid lines show same signals for a slow CPU 0 clock (CPU 0 = 01).
2. Propagation delay is omitted in this diagram unless timing symbol is shown.
3. Data would become non-valid after the address becomes non-valid if address latches were not used.

Figure 4. DMA IN TIMING

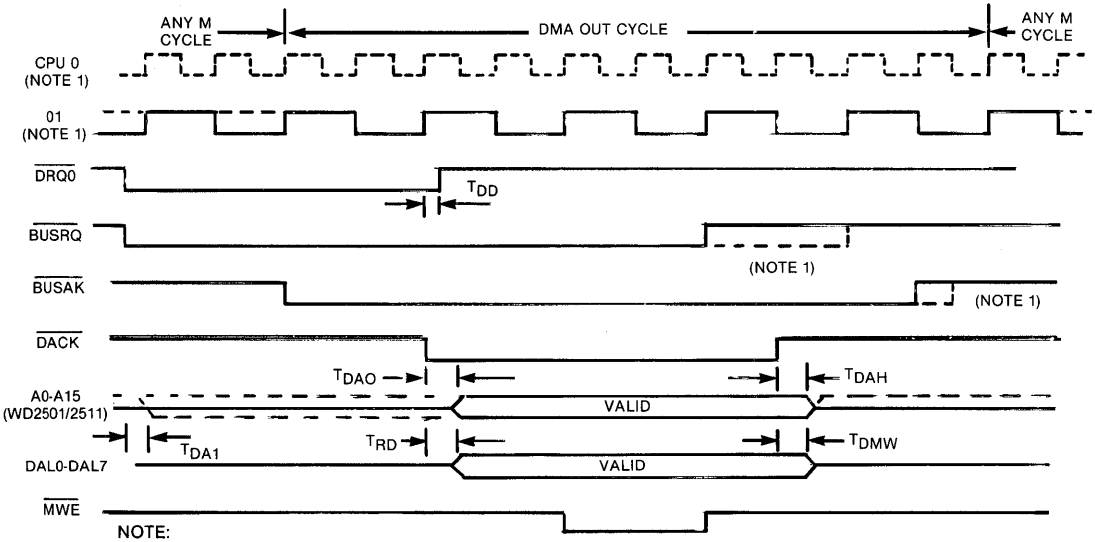

NOTE:
1. These dashed lines show the respective signals when a fast CPU 0 clock is used. Solid lines show same signals for a slow CPU 0 clock (CPU 0 = 01).
2. Propagation delay is omitted in this diagram unless timing symbol is shown.

Figure 5. DMA OUT TIMING

3.3 SERIAL INTERFACE

The receiver and transmitter sub-systems are completely independent of each other, the CPU Read/Write functions and the DMA In/Out functions.

The serial data is synchronized by the externally supplied $\overline{TC}$ clock and $\overline{RC}$ clock. The falling edge of $\overline{TC}$ generates new transmitted data and the rising edge of $\overline{RC}$ is used to sample the received data.

After initialization and before the first frame is sent, the TD output sends Idles (continuous 1s).

After the first frame is sent or the ACTIVE/$\overline{PASSIVE}$ bit is set, continuous flags are sent in between frames.

For detailed information on what type of frames are sent for certain conditions, see the WD2501 or WD2511 specifications.

4.0 SOFTWARE

Initialization of the WD2501 and I-field data processing (level 3) is accomplished by user written software. This software need not be realtime, since the WD2501 responds to link exceptions and overhead functions on its own.

Configuring the WD2501 to certain test functions, modes, timer values, location of initial memory pointers, chain buffer lengths and link level addresses is performed via the sixteen I/O registers.

All buffer management support, buffer chaining and free/busy flags, occur in user memory. Here two look-up tables (TLOOK/RLOOK), located in the user memory, contain pointers/counters for up to eight outstanding transmit/receive packets. The WD2501 contains only one address pointer, which is the starting address of Segment #0 in the TLOOK table. Segment #0 in the RLOOK table always begins 40(Hex) bytes after TLOOK, Segment #0, byte #0. See section "Memory Access Scheme" in the WD2501 specifications.

Link monitoring is done by use of the I/O registers and the memory buffers. The WD2501 indicates to the system CPU that a certain event has occured by setting a bit in status register 1 and setting the interrupt flag. This indicates whether a packet has been received, a transmitted packet has been acknowledged, a non-recoverable error condition or some other condition needs the attention of the CPU.

In this section, a flow-chart is given to show the user how to program the WD2501. For more details refer to the data sheets.

The flow for programming/monitoring the WD2501 for transmitting or receiving a packet(s) or for a loop-back test is shown in the flowchart below. The flow starts at START1 if a power-up was just done and/or if no data communication environment programming (initialization) has been done.

If initialization is complete, the flow starts at START2 when the WD2501 is to be enabled to receive a packet(s).

If a packet/s is to be transmitted and initialization is complete, the flow starts at START3.

WD2501/2511 PROGRAMMING FLOWCHART

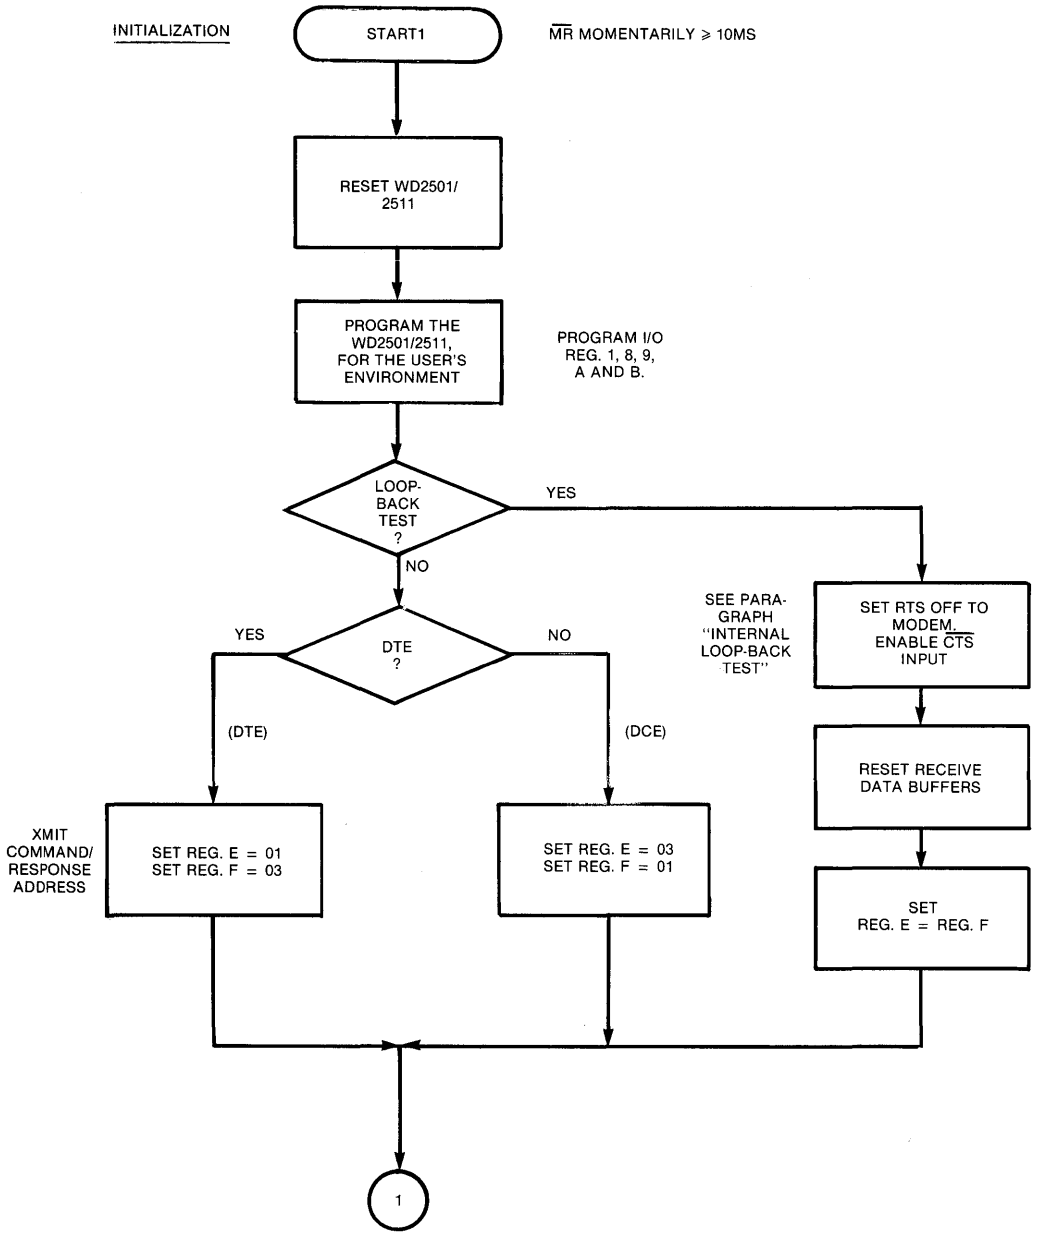

ENABLE RECEPTION OF PACKET(S)

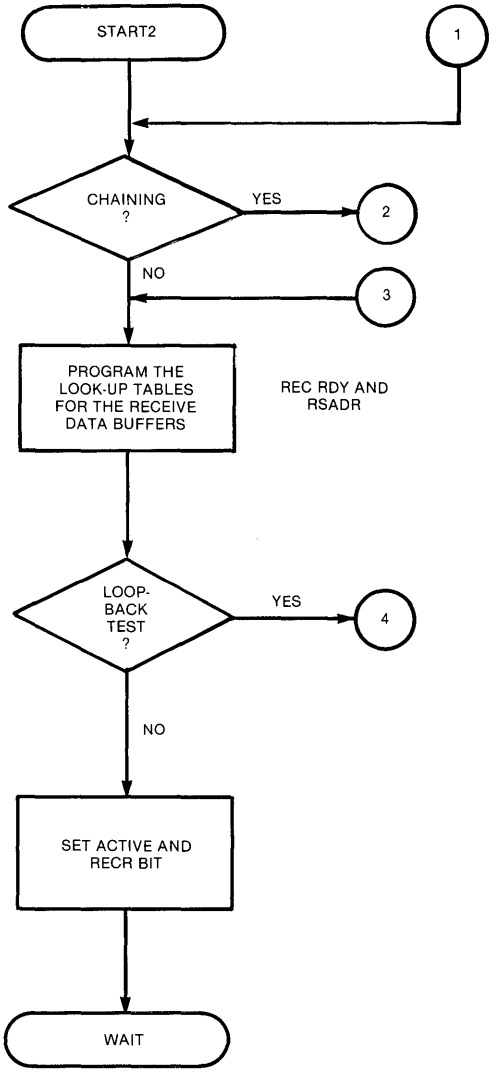

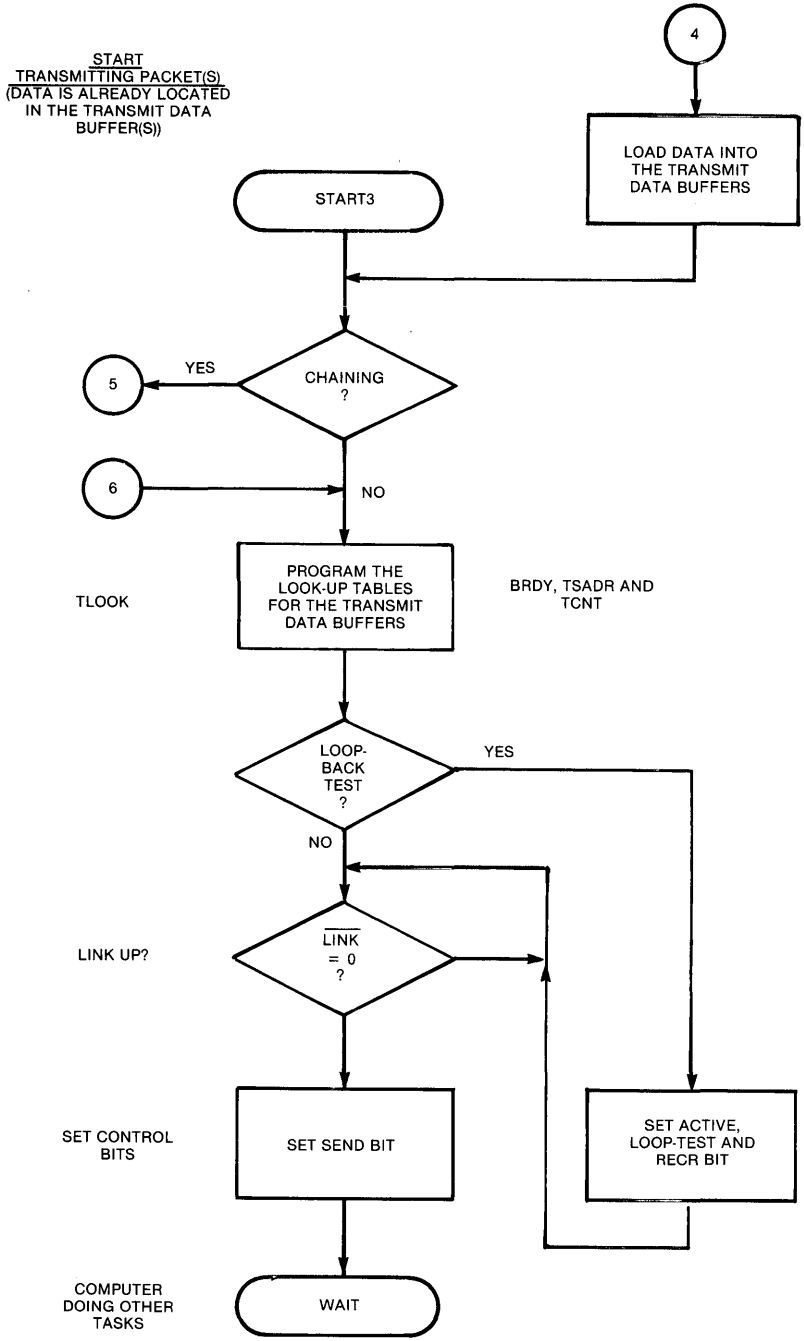
START
TRANSMITTING PACKET(S)
(DATA IS ALREADY LOCATED
IN THE TRANSMIT DATA
BUFFER(S))
4
LOAD DATA INTO
THE TRANSMIT
DATA BUFFERS
START3
5
YES
CHAINING
?
6
NO
TLOOK
PROGRAM THE
LOOK-UP TABLES
FOR THE TRANSMIT
DATA BUFFERS
BRDY, TSADR AND
TCNT
LOOP-
BACK
TEST
?
YES
NO
LINK UP?
LINK
= 0
?
SET CONTROL
BITS
SET SEND BIT
SET ACTIVE,
LOOP-TEST AND
RECR BIT
COMPUTER
DOING OTHER
TASKS
WAIT

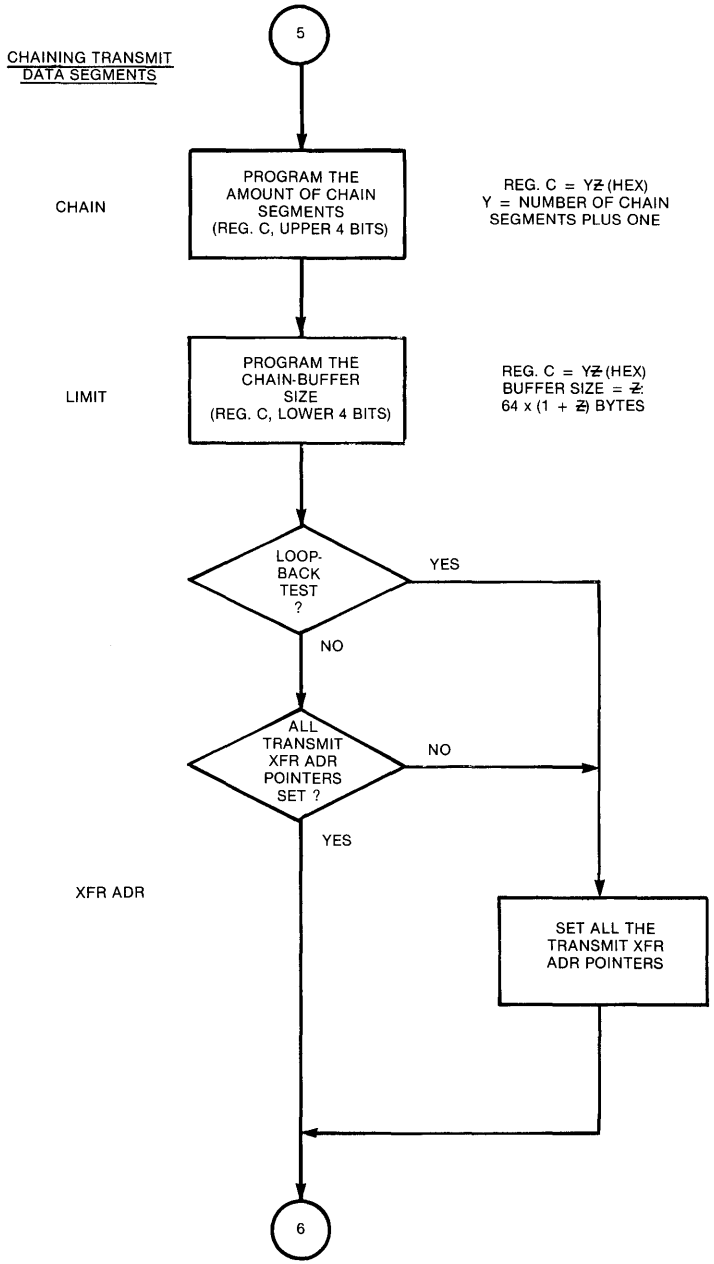
CHAINING TRANSMIT DATA SEGMENTS
5
CHAIN
PROGRAM THE AMOUNT OF CHAIN SEGMENTS (REG. C, UPPER 4 BITS)
REG. C = YZ (HEX)
Y = NUMBER OF CHAIN SEGMENTS PLUS ONE
LIMIT
PROGRAM THE CHAIN-BUFFER SIZE (REG. C, LOWER 4 BITS)
REG. C = YZ (HEX)
BUFFER SIZE = Z:
64 x (1 + Z) BYTES
LOOP-BACK TEST ?
YES
NO
ALL TRANSMIT XFR ADR POINTERS SET ?
NO
YES
XFR ADR
SET ALL THE TRANSMIT XFR ADR POINTERS
6

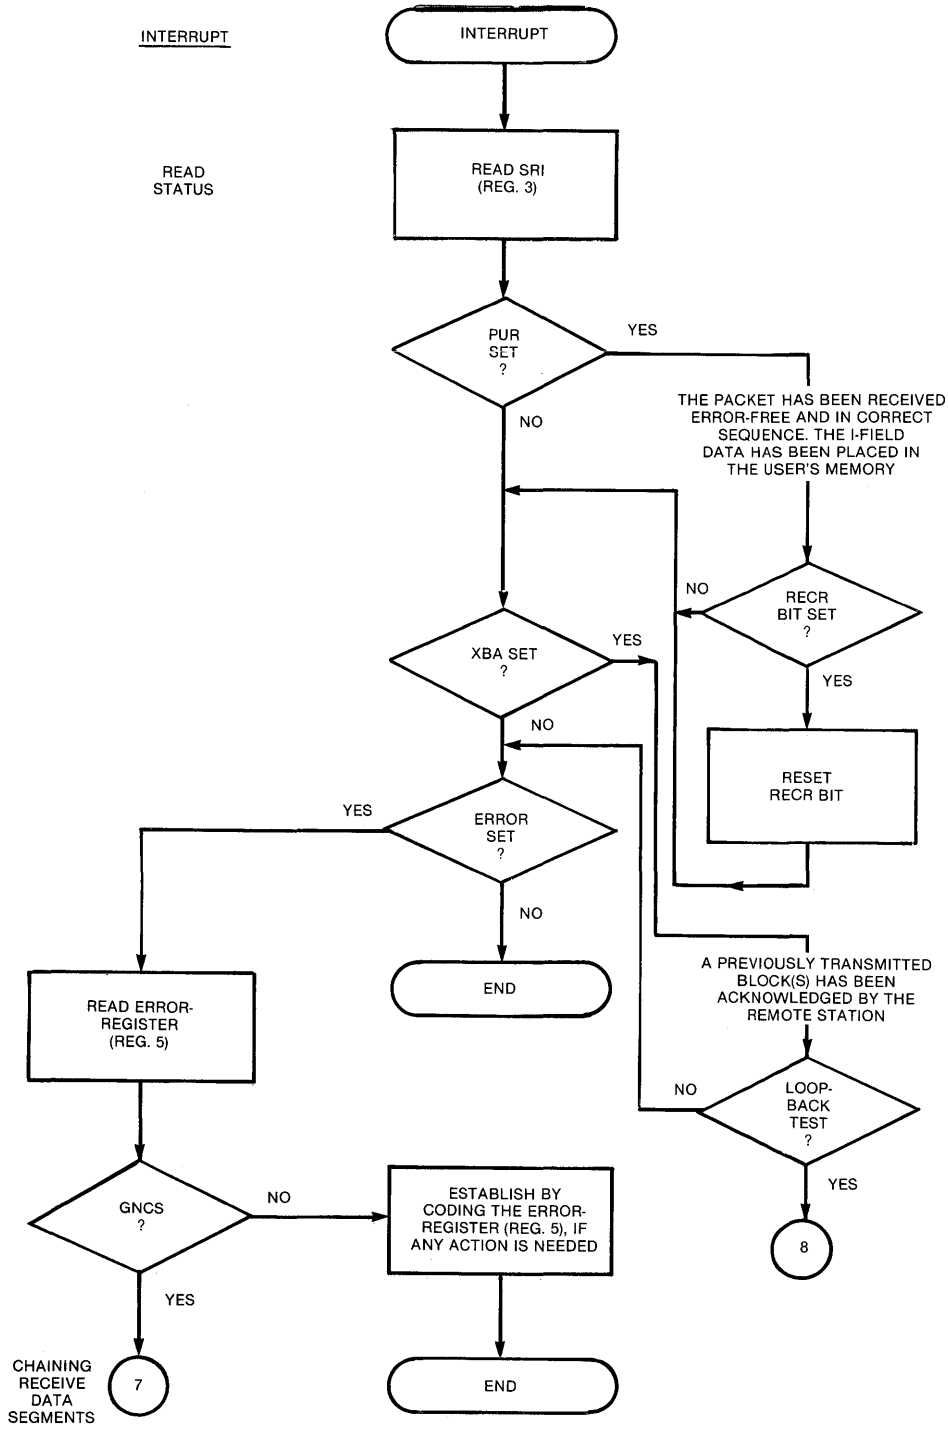
INTERRUPT
INTERRUPT
READ STATUS
READ SRI (REG. 3)
PUR SET ?
YES
NO
THE PACKET HAS BEEN RECEIVED ERROR-FREE AND IN CORRECT SEQUENCE. THE I-FIELD DATA HAS BEEN PLACED IN THE USER'S MEMORY
RECR BIT SET ?
NO
YES
XBA SET ?
YES
NO
RESET RECR BIT
ERROR SET ?
YES
NO
END
READ ERROR-REGISTER (REG. 5)
A PREVIOUSLY TRANSMITTED BLOCK(S) HAS BEEN ACKNOWLEDGED BY THE REMOTE STATION
LOOP-BACK TEST ?
NO
YES
8
GNCS ?
NO
ESTABLISH BY CODING THE ERROR-REGISTER (REG. 5), IF ANY ACTION IS NEEDED
YES
CHAINING RECEIVE DATA SEGMENTS
7
END

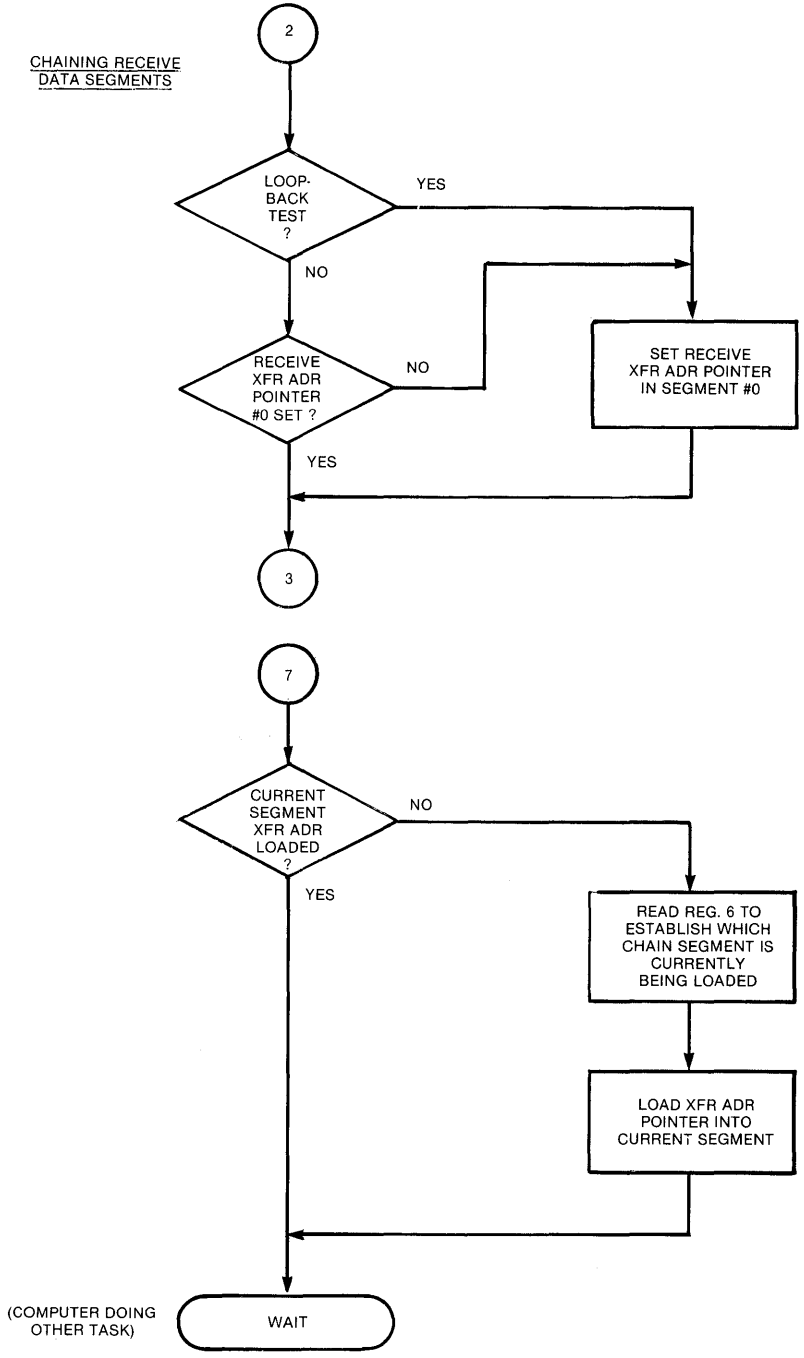
CHAINING RECEIVE DATA SEGMENTS
2
LOOP-BACK TEST ?
YES
NO
RECEIVE XFR ADR POINTER #0 SET ?
NO
SET RECEIVE XFR ADR POINTER IN SEGMENT #0
YES
3
7
CURRENT SEGMENT XFR ADR LOADED ?
NO
YES
READ REG. 6 TO ESTABLISH WHICH CHAIN SEGMENT IS CURRENTLY BEING LOADED
LOAD XFR ADR POINTER INTO CURRENT SEGMENT
(COMPUTER DOING OTHER TASK)
WAIT

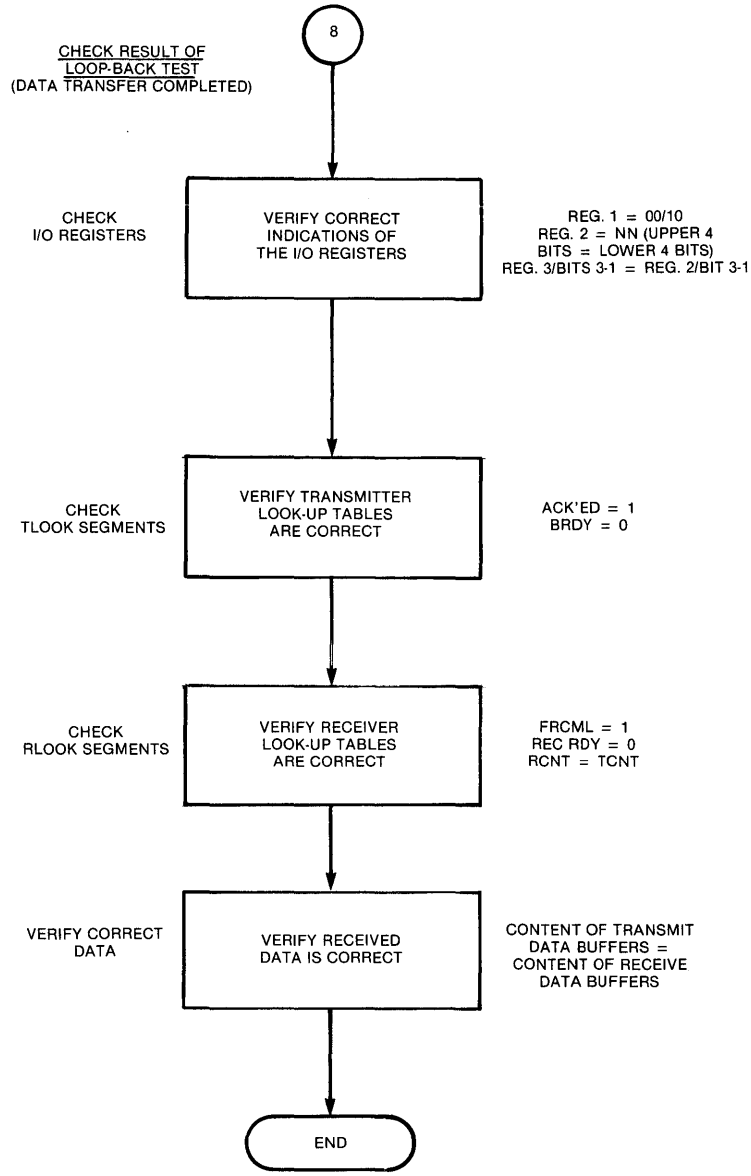
8
CHECK RESULT OF LOOP-BACK TEST (DATA TRANSFER COMPLETED)
CHECK I/O REGISTERS
VERIFY CORRECT INDICATIONS OF THE I/O REGISTERS
REG. 1 = 00/10
REG. 2 = NN (UPPER 4 BITS = LOWER 4 BITS)
REG. 3/BITS 3-1 = REG. 2/BIT 3-1
CHECK TLOOK SEGMENTS
VERIFY TRANSMITTER LOOK-UP TABLES ARE CORRECT
ACK'ED = 1
BRDY = 0
CHECK RLOOK SEGMENTS
VERIFY RECEIVER LOOK-UP TABLES ARE CORRECT
FRCML = 1
REC RDY = 0
RCNT = TCNT
VERIFY CORRECT DATA
VERIFY RECEIVED DATA IS CORRECT
CONTENT OF TRANSMIT DATA BUFFERS = CONTENT OF RECEIVE DATA BUFFERS
END

4.1 INTERNAL LOOP-BACK TEST (Example 1)

The loop-back test feature is an internal programmable loop-back of data, enabling the user to make an almost complete test of the WD2501. It allows diagnostic testing of the WD2501 and the interfacing circuitry. In this mode, transmitted data to the TD pin is internally routed to the received data input circuitry, thus allowing this WD2501 to set-up a link, send a number of packets to itself and then reset the link.

The $\overline{RC}$ clock is internally connected to $\overline{TC}$ clock. $\overline{CTS}$ input however, must be connected externally to GND or the $\overline{RTS}$ output.

The loop-back test allows the verifying of proper operation of practically all the various functions of the WD2501. The features tested here, the addresses and values of the variables chosen are only used as examples and are as follows:

TLOOK segments starting address = 0800H
Transmit Data buffer no.0 starting address = 1000H
Received Data buffer no.0 starting address = 1800H
Number of packets transferred = 1
Number of I-field bytes per packet = 1024
Number of residual bits = 0
T1 = 101H
N2 = 20H

Chaining is used in this example. The 1024 bytes are divided into 256 byte chain segments. Five segments are needed for this operation with 254 bytes of I-field data and two XFR ADR bytes per segment in the first four chain-segments. The rest of the I-field data (8 bytes) are located in the fifth chain-segment.

Programming:

CHAIN = 4 = number of CHAIN segments − 1
LIMIT = 3 = (number of bytes per segment divided by 64) − 1

For buffer management programming, see memory access scheme in Figure 7.

XMIT Command Address and XMIT Response Address (REG. E and F) must be the same value.

In some applications, it is necessary to keep the $\overline{RTS}$ signal to the modem in the Off condition during internal loop-back test. Also, to accomplish the most complete test, $\overline{RTS}$ output should be connected to $\overline{CTS}$ input externally (not done internally). Figure 6 shows one example of how to implement these two functions. The ILOOP signal is connected directly to a PIO output.

In the loop-back test example shown in this section, the logic in Figure 6 is used and contains the Z80 CPU, programmable I/O (PIO) etc., as shown in Figure 3.

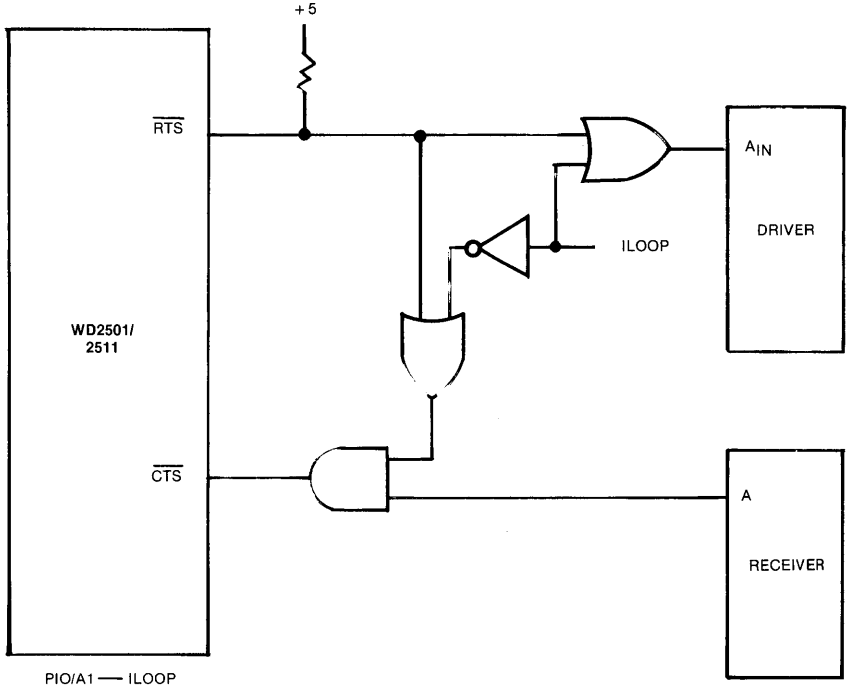

Figure 6. LOGIC FOR INTERNAL LOOP-BACK TEST

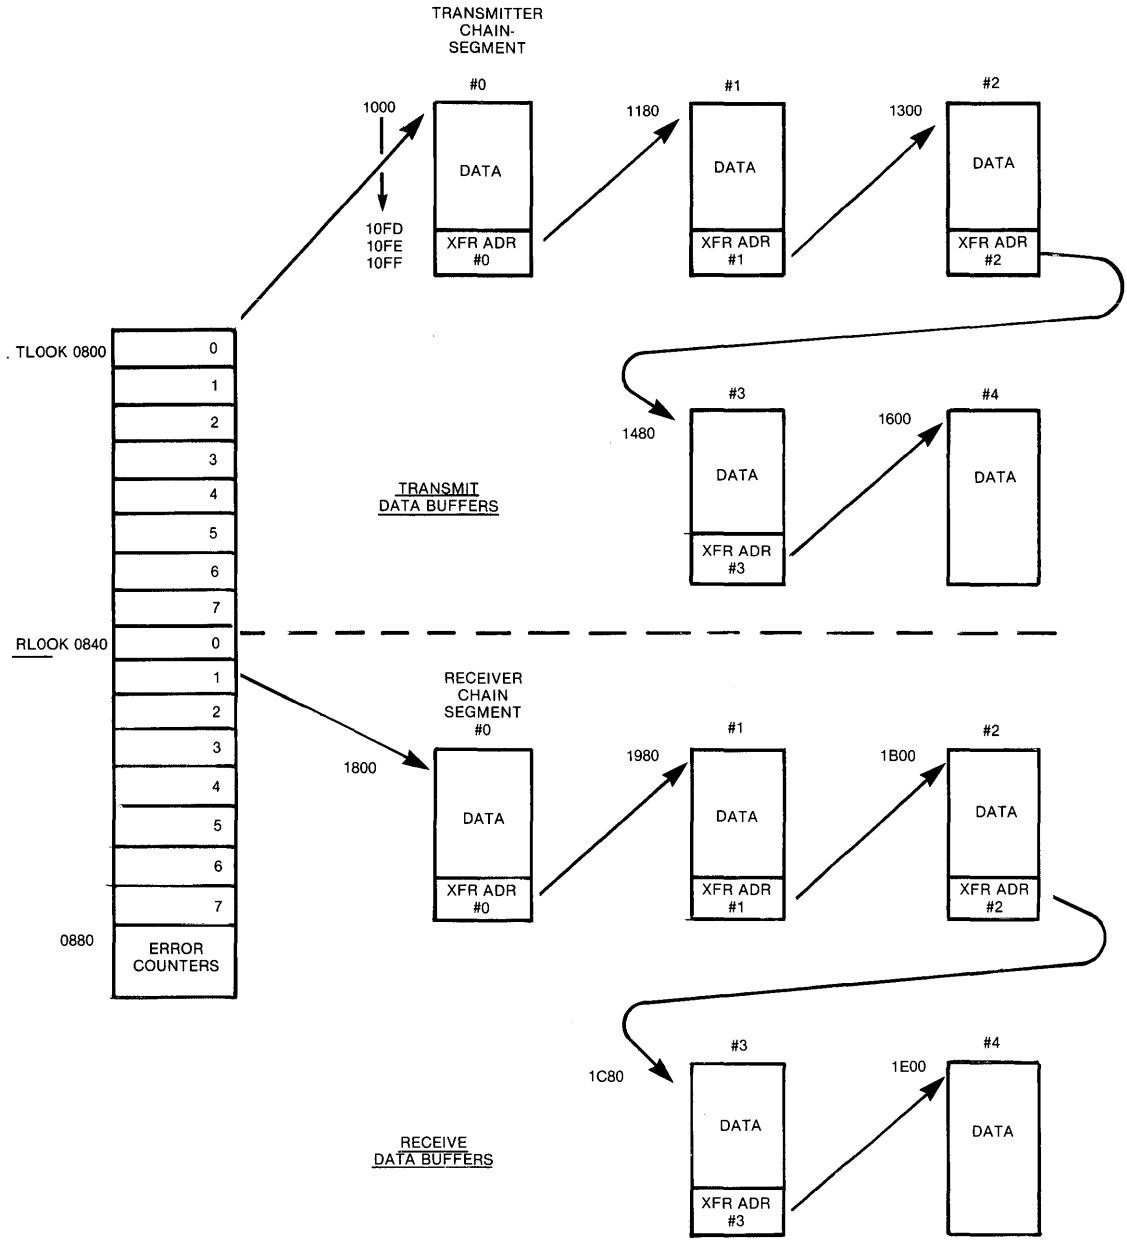

Figure 7. MEMORY ACCESS SCHEME FOR LOOP-BACK TEST (Example 1)

APPENDIX A

GLOSSARY OF DATA COMMUNICATIONS TERMS

The following is a list of industry-accepted data communications terms that are applicable to this specification.

ABM	Asynchronous Balanced Mode
ADCCP	Advanced Data Communications Control Procedure (ANSI BSR X3.66)
ANSI	American National Standards Institute
ARM	Asynchronous Response Mode
CCITT	International Consultative Committee for Telegraphy and Telephony
CMDR	Command Reject. A U-Frame
DCE	Data Circuit Termination Equipment (the network side of the DTE/DCE link)
DISC	Disconnect. A U-Frame
DTE	Data Terminal Equipment
DM	Disconnect Mode. A U-Frame (LAPB, only)
ECMA	European Computer Manufacturers Association
FCS	Frame Check Sequence
FDX	Full Duplex (also called "two way simultaneous")
FRAME	Basic Serial Block of Bit-Oriented Data. Includes leading and trailing flags, address field, control field, FCS field, and an optional information field.
FRMR	Frame Reject. A U-Frame (LAPB, only)
HDLC	High-Level Data Link Control (ISO 3309)
HDX	Half Duplex (also called "two way alternate")
HOST	Another name for a DTE
I-Frame	Information Frame. Control field bit 0 is 0. In X.25 an I-frame is a packet.
ISO	International Standards Organization
LINK	The logical and physical connection between two data terminals
LAP	Link Access Procedure
LAPB	Link Access Procedure Balanced
N2	Maximum number of retransmissions of a frame. (Also called retransmission count variable.)
NODE	Another name for a DCE.
N(R)	Sequence number of next frame expected to be received.
N(S)	Sequence number of current frame being transmitted.
OCTET	An 8-bit byte
P(R)	Receive Data, Packet count of next packet expected to be received
P(S)	Send Data, Packet count of current packet being transmitted
PACKET	An I-Frame in X.25
PAD	Packet Assembly/Disassembly
REJ*	Reject. An S-Frame
RNR*	Receiver Not Ready. An S-Frame
RR*	Receiver Ready. An S-Frame
S-Frame	Supervisory Frame. Control field bit 0 = 1 and bit 1 = 0
SARM	Set Asynchronous Response Mode. (LAP, only)
SABM	Set Asynchronous Balanced Mode. (LAPB, only)
SDLC	Synchronous Data Link Control (IBM document GA27-3093)
SNAP	Standard Network Access Protocol (Trans-Canada)

T1	A Primary Timer for a delay in waiting for a response to a frame
U-Frame	Unnumbered Frame. Control Field bit 0 = 1 and bit 1 = 1
UA	Unnumbered Acknowledge. A U-Frame
X.25	Recommendation by CCITT on Packet Switching Networks
X.3, X.28, X.29	Recommendations by CCITT involving PAD facilities
ACS	Advanced Communications Service (AT&T)

*There are also RR, RNR, and REJ packets which are not the same as the S-frame RR, RNR and REJ discussed in this document.

APPENDIX B

THE DIFFERENCE BETWEEN LAP AND LAPB

In March 1976, the CCITT adopted Recommendation X.25 as an interface standard for public packet-switching networks. The link level procedure adopted was called Link Access Procedure (LAP) and used the HDLC Asynchronous Response Mode (ARM). However, LAP was not fully in conformity with HDLC (ADCCP is essentially the same. See ANSI X3.66, 1979). Therefore, in 1977, when Provisional Recommendation X.25 was adopted, a procedure called LAPB was added. LAPB is Link Access Procedure-Balanced and operates under the HDLC Asynchronous Balanced Mode (ABM). Unfortunately, the 1977 LAPB lacked good symmetry between the DTE and DCE, and was unworkable.

In the April 1979 CCITT meeting, the LAPB was greatly enhanced, especially in the DTE/DCE symmetry. This enhanced version was approved in the February 1980 Plenary meeting of the CCITT. We now have a good, workable LAPB standard.

The LAP, because of not fully complying with HDLC, has subtle problems. LAPB is a superior procedure. The usage of LAP will eventually be replaced with LAPB.

LAP and LAPB are different only in the link set-up, disconnect, reset and receiver-not-ready procedures.

1.0 Link Set-Up

In LAP, a link is set-up when both ends exchange a SARM and UA. In LAPB, only one end will send a SABM, and when the other end sends a UA, the link is up.

2.0 Link Reset

A link, which is up, may be reset.

In LAP, a link may be reset in one I-frame direction by sending a SARM. When the other end acknowledges the SARM with a UA, the end that send the SARM clears its send sequence number. However, I-frames coming from the other direction are not reset. This could be a problem. Suppose a link is up, but one end momentarily loses power. When that end tries to set-up the link by sending a SARM, the other end sends a UA, but "thinks" the SARM was a link reset. Thus, a SARM is never returned. (The WD2501 gets around this by sending a DISC before trying to bring a link up.)

In LAPB, a link reset SABM resets the link in both directions. The ambiguity between a reset command and a link set-up command is much less than for LAP. (It is still a good practice to send a DISC before trying to bring a link up in LAPB. This will insure that Level 3 software is fully aware of a disconnect situation, as opposed to a reset.)

3.0 Link Down

In LAP, if a link is down, a station may respond only to received SARM's or DISC's. The response is a UA. Any other received command is disregarded.

In LAPB, a down station may respond to a SABM or DISC with a UA. When receiving any other command frame with the P-bit set to 1, the DCE will transmit a DM response with the F bit set to 1. This is an advantage in a down station which does not initiate link set-up. If the other station "thinks" the link is up, it will eventually send a packet with P = 1. When this happens, it is informed, by the DM response, that the link is down.

4.0 Supervisory Commands

In LAP, S-frames are responses, only.

In LAPB, an S-frame may be a command or response.

Suppose two stations are operating over a set-up link (station A and station B). Suppose B sends an RNR. In LAP, the only way that A can test B to see if B is ready again is to send the last unacknowledged I-frame with P = 1 at T1 intervals. This forces B to reply with either an RNR or RR (F = 1). An I-frame is used here as a request for status, and it could be a long I-frame.

In LAPB, station A could send an RR command with P = 1 at T1 intervals. B would respond with either an RNR or RR (F = 1), and this procedure is more efficient.

5.0 CMDR/FRMR

Suppose a command or response is received error-free (FCS good), but there is an invalid condition. This could be an acknowledgment to a packet never sent, an unrecognized command or response, an I-field where one is not allowed, etc.

In LAP, received invalid commands may be rejected except for packets with an invalid N(R). The receiving station would send a command reject (CMDR) along with the rejecting cause and its current V(R) and V(S) counts. These counts allow for re-initialization at a higher level to prevent duplicate packet transmission (i.e., sending an already acknowledged packet). However, invalid responses or packets with an invalid N(R) may not be rejected. Only a link resetting SARM may be sent. Thus, the reason for the reject is not communicated to the other end. Even worse, duplicate packets could be exchanged since the V(R), V(S) counts are not transferred.

In LAPB, all invalid responses and commands are rejected with a frame reject (FRMR). An FRMR is the same as a CMDR except for one bit which tells whether the rejected frame was a response or command. Thus, the rejecting cause and V(R), V(S) counts are communicated for both invalid commands and responses.

LAP COMMANDS AND RESPONSES

(Bit 0 is transmitted first) Only the CMDR and I-frame contain I-fields

FRAME TYPE	COMMAND	RESPONSE	CONTROL FIELD BIT # 7 6 5	4	3 2 1	0	
I-FRAME	I-FRAME		N(R)	P	N(S)	0	
S-FRAME		RR	N(R)	F	0 0 0	1	RECEIVER READY
		RNR	N(R)	F	0 1 0	1	RECEIVER NOT READY
		REJ	N(R)	F	1 0 0	1	REJECT
U-FRAME	SARM		0 0 0	P	1 1 1	1	SET ASYN-CHRONOUS RESPONSE MODE
	DISC		0 1 0	P	0 0 1	1	DISCONNECT
		UA	0 1 1	F	0 0 1	1	UNNUMBERED ACKNOWL-EDGE
		CMDR	1 0 0	F	0 1 1	1	COMMAND REJECT

LAPB COMMANDS AND RESPONSES

FRAME TYPE	COMMAND	RESPONSE	CONTROL FIELD BIT # 7 6 5	4	3 2 1	0	
I-FRAME	I-FRAME		N(R)	P	N(S)	0	
S-FRAME	RR	RR	N(R)	P/F	0 0 0	1	RECEIVER READY
	RNR	RNR	N(R)	P/F	0 1 0	1	RECEIVER NOT READY
	REJ	REJ	N(R)	P/F	1 0 0	1	REJECT
U-FRAME	SABM		0 0 1	P	1 1 1	1	SET ASYN-CHRONOUS BALANCED MODE
	DISC		0 1 0	P	0 0 1	1	DISCONNECT
		DM	0 0 0	F	1 1 1	1	DISCONNECT MODE
		UA	0 1 1	F	0 0 1	1	UNNUMBERED ACKNOWL-EDGE
		FRMR	1 0 0	F	0 1 1	1	FRAME REJECT

Only the FRMR and I-frame contain I-fields
P = Poll Bit F = Final Bit

APPENDIX C

OPERATING THE WD2501 OR WD2511 IN A MULTIPOINT CONFIGURATION

A typical multipoint consists of a primary station controller connected to one, or more, secondary controllers by means of a four wire connection as shown in Figure 1. One wire pair carries serial data from the primary for broadcast to all secondaries. One wire pair carries serial data from all secondaries to be received by the primary. Thus, the primary communicates with secondaries, and secondaries never communicate directly with one another. Also, the primary can only "talk" to one secondary at a time.

ADCCP (Advanced Data Communication Control Procedures, ANSI X3.66) specified three modes: Normal Response Mode (NRM), Asynchronous Response Mode (ARM), and Asynchronous Balanced Mode (ABM). NRM is strictly an association between a primary and one, or more, secondaries. Therefore, NRM has been specifically designed for use with multipoint, and is the best choice for multipoint. ABM and ARM are not as well suited for multipoint as NRM, but can none-the-less be used in multipoint.

The WD2501 is LAP which is an application of ARM. The WD2511 is LAPB which is an application of ABM. For multipoint, the primary will establish a link with one secondary, and communicate I-frames with that secondary. After this "session," the primary will initiate link disconnect, and go on to another secondary.

The control of the WD2501 and WD2511 is by means of Registers E and F which control the command/response definitions for the A field. This may be handled in one of two ways: First, make Register E the same for all secondaries, and this value will be used in Register F of the primary. Register F will be unique for each secondary, and the primary will modify Register E for each session with an individual secondary. Second, Register E and F could be equal in each secondary, but unique for that secondary. The primary would modify Registers E and F for each session with a given secondary. Neither of the above methods appears to have an advantage over the other.

By way of example, suppose the second method is chosen, and suppose there are five secondaries. Register E and F for the five will be 1, 3, 5, 7, and 9. Each secondary will be in an idle state with CR0 bits 5 and 1 set (hex 22). The $\overline{RTS}$ pin will be off for the secondaries. As soon as a secondary receives a DISC with A field which matches Register F, $\overline{RTS}$ will be asserted. If the A field does not match Register F, the frame is discarded, and $\overline{RTS}$ remains off.

When the primary needs to communicate to #1, Register E and F in the primary are set to 1. The primary will set CR0 bits 4 and 1 (hex 12). The primary will initiate a link set-up procedure wih #1. After the link is brought up with #1, there will be a link-up interrupt with ER0 = hex 21 in both the primary and secondary. If, however, #1 is off-line, there will be no link-up time-out from the WD2501/2511 for failure to bring-up a link. The user must provide this time-out.

After the link is up, if a WD2511 is being used, the user may wish to clear CR0 bit 5 (does not need to do so for the WD2501). For the WD2511, this is a half duplex bit, and will become full duplex when the bit is cleared.

To discontinue a session, the primary will set bit 0 in CR0 (hex 01). This will cause a DISC to be generated to the secondary. The user may change Registers E and F for the next secondary after waiting a few milliseconds (2 or 3) after the $\overline{LINK}$ bit becomes set.

When the secondary receives the DISC, the secondary will generate a disconnect interrupt (ER0 = hex 30), and transmit a UA.

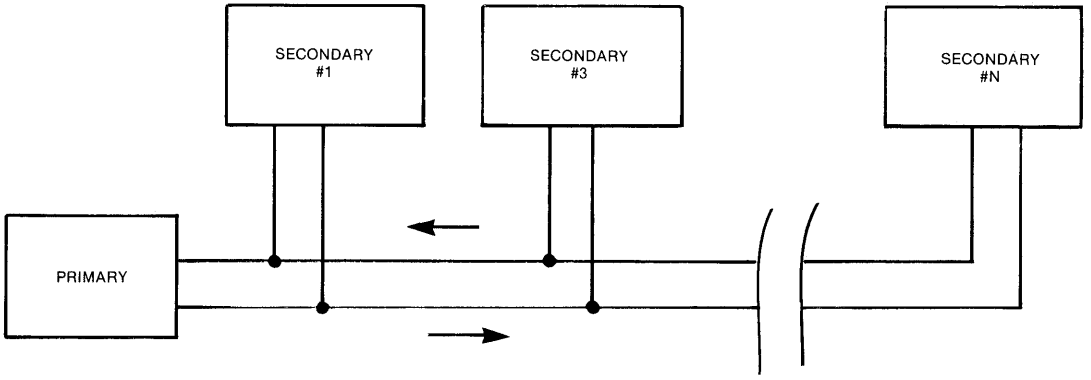

Figure 1. TYPICAL MULTIPOINT CONFIGURATION

Printed in U.S.A.

50 YEARS Electronics®
A McGRAW-HILL PUBLICATION

WD2511 LSI circuit simplifies packet-network connection

by Geary L. Leger, *Western Digital Corp., Newport Beach, Calif.*

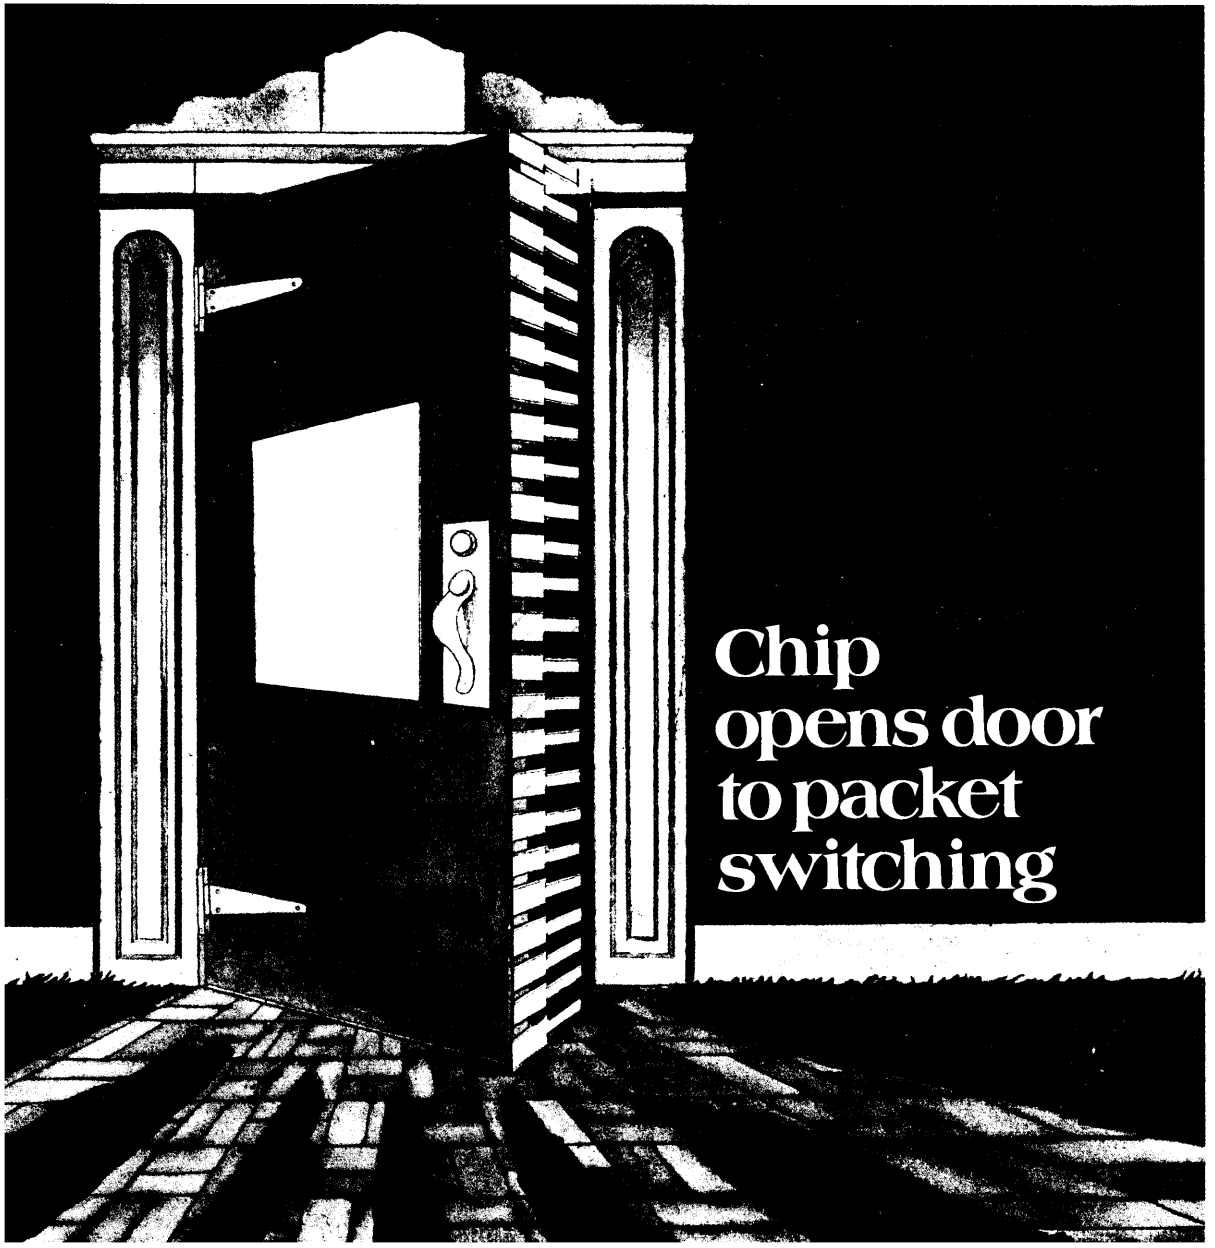

LSI ready to make a mark on packet-switching networks

New chip's link-control capabilities ease connection to terminals: Part I

by Geary L. Leger, *Western Digital Corp., Newport Beach, Calif.*

☐ Packet-switching networks are prime targets for the application of large-scale-integrated circuit technology. In fact, sometime during the first quarter of next year, this useful and expanding approach to data communications will have its first dedicated LSI circuit, one designed to take advantage of LSI's potential for lower cost and greater reliability.

The circuit is the Micro Packet Interface chip, or μPAC, being developed by Western Digital. It will handle Level 2 control of the link between a data terminal and a network node as set forth in the X.25 protocol established by the Consultative Committee for International Telephony and Telegraphy (CCITT). Because these

This is the first of two articles. It deals with the overall characteristics of packet-switching networks. Part 2, which starts on page 95, describes an LSI chip being developed for Level 2 control per the X.25 protocol for packet networks.

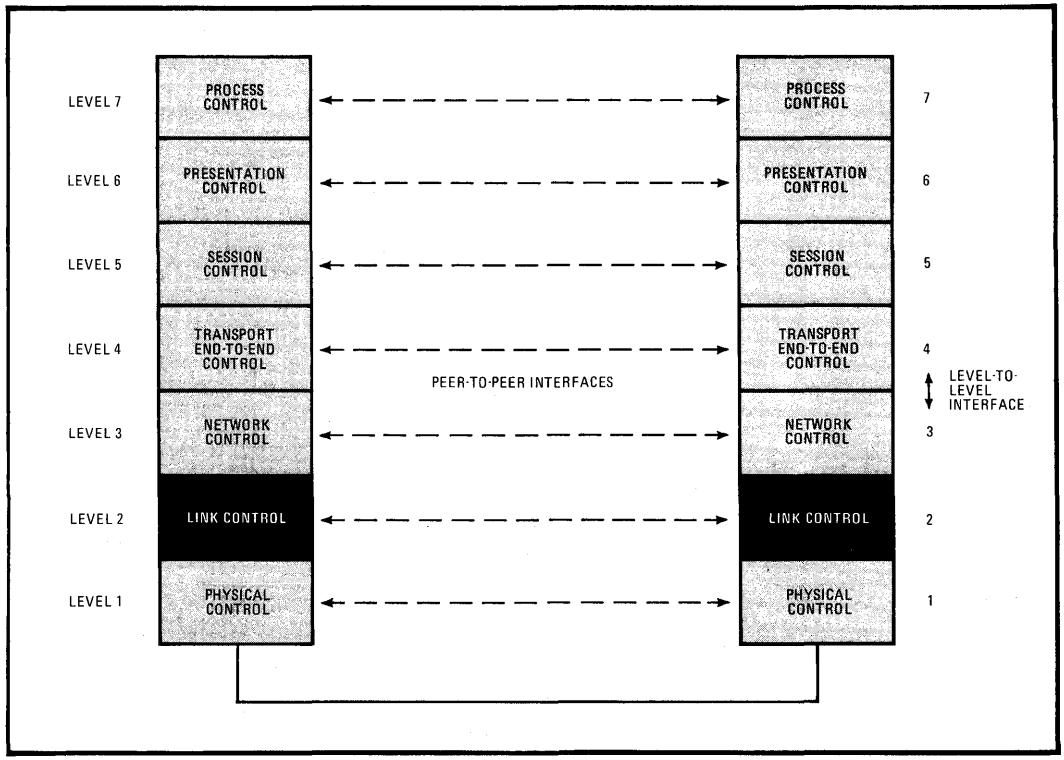

1. Layered architecture. Independence among system levels allows changes to be made to one level without disrupting the operation of other levels; and adjacent level is affected only if the changes affect the interface to that level. Standards apply to peer-to-peer interfaces.

Recent efforts in packet switching

The security, survivability, and economic advantages of packet-switching data-communications networks have not gone unnoticed in either the military or the corporate sectors. Though still in its infancy, this type of communication is growing rapidly.

According to Defense Advanced Research Project Agency director Eugene H. Kopf, the latest military packet effort aims to find the optimum architecture for a command and control network of 5,000 to 10,000 small packet radio relay terminals whose purpose would be to insure survivable control over strategic weapons. The network voice and data-packet radios would provide line-of-sight communications throughout the continental U. S. after an attack.

The civilian sector has continued the development of packet networks from their modest beginnings. In 1972, Bolt Beranek & Newman Inc. founded Telenet, a public packet-switching network taken over this year by General Telephone & Electronics Corp. GTE Telenet Communications Corp., Vienna, Va., completed installation this month of a packet-switching exchange in San Juan, Puerto Rico, for ITT World Communications Inc. The new service allows businesses and industrial organizations to link to the ITT gateway and transmit and receive data over shared transmission lines to data terminals or computers on the U. S. mainland.

Tymnet Inc., Cupertino, Calif., is the largest public packet-switching network in the U. S. It has so many customers that it issues a 34-page directory describing 200 data bases accessible through its network. Tymnet now serves 250 computers in the U. S.; it recently added New Zealand to its list of countries served, bringing the total to 26.

The Japanese have not been idle in adapting packet technology to their needs. Nippon Telegraph and Telephone Public Corp. started work on its digital data-exchange system in 1971 and had installed packet-network equipment in seven cities by late 1978. This commercial packet-switching network (called D50) is expected to go into full service this year. D50 conforms fully to Recommendation X.25 of the Consultative Committee for International Telephony and Telegraphy.

The packet network industry has come a long way since the first operational system (Arpanet) was installed by the Defense Advanced Research Projects Agency in 1969. In 1976, the CCITT adopted X.25 as a standard three-level protocol for interfacing terminals to public packet networks—a major step for the industry. Even though it has been criticized as being too complicated, X.25 has stimulated interest in packet networks.

However, packet switching is not the answer to all data and voice communications problems, as some have claimed. Gino J. Coviello of the Defense Communications Agency in Arlington, Va., concluded in a recent study that the number of channels traversing a particular transmission link and the network topology and architecture have a significant impact on the cost-effectiveness of a packet-switching network. Ray W. Sanders, president of Computer Transmission Corp., says, "Packet switching will take its rightful place alongside circuit switching." A hybrid approach combining features of both circuit and packet switching "provides the best of all possible worlds," according to Sanders. **-Harvey J. Hindin**

networks are relatively new, familiarity with the Level 2 link control and other details of their operations is not widespread. Yet the purpose of the μPAC is intelligible only in the context of such an understanding.

To date, most data-communications systems use circuit-switching techniques. A physical circuit is assigned either permanently (a private, leased line) or for the duration of the call (a dial-up line). But of a given line's total available time, only a small percentage is actually taken up by data transmission. A system for dynamic allocation of the physical circuits, in contrast to static circuit-switching allocation, requires the logic and memory capabilities of computers.

Prior to the late 1960s, static circuit allocation was more economical than using computers in a dynamic allocation system. The low cost of today's minicomputers and microprocessors and the dramatic drop in the cost of memory, however, make dynamic allocation more economically feasible in many cases. It is most suitable in multipurpose applications—digital communications systems linking various types of data terminals such as facsimile machines, computerized data bases, interactive keyboard printers, or cathode-ray-tube terminals.

Historically, communications systems have been developed to satisfy one application at a time. The wide variety of computers, terminals, and technologies has led to the development of many incompatible networks. A time-sharing network may connect many asynchronous interactive keyboard printers on dial-in lines at 110 or 300 bits per second. A clustered CRT application such as IBM's 3270 may operate synchronously at 1,200, 2,400, or 4,800 b/s.

The incompatibility of different types of equipment and the protocols they use for communicating greatly reduces the reliability and efficiency of data communications as a whole. A single corporation may, for example, use several incompatible networks.

Security and survivability

The technique of sending a digital data in short packets, rather than in a continuous stream, was first suggested by Paul Baran of the Rand Corp. more than a decade ago. The packets are transmitted between intermediate points in the network, called nodes, or DCEs, for data-circuit-terminating equipment (see table).

This dynamic-allocation technique has two major inherent advantages over circuit-switching methods. It increases data security, since the message is broken up: all the packets would have to be picked up and combined by an intruder before he could use the data. And system survivability and reliability are enhanced by the large number of linked nodes. Alternate routes will get a message through if some of the nodes or links are malfunctioning or destroyed. The security and survivability are of great interest to the military. The commercial sector is also developing this type of system (see "Recent efforts in packet switching," above).

The problem of equipment and communication-proto-

col incompatibility is an international one, since data communications is international. This question is addressed at the global level by the CCITT (see "Setting the standard," p. 93).

The difficulty of expanding, modifying, or upgrading existing data-communications networks is another problem that is tackled by the new packet-switching systems. Any communications system represents a large capital investment, and down time can be disastrously expensive. A user cannot simply tear down an old network and substitute a new one with the latest advances. New terminals and technologies must be phased into the existing structure without interrupting operation. This calls for a degree of system flexibility.

Layered architecture

Packet-switching networks achieve this flexibility through layered (or multilevel) architecture [*Electronics* May 24, 1979, p. 111]. Several standards organizations have been working on the specifics of this concept.

The importance of layered structure is easy to understand. Suppose Mr. Jones, an executive, wishes to talk to Ms. Smith, another executive. Jones (level 4) tells his secretary (level 3) to get Smith on the line. Jones' secretary dials the number (level 2). An electromechanical switching mechanism connects the two phones (level 1). Smith's phone rings (level 2). The two secretaries converse (level 3) and pass on the information that the call is ready to their bosses. Smith and Jones now communicate (level 4).

Jones was never concerned with the electromechanical switching mechanism, nor with Smith's telephone number; he was primarily concerned with talking to Smith (peer to peer) and secondarily with talking to his secretary to get the call set up.

Multilevel communications systems are structured in a similar fashion. The protocol standards are prepared for connection in the peer-to-peer layer. Standards do not define the interface between adjacent layers. This is intentional: terminal manufacturers are thus left free to design the adjacent-layer interface in their own way. This enhances system flexibility. If a layer is changed or upgraded, nonadjacent layers are not affected.

X.25 defines and standardizes three levels. There are as many as four more definable levels (Fig. 1), but much work remains to be done to standardize these higher levels. Level 1 may be viewed as a data-exchange mechanism serving Level 2. Level 2 is a data-exchange mechanism serving Level 3, and so on.

Three standardized levels

Level 1 concerns itself with the link's physical interfaces. Level 2 deals with link control. It includes setting up and disconnecting a link, the control of flow between data generators and data receivers, and bit-oriented frame structure. The μPAC from Western Digital is designed to perform the Level 2 functions. Level 3, network control, includes the procedures for establishing and disconnecting the virtual circuit and for controlling the flow of data packets in the network.

In a packet network, the sender or receiver has a terminal (commonly called DTE, for data-terminal equipment) with a distinct address. Part of the gear at a network node might also be called data-terminal equipment. The packets of data are transferred from node to node and finally to the receiver's terminal.

When a node receives a packet, it stores the packet, decides where and when to forward it on the basis of the packet's destination and priority and the load conditions of the network, and then does so. This store-and-forward facility is the key to the network's ability to allocate circuits dynamically. Packets going from terminal A to terminal B in Fig. 2 could follow the node path 7-12-11-10, 7-5-3-1-10, or any of a number of others.

Dynamic routing within the network is transparent to the users at their terminals. The path data takes is called a virtual circuit between A and B: the terminals communicate as if a dedicated circuit joined them. In order to

PACKET-SWITCHING NETWORK TERMINOLOGY

ADCCP	Advanced Data-Communications Control Procedure (ANSI X3.66)
ANSI	American National Standards Institute
ATDM	asynchronous time-division multiplexing
BOP	bit-oriented protocol
CCITT	Consultative Committee for International Telephony and Telegraphy
DCE	data-circuit-terminating equipment (network node)
DTE	data-terminal equipment (user's terminal)
FCS	frame check sequence
HDLC	High-level Data-Link Control protocol (ISO 3309)
ISO	International Standards Organization
ITU	International Telecommunications Union
I frame	information frame, known as a packet under X.25
LAP	Link-Access Procedure (X.25)
LAPB	Link-Access Procedure, Balanced (X.25)
Link control	X.25 Level 2 control for linking DTE and DCE, including link initialization, establishment, and disconnection, and control of data flow on the link
Network control	X.25 Level 3 control of virtual circuits in network, including circuit establishment, disconnection, and reset, and the control of packet flow
N1	maximum number of bits in a packet
N2	maximum number of command retransmissions
PAD	Packet Assembly/Disassembly facility (defined in CCITT recommendations X.3, X.28, and X.29)
Physical interface	X.25 Level 1 specifications for the physical connection of DTE and DCE, including electrical parameters and transmission rate
S frame	supervisory frame
SDLC	Synchronous Data-Link Control protocol
T1	time minimum before retransmission of unacknowledged command
U frame	unnumbered frame
X.25	CCITT recommendation for packet-switching network protocols (others include X.3, X.28, and X.29)

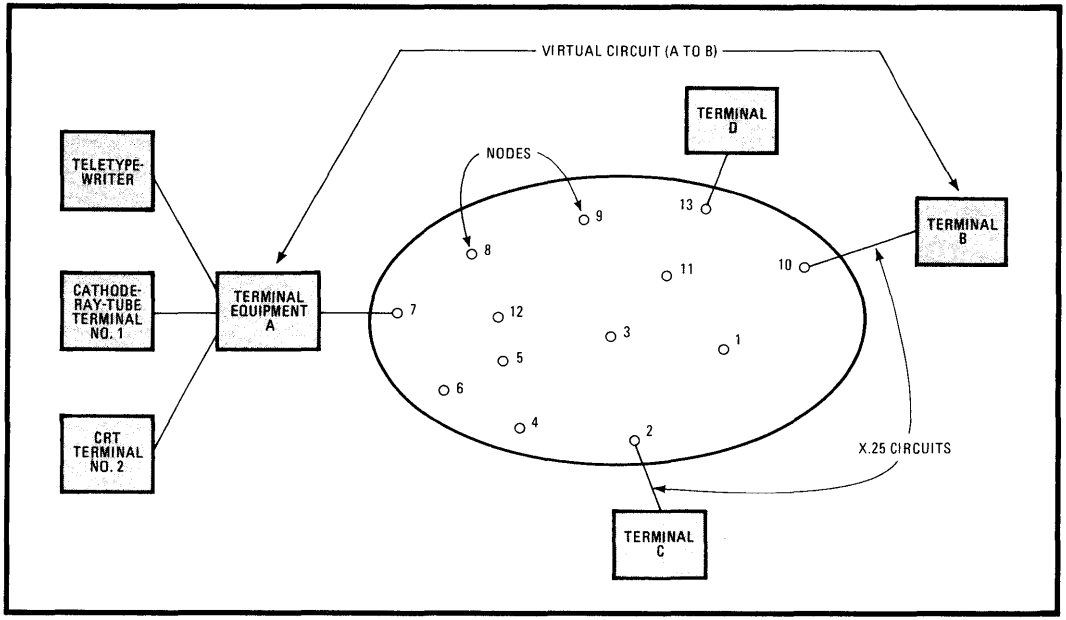

2. Many possible paths. The user of a packet-switching network at his terminal sees no difference between the virtual circuit and an ordinary physical link. Network control may send the data packets through a changing series of nodes as system traffic conditions change.

establish a virtual circuit, terminal A transmits a call-request packet that includes the caller's address and the address of terminal B, the destination. Terminal B accepts the request by returning a call-accepted packet to A, and the virtual circuit is set up.

Circuit sharing

Several simultaneously active virtual circuits can be set up by interleaving packets. This asynchronous time-division multiplexing (ATDM) exploits the fact that a typical virtual circuit carries data for only a small percentage of the time it is set up. It differs from other time-division multiplexing schemes in that a dedicated time slot is not provided for each virtual circuit being multiplexed.

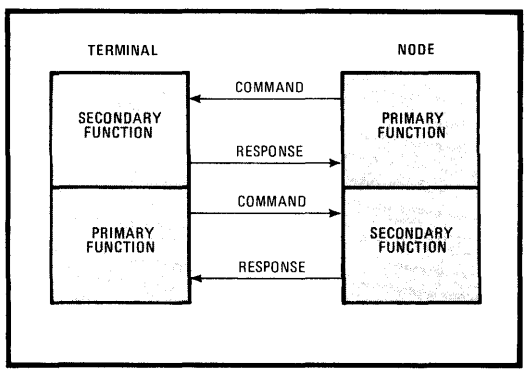

3. Addressable functions. A terminal or node has a primary function that sends commands and receives responses. Its secondary function, which has a different address, responds to received commands. Arrows represent system logic, not physical wires.

In the multilevel packet-switching architecture, Level 2 (also known as the link level or the frame level) involves the point of contact between the subscriber's terminal and the network node it is linked to directly.

Each station, be it terminal or node, has two logical functions needed for addressing and signal implementation, called primary and secondary (Fig. 3). The primary function transmits commands and receives responses; the secondary function does the reverse—it receives commands and transmits responses.

The structure of the data frames used for this communication is common to all bit-oriented protocols (BOPs)—High-level Data-Link Control (HDLC), the essentially similar Advanced Data-Communication Control Procedures (ADCCP), and the Synchronous Data-Link Control (SDLC) protocol worked out by IBM [*Electronics*, Jan. 18, 1979, p. 137]. The Level 2 protocol defined by X.25 is an outgrowth of HDLC.

The frame is simply a block of serial data exchanged between two terminals or a terminal and a node. It consists of a flag, an address field (or A field), a control field (or C field), an information field (or I field), a frame-check sequence (FCS), and another flag. Depending on the frame type, the information field may or may not be included.

There is a flag at either end of a frame; a single flag may close one frame and open the next. Data transparency is provided within the frame by the transmitting station: a logic 0 is inserted after all sequences of 5 contiguous logic 1 bits, so that no transmitted data is inadvertently read as a flag, which has the binary form

Setting the standard

International interface standards are vital to the development and growth of packet-switching networks. Standards lead to lower costs for equipment bought by network users, since this equipment can be manufactured in much larger quantities. The user also benefits from the interchangeability of gear from different vendors. Manufacturers reap the rewards of a global market rather than a local one, and network organization is made vastly easier.

A number of U. S. and international standards organizations are working together to set up interface rules. The International Telecommunications Union (ITU), formed in 1865, operates under the auspices of the United Nations. Under the ITU is the Consultative Committee for International Telegraphy and Telegraphy (CCITT), which is primarily an organization of carriers.

Study Group VII is a CCITT organization that handles public data networks. SG VII is responsible for publishing a number of standards or recommendations for packet-switching networks. The best known of these is Recommendation X.25.

The International Standards Organization (ISO), which is composed of representatives from the manufacturing and user community, works closely with the CCITT; the ISO also has a group under its wing with responsibility for public data networks.

In the U. S., the American National Standards Institute (ANSI) is a clearinghouse that coordinates activity for voluntary standards. Its X3S37 committee, which has the responsibility for public data networks, does liaison work as well as coordination. This committee represents a cross section of U. S. industry: manufacturers, users, and carriers. It offers inputs to both the ISO and the CCITT.

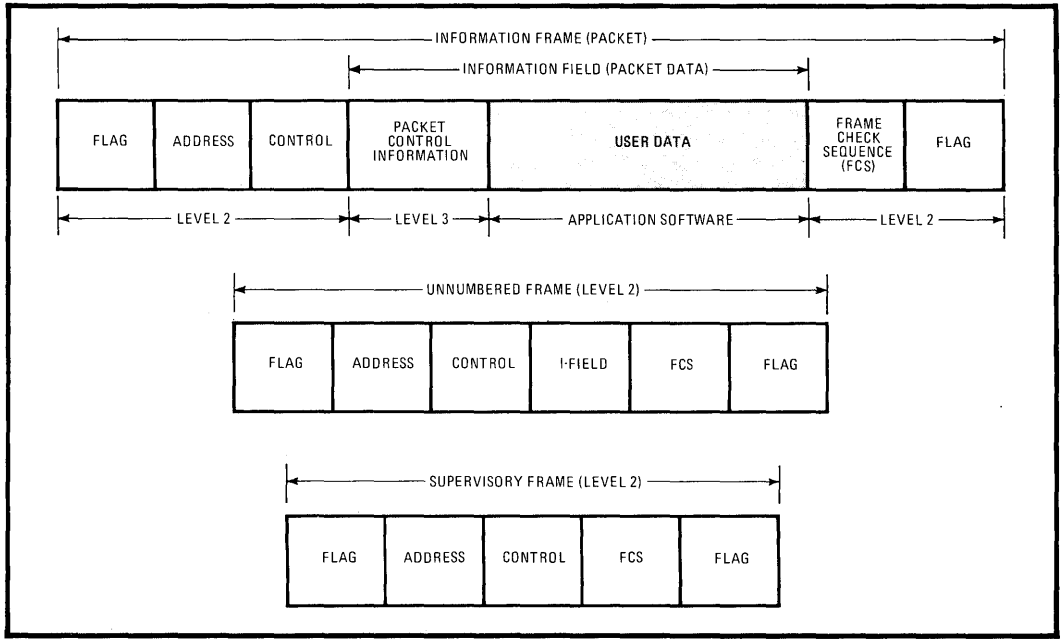

4. Standard frames. Three types of data frames may be sent over a packet network. All data except the user data in the information field of an information frame is system overhead required for synchronization, data checking, verification, and bookkeeping functions.

01111110. The receiving station automatically deletes the inserted 0s from the data.

The frame-check sequence is the last 16 bits before the closing flag. They are produced by a calculation that checks all data between the opening flag and the first bit of the FCS. The logic 0s inserted for data transparency are not checked.

A frame may be one of three types (Fig. 4): a supervisory frame (or S frame), an unnumbered frame (or U frame), or an information frame (or I frame).

Level 2 control does not involve itself with the data within the information field of an information frame. It simply encloses the packet data in an HDLC frame and sends it out onto the network.

Supervisory frames are used to perform supervisory control of a link, such as acknowledging packets, requesting retransmission of packets, and requesting temporary suspension of transmission. Unnumbered frames are used to set up, disconnect, and reset links.

The Level 2 protocol may take one of two forms: Link-Access Procedure (LAP) and Link-Access Procedure, Balanced (LAPB). When it was originally written in 1976, Recommendation X.25 contained LAP only. LAPB has been added since that time, offering some improve-

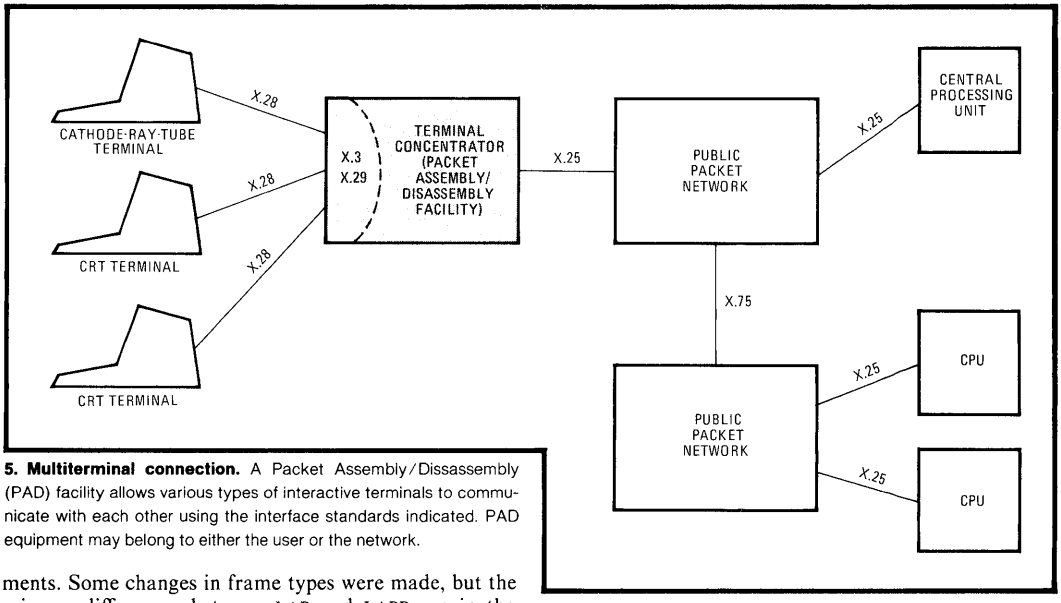

5. Multiterminal connection. A Packet Assembly/Dissassembly (PAD) facility allows various types of interactive terminals to communicate with each other using the interface standards indicated. PAD equipment may belong to either the user or the network.

ments. Some changes in frame types were made, but the primary differences between LAP and LAPB are in the functions that set up, disconnect, and reset links. (Two models of the μPAC, the WD 2501 and the WD 2511, are geared to the LAP and the LAPB, respectively.)

There are four system parameters defined by the X.25 Level 2 protocol: T1, N2, N1, and k. T1 is the time limit set for the primary timer; when T1 runs out, an unacknowledged command may be retransmitted. N2 is the limit set for a counter that is incremented each time a command is retransmitted because time T1 ran out without its being acknowledged. N1 is the maximum number of bits in a packet; it depends on the maximum length of the information field. And k is the maximum number of sequential packets that a terminal or node may have outstanding (transmitted but unacknowledged) at any given time. In the μPAC, T1, N2, and N1 are programmable. The number k can never exceed seven under X.25, and it is fixed at seven in the μPAC.

Multiplexing terminals

Since each user of the packet network typically has many different types of data generators and receivers, multiplexers must connect the network to the existing equipment. This multiplexer has been defined by the CCITT as the Packet Assembly/Disassembly (PAD) circuit (Fig. 5). The PAD is specifically for use with asynchronous terminals; it combines or separates the multiple signals that are sent to or received from the network.

CCITT protocol standards X.3, X.28, and X.29 are used together to define a PAD interface. A PAD facility may be viewed as a terminal concentrator that connects several asynchronous terminals to a single X.25 link. The PAD circuit is sometimes called an interactive-terminal interface because in practice most terminals connected to PAD interfaces require human interaction via keyboards and CRT displays or printing equipment.

When a PAD interface is used between the packet network and the terminals, two stations that are incompatible by themselves can communicate. They need only be able to talk to the PAD. The μPAC chips will allow them to do this. Another advantage of this approach is that new types of equipment added at a terminal are transparent to the network.

On the other hand, changes and improvements within the packet network are transparent to the user. These improvements could include increasing node-to-node communication speed, increasing the number of nodes, and changing node-to-node connections to fiber optics.

Variations on the theme

Many packet systems are available; they vary according to the network organization. Several networks, such as Montreal-based Bell Canada's Datapac, offer (in addition to the standard virtual circuit) a permanent virtual circuit that requires no call for link establishment and is continually available.

Another possible service, Datagram, when made available, will not require the initial establishment of a virtual circuit. In this approach a packet is merely put out on the line—typically by users of so-called transaction-based networks. There is no call procedure, and duration of connection is not of concern for billing purposes. Users may, for example, pay a flat fee. Short, independent data bursts will ultimately work their way through the network to their destinations.

A closed user group, available from Datapac and others, is like a private network. Users in a group, actually connected to a public network, can communicate with one another, but access is barred to and from all other users of the network. AT&T's proposed Advanced Communications Service includes this feature; the company calls it a virtual subnetwork. □

WD2511 LSI circuit simplifies packet-network connection

48-pin chip replaces entire board and thousands of lines of software

by Geary L. Leger, *Western Digital Corp., Newport Beach, Calif.*

☐ Packet-switching data-communications technology can now claim its first dedicated large-scale integrated circuit. Called the Micro Packet Network Interface chip, or μPAC for short, it is a complete X.25 Level 2 controller with on-chip bidirectional direct-memory-access facilities. This n-channel silicon-gate MOS chip in a 48-pin package replaces a board full of electronics.

The μPAC goes way beyond the functions performed by the bit-oriented-protocol (BOP) control chips currently in widespread use. It includes the circuitry of a BOP chip. But it handles many other operations, eliminating the need for separate DMA circuits and associated address latches, timing chips, and the system software (more than 1,000 lines of code) required until now to perform Level 2 control of the link between a data terminal and a node of a packet-switching network. It has an 11-K read-only memory and the equivalent of three microprocessors: one to handle data-transmission operations, another for dealing with received data, and a third central processor to coordinate all chip functions. Sample quantities of the controller will be available from Western Digital Corp. in the first quarter of 1980.

The data-link controllers already on the market (Western Digital 1933, Signetics 2652, Intel 8273, Zilog SIO, and others) handle BOP frame structure in a broad range of applications. For example, the WD 1933 can be used with the High-level Data Link Control (HDLC) and Synchronous Data-Link Control (SDLC) protocols, including the SDLC loop mode. This chip and others like it handle zero-bit insertion and deletion, the frame-check sequence (FCS), and the flags that define the beginning and end of a data frame.

The μPAC trades some of this protocol flexibility for the sake of greatly enhanced usefulness within its area of application (Fig. 1). It is restricted to the Level 2 packet-switching protocol defined in Recommendation X.25 from the Consultative Committee for International Telephony and Telegraphy (CCITT), a protocol developed from HDLC. But other BOP chips do not set up, disconnect, or reset the link; they do not automatically retransmit up to seven information frames (I frames); nor do they have a timer for retransmission control. These are

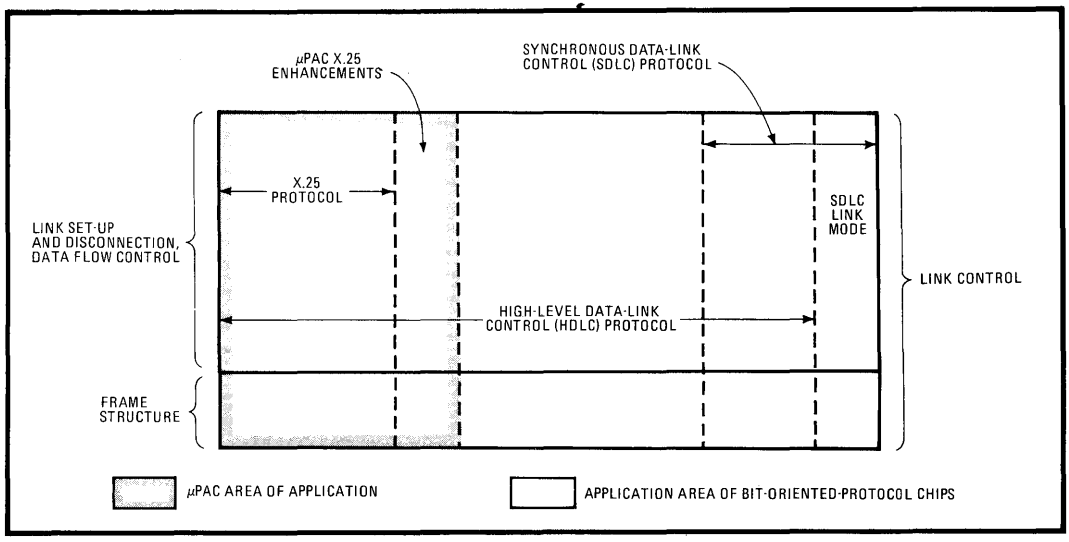

1. Targeted. The Micro Packet Interface (μPAC) chip is the first large-scale integrated circuit designed specifically for packet-switching applications. The application range of other chips that handle bit-oriented frame structure is wider, but the μPAC does much more in its area.

TABLE 1: COMPARISON OF FEATURES OF BIT-ORIENTED-PROTOCOL CHIPS

HDLC/ADCCP protocol feature	X.25 Level 2	Bit-oriented-protocol chips	μPAC
Basic bit-oriented frame structure	yes	yes	yes
Retransmission of up to 7 I frames (modulo 8)	yes	no	yes
Asynchronous response mode	yes, LAP	no	yes, 2501
Asynchronous balanced node	yes, LAPB	no	yes, 2511
Control of S, U frames	yes	no	yes
Link set-up, disconnect, and reset procedures	yes	no	yes
Time-out recovery	yes, T1/N2	no	yes, T1/N2
Multipoint operation	no	no	yes
Normal response mode (NRM)	no	no	no
Level 2 modulo 128	FS	no	no

FS = item for further study by the CCITT

all features of the μPAC chip (see Table 1).

Two versions of the μPAC will be made available. One, the WD 2501, uses the Link-Access Procedure (LAP) defined in the first version of X.25. The WD 2511 is for networks using the Link-Access Procedure, Balanced (LAPB) added to X.25 subsequently. The two chips differ only in the program stored in ROM. They are pin-compatible and interchangeable without hardware or software modifications. Both may be used either in a terminal (DTE, data-terminal equipment) or in a network node (DCE, data-circuit–terminating equipment).

Direct memory access

Because of the HDLC feature that allows up to seven packets (I frames) to be outstanding (transmitted but unacknowledged) at any time, the μPAC has information-field data (the I field of an information frame) buffered for up to eight packets both when transmitting and when receiving. In other words, the μPAC may have to retransmit up to seven packets. It must therefore be able to retrace its steps through as many as seven of its eight buffers.

DMA circuitry, included in the μPAC, is the best way to achieve this. A number of other control chips (floppy-disk controllers and data-link controllers) are DMA-compatible, but they do not actually include DMA. General-purpose microprocessors that have their own DMA, such as the Intel 8089, are not in the same category as the μPAC.

DMA control on the μPAC is simple, requiring only three pins ($\overline{\text{DRQW}}$, $\overline{\text{DROR}}$, and $\overline{\text{DACK}}$) for handshaking with the central processing unit's bus (Fig 2.). There are 16 address-output pins (A0 through A15) that are separate from the eight data pins (DAL0 through DAL7). This means that the DMA transfers are fast—they occur in a single cycle. Unlike the μPAC, DMA chips such as Western Digital's 1883 or Intel's 8257 require external address latches. This means that some or all of the address must come through the data bus and two or three cycles are required for data transfer.

In general, DMA control is either of the block-transfer type or the transparent type. In block-transfer DMA control, the DMA controller transfers several bytes of data while the CPU is disabled from using the bus. If transparent, the DMA control is imbedded in the CPU's clock cycle in such a way that the transfers are invisible, or transparent, to the CPU. Since the μPAC must be able to transmit and receive data on two DMA channels at once (for full-duplex operation), the only logical choice for the μPAC is transparent DMA, since block-transfer DMA would restrict operation to half-duplex.

All Level 2 data is appended and checked automatically by the μPAC. The I-field data is accessed via DMA channel. All supervisory frames (S frames) and unnumbered frames (U frames) are automatically transmitted and checked by the μPAC. The user's CPU operates only on the I field of I frames.

Keeping track of packets

The DMA uses two lookup tables—one for transmitted frames (TLOOK) and another for received frames (RLOOK). These contain addresses and control bytes for the individual packets. Thus packet data is addressed indirectly. This method is best suited for most software applications.

The 16-bit starting address for TLOOK is loaded into the μPAC by the CPU. RLOOK must follow immediately, and both TLOOK and RLOOK are stored in random-access memory external to the μPAC.

There are a total of eight segmented control sections for each table. Each section contains 8 bytes, 4 of which are used for memory starting address and length. The rest are for control.

In the transmit mode, the μPAC reads (from TLOOK) the starting address and length of the first packet to be transmitted. The chip then automatically transmits the flag, address, and control fields. Next, the information-field data is transmitted using DMA and the memory location called "send #0 packet." At the end of the information field, the μPAC automatically sends the FCS and closing flag. It then moves on to the next packet.

If retransmission of one or more (up to seven) packets becomes necessary, the chip automatically retraces the previous transmissions through the TLOOK table. The user's CPU software does not become involved in the retransmission. An error counter is incremented.

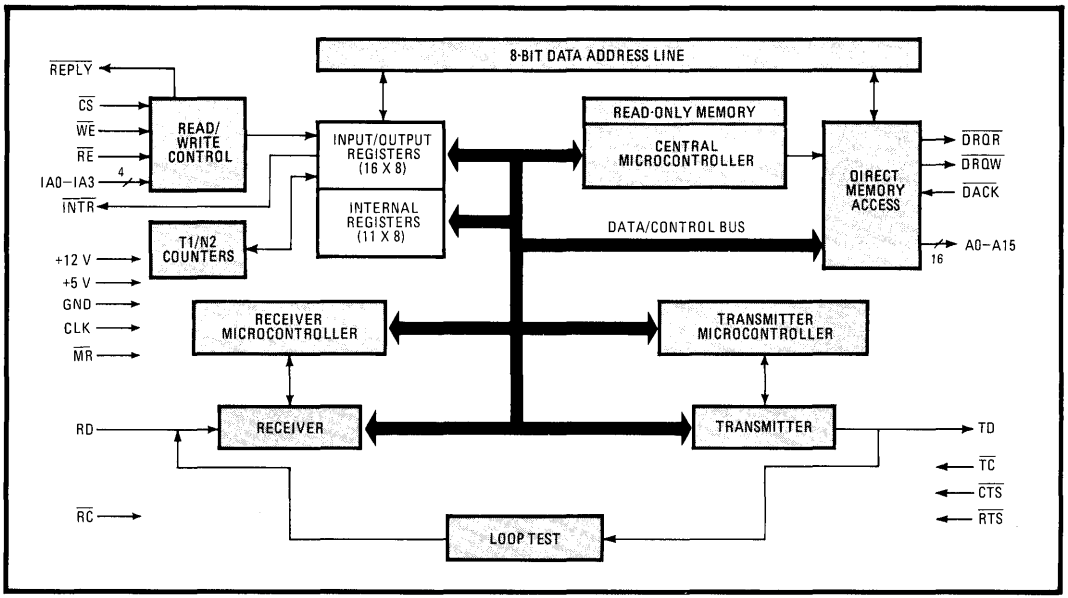

2. Inside the μPAC. The Level 2 controller has its own timer and direct-memory-access circuitry and is the logical equivalent of three microprocessors. Routines stored on the chip allow it to relieve the network user's central processor of a large software overhead burden.

Each received frame is checked for correct address and FCS fields and for type of control field. If the frame is an I frame, the I field is placed in the assigned memory location using a method similar to that used in transmission. After the packet is received error-free and in proper sequence, an interrupt is generated and the μPAC is ready for the next packet, which will be placed in the next location.

Ten 8-bit error counters follow RLOOK in the external RAM. These counters do not cause an error interrupt, but maintain a running count of error activity. The contents of the counters include: the number of frames received with FCS error; the number of times T1 (the time minimum set for a timer that allows retransmission of an unacknowledged packet) ran out; and the number of packet retransmissions.

Control bits are included in TLOOK, RLOOK, and the μPAC to ensure orderly transfer of data blocks. For example, the control bits are designed to prevent what is known as "deadly embrace," a situation in which the μPAC and the user's computer are waiting for one another to start.

Self-testing

Self-testing features are critical to proper operation. The μPAC does a comparison test, an internal RAM register test, and a loop-back test. All three are suitable for use during manufacturing and inspection. The internal RAM and loop-back tests are also useful for system diagnostics and troubleshooting.

The comparison test requires a device known to be good or a stored list of known good responses. The program location counter (PLC) for the main ROM is halted so it may be incremented under external control.

TABLE 2: MICRO PACKET INTERFACE CHIP (μPAC) TERMINOLOGY

ABM	asynchronous balanced mode
ARM	asynchronous response mode
CMDR	command reject (U frame, LAP only)
DISC	disconnect (U frame)
DM	disconnect mode (U frame, LAPB only)
FRMR	frame reject (U frame, LAPB only)
REJ	reject (S frame)
RNR	receiver not ready (S frame)
RR	receiver ready (S frame)
SABM	set asynchronous balanced mode (LAPB only)
SARM	set asynchronous response mode (LAP only)
UA	unnumbered acknowledgement (U frame)

All jumps stored in ROM are disabled so that each location of the PLC may be counted. As the PLC is incremented, the responses of the output pins and status registers are compared to the known good responses.

There are 11 8-bit registers in the μPAC that are not directly accessible by the user's CPU, which complicates testing somewhat. The internal RAM register test provides a means of checking these registers. The contents of register A are placed in six even internal registers and the contents of register B in five odd internal registers. The 11 registers are then added together without carry and the result is placed in status registers. This test is initiated by a control bit in the μPAC. The loop-back test is discussed later.

For the purposes of discussing link establishment procedures, it will be assumed that there is a 2501 μPAC at each end of the link. In practice, the 2501 can

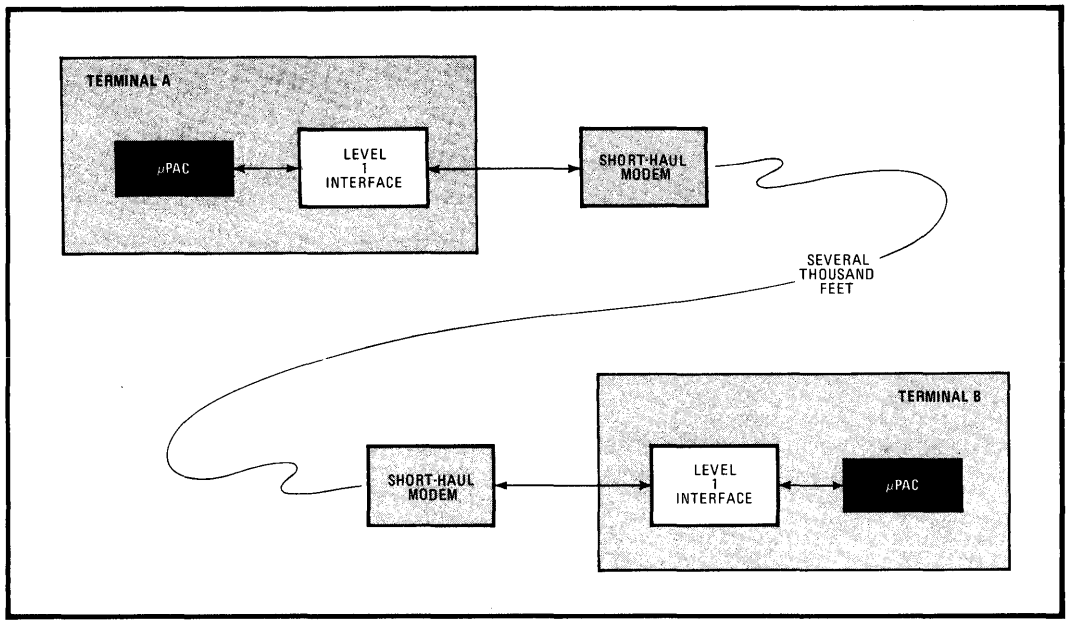

3. Off the network. The μPAC is also useful in non-network applications that use bit-oriented protocols. It provides full-duplex capability, does error detection and recovery, and gives systems the option of hooking directly to a packet-switching network at some future date.

communicate with any device meeting X.25 Level 2 specifications.

When a link is set up, it is said to be in the information-transfer phase. This means that the terminal and node will accept and transmit I and S frames. When a link is logically disconnected, only U frames—DISC, SARM, or UA (disconnect, set asynchronous response mode, and unnumbered acknowledge; see Table 2)—will be accepted or transmitted.

Link supervision

A link-connect frame is not the same as a link-reset frame. A link in the information-transfer phase may be reset in one direction by a SARM transmission. A link is up after both ends send a SARM command and receive a UA response.

Since a SARM can be either a command to reset or set up a link, misinterpretation by the receiver of a SARM is possible. This could happen when a link is established if one end momentarily loses power. When that end tries to bring the link up by sending a SARM, the other end may interpret the command as a link-reset.

There are two ways to get around this problem. Suppose a terminal or node attempting to bring a link up sends a SARM command and receives a UA. After time T1, if the station does not receive SARM, it assumes that the other end considered the link up. It will then disconnect the link by sending DISC and receiving a UA, and attempt to set up a link a second time.

The other way around the problem is the method used by the 2501. The 2501 will always send DISC and receive a UA before attempting to bring the link up. This will assure a logically disconnected link so that it may attempt to set one up. Immediately after the link is up, the 2501 generates an interrupt.

It is possible to recover a single error on a packet with μPAC control. The error makes the received FCS bad, so B does not recognize A's first transmission of frame 1. When B receives frame 2, something is wrong since the last successfully received packet was frame 0. Thus, at the next opportunity, B sends a REJ (reject—an S frame) asking A to retransmit frame 1. This opportunity comes after B completes sending its frame 2.

When A receives the REJ frame, it is sending frame 3. There is no need to continue with frame 3, so A aborts transmission of frame 3 and goes back and retransmits frame 1. After retransmitting frame 1, A will retransmit frames 2 and 3. Finally, A will continue transmitting other frames.

Loop-back

A loop-back condition exists when a station receives the same serial information it has transmitted. In the loop-back test, the serial-transmit output is connected to the serial-receive input in order to test the transmitter and receiver channels. Each station has both primary and secondary functions, so there are two logical primary-to-secondary associations on a terminal-to-node link, and each association is identified by a different address field. This makes loop-back testing impossible when a strict X.25 connection is made. Commands will have the A field of a response and vice versa. One way around this is to make the A fields of the two associations equal for the duration of the loop-back test. (The A fields are programmable in the μPAC.)

Another problem with loop-back testing is the actual

detection of the condition and the detection of the condition's removal. There is no simple way around this problem, and the μPAC gives only limited assistance.

First, detecting the existence of a loop-back condition is the responsibility of the CPU driving the μPAC. If the CPU sees that a link cannot be brought up, or if a link is up and suddenly has excessive link resets and CMDRs (command reject, a U frame), the CPU could assume the presence of a loop-back condition. After making the two A fields the same, if a disconnected link is successfully brought up, then the loop-back condition exists.

To detect the removal of this condition, a particular control bit (RRT1) in the μPAC may be used. It causes the μPAC to send an RR (receiver ready, an S frame) or an RNR (receiver not ready, also an S frame). These frames are sent at T1 intervals as long as the μPAC is not commanded to send a packet. As long as the μPAC receives those S frames, the loop-back condition exists. However, if the μPAC fails to receive an S frame for a time equal to T1 × N2, an interrupt is generated, signalling that the loop-back condition has been removed.

Modified X.25

The original design intention was to use the μPAC in a strict X.25 terminal-to-node application, the only application covered by X.25. However, by taking advantage of the terminal-node symmetry of the μPAC (the fact that it can be used in both DTE and DCE), other applications are possible that use its built-in features.

For instance, the user does not need to develop the software for error recovery since this is a μPAC feature. For another, using a μPAC makes it possible to connect a non-packet terminal to an X.25 link at a future time. And lastly, the chip's protocol is bit-oriented. It has a number of advantages over older, character-oriented protocols, such as code transparency, full-duplex capability, flexibility, and modularity [*Electronics*, Jan. 18, 1979, p. 137].

One possible application is the connection of two terminals at Levels 1 and 2 (Fig. 3). How much of Level 3 is used would depend upon the individual application; the more of Level 3 used, the better standardized the interface is. One of the terminals in Fig. 3 could be a terminal concentrator (a Packet Assembly/Disassembly facility, or PAD, as defined by Recommendations X.3, X.28, and X.29) on a factory floor, and the other could be a host computer in a data-processing center.

Modified X.25 could also be used in a multipoint system (Fig. 4). Idle terminals in this type of system must transmit an "idle" sequence, not continuous flags. The terminal addresses (AD3, AD5, AD7, and so on) correspond to the transmitted response A field. The transmitted command A field is the same for all terminals and is chosen to be hexadecimal 01 in this case. All A fields are selected with odd values (least-significant bit transmitted first) to conform to the extended-address format of the Advanced Data-Communication Control Procedures (ADCCP).

Two terminals on the multipoint line may establish and discontinue communications by exercising X.25 procedures for setting up and disconnecting a link. But only two terminals can communicate at any one time.

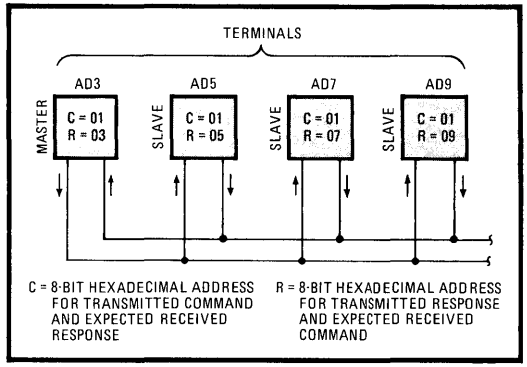

4. Multipoint line. The programmed features of the μPAC chip enhance the flexibility of a system comprising one master terminal and up to 128 slave terminals. Hardware and software savings are possible when the μPAC is used in this off-network context.

Suppose that AD3 wishes to communicate with AD7. AD3 will first make sure that its receiving line is idle (a status bit in the μPAC). Next, AD3 will change its transmitted command and response A fields to be the reverse of AD7 (command field is set to 07, response field is set to 01). Then AD3 will initiate link establishment by setting a control bit, called "active," in the μPAC. Once the link has been established (the μPAC generates an interrupt when the link is first set up), AD3 and AD7 may exchange I frames. To discontinue the session, either AD3 or AD7 will set the mandatory-disconnect control bit in its μPAC. This will cause that terminal to initiate a logical-disconnect procedure.

Contention and roll-call methods

The multipoint system may be implemented by either contention or roll-call polling. In the roll-call method, the master terminal will initiate link establishment with one of the slave terminals, communicate with that slave, discontinue the session (disconnecting the link), and go on to the next slave. This process continues until all slaves are polled and then starts over. One advantage of the roll-call method is that the master has tight control over the line for efficient operation.

A disadvantage is that slaves must be queried (polled) before sending data, and the more slaves on the line, the longer it takes for the master to poll them. Therefore it is essential that each slave be designed to exchange a relatively small amount of data with the master in a single session, lest it tie up the line for long periods. Large amounts of data should be broken up and exchanged in more than one session. This method is suited to applications where the multipoint line has a high usage.

In the contention method, any terminal may initiate a session at any time. This is similar to a party telephone line and is suited to applications where line usage is low. All sessions are between the master and one of the slaves, but unlike the roll-call method, a slave may initiate the session. The terminal that initiates a session must send an I frame with its unique address immediately after the link is set up. □

LOCAL NETWORK ACCESS TRADEOFFS

Cost/complexity tradeoffs are examined in CSMA/CD and token passing techniques for accessing local area networks

by Mark Stieglitz

Local networks are characterized by problems that are very similar to those encountered in conventional data communications networks.[1] Local networks, however, generate new problems and opportunities that require reconsideration of tradeoffs in system cost/complexity. A fundamental point of decision in local network design is the choice of access method. Chief contenders among access techniques are carrier sense multiple access and token passing.

What is a local network?

The current controversial nature of local area networks (LANs) is highlighted by their many definitions. A common theme in these is that the LAN be privately owned and/or administered by the user. An LAN need not be considered only as a high speed data transfer mechanism; current private branch exchanges also meet the definition of a private system. The opportunity to optimize the network for a particular user's application, therefore, becomes a key feature of the network. In this discussion we assume the following: that a local network is a privately owned communication system; it usually runs at data rates of 100k bits/s and above; and it is usually restricted geographically (100 to 25,000 m).

It is often asked if the X.25 protocol can be used in LAN applications, especially now that X.25 large scale integration (LSI) controllers are available. This question can be more readily answered by comparing LAN and X.25 protocol functions using the International Standards Organization Open Systems Interconnection (ISO/OSI) reference model.[2,3] The model was developed to help conceptualize the relationships of various elements in a communications protocol. The access function resides between physical and link level functions, often referred to as a link layer sub-layer (Fig 1). The primary difference is that the concept of a shared medium is foreign to X.25. Addresses at the link level are actually command/response indicators, since it is assumed that pairs of stations have point to point links between them.

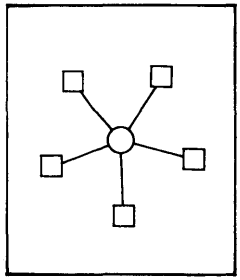

The local network access layer implements both the device arbitration and addressing necessary for shared medium operation. Once this layer is chosen and implemented, it is expected that the remaining layers may be used in this new application with little change.

Network topologies

In simple terms, topology is the way in which networks are tied together (Fig 2). Many networks are wired in ring or star configurations in order to eliminate the contention problems that occur when more than one connected device tries to send data at the same time. The primary advantage of the bus topology is easy reconfigurability, more important, perhaps, than its reliability advantage. Costs of improving reliability in a star or ring network, eg, adding redundant subsystems, can be much less than reconfiguration costs of the same network over its lifetime. Reconfiguration is labor intensive, and the cost of labor is increasing at a faster rate than that of reliable electronics.

The security of a broadcast bus system is often questioned by users who are apprehensive of the party line

Mark Stieglitz is local networks program manager at Western Digital Corp, 2445 McCabe Way, Irvine, CA 92714. He is responsible for planning and developing the company's LSI and system level local network products. Active in IEEE standards activities, he is currently chairman of the IEEE-802 Token Access Working Group.

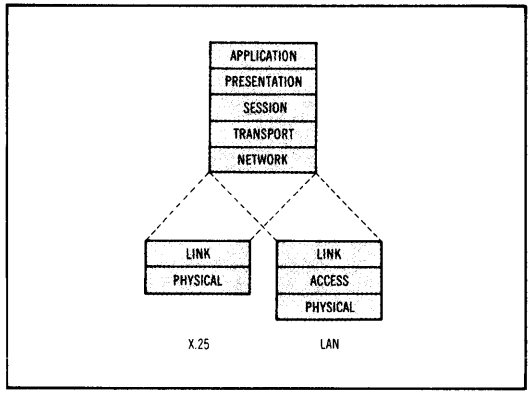

Fig 1 ISO/OSI reference model applicability. The model directly applies to local networks with addition of access layer

concept, where they share a network with diverse user groups. This problem is readily overcome by encrypting the appropriate data on the network. This alternative was at one time unfeasible because of high costs. Now several solutions are made possible by extensive LSI implementation of the National Bureau of Standards data encryption standard, which resolves this obvious problem in bus topology.

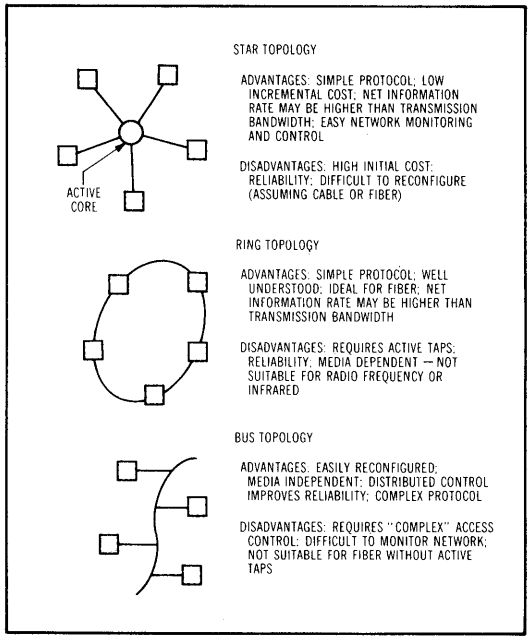

Fig 2 Typical local network topologies. Each has fundamental strengths and weaknesses. Bus topology's efficiency, maintainability, and cost are heavily dependent on access method used

Access methods

Currently the most controversial open question in the local network area is the choice of access methods in LAN buses. An access method is that part of a protocol that coordinates bandwidth use among all network subscribers. It ensures that only one station transmits at a given time, or, if more than one, that proper recovery action is taken to provide correct data transmission. Two common methods for allowing multiple transmission sources on a broadcast medium are frequency division and time division multiplexing (FDM and TDM). Both are fixed assignment schemes and require some centralized network intelligence to assign channels (FDM) or time slots (TDM). There are cost and reliability drawbacks to this centralized scheme. Also, it is difficult to effectively use the communications bandwidth where there are many sporadic data sources, such as word processing terminals. The solution to this lies in a demand access scheme, two of which are currently being promoted.

Carrier Sense Multiple Access (CSMA). In this method a station wishing to transmit listens first for channel clear, and transmits if such is the case. When two stations hear that the channel is clear and transmit simultaneously, a collision occurs. This must be detected and recovered by the CSMA protocol. The simplest type of collision detection requires a higher, usually link level, intelligence to note that a frame has been lost on the network. All frames would be buffered until acknowledged and retransmitted if no timely ACK is received.

Carrier sense multiple access with collision detection (CSMA/CD) is a CSMA implementation that can detect transmission collisions while the data are being transmitted. This enhancement greatly minimizes bandwidth wastage during collisions, but imposes a minimum size restriction on every frame to ensure that collisions are detected (Fig 3). A more serious drawback in collision detection is in its actual implementation. It must detect two simultaneous transmissions (a station's receiver must "listen" for others while its own transmitter is "talking"). Transceiver design is critical. Special cable and cable taps are often needed to minimize noise and impedance problems. Special installation and grounding practices that have been developed may necessitate additional training of cable installers and modifications to building codes. All these constraints have recurring cost implications. So, while many solutions have been implemented, some are costly, and each is media/speed dependent. Several different systems using CSMA/CD are commercially available. The most notable is Ethernet, a joint offering of DEC, Intel, and Xerox.[3,4]

The CSMA scheme is comparatively simple and has enjoyed much academic research, but it has some shortcomings. In the pursuit of simplicity, visibility of network errors and the potential for future upgrade have been sacrificed. Since CSMA allows and expects collisions on the transmission medium, it is difficult for diagnostic equipment to distinguish expected errors from those that are induced by noise or faults. Determinism, or the ability to guarantee the successful (no

collision) access of a station within a fixed time interval, cannot be accommodated in a CSMA environment.

Office automation, which is not real time and therefore CSMA compatible, is a major local network market. Process control, the other major application category, requires absolute delay limits and reliability guarantees. Both markets can be addressed with the same "standard" protocol and access method only if the needs of both are met. The token access method is a way to accomplish this.

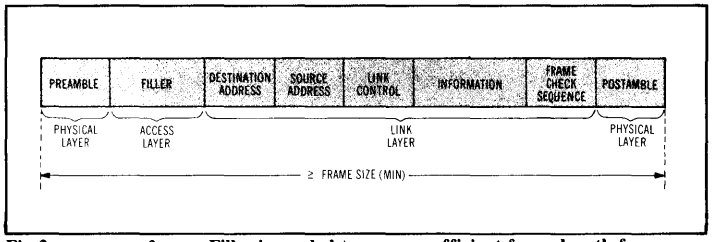

Fig 3 CSMA/CD frame. Filler is needed to ensure sufficient frame length for collision detection. Minimum frame length is function of propagation delay through maximum length of medium

Token passing. A token is an exclusive right, held by exactly one station at any given instant, to initiate transactions on the medium. Distributed network intelligence passes this access right around the network in a logical ring, resulting in an ordered and controlled access method. In the token passing scheme, sometimes referred to as "baton passing," each station sends a message to its access successor when it has finished its transactions (Fig 4).

Control messages are sent in the same format as is information, in frames. At first glance the token access scheme's frames look much like those of CSMA systems as shown in Fig 3. The similarities are purposely at the physical and link layers (Fig 5). The similarity ends with the access field; the required filler in CSMA/CD systems is replaced in the token passing frame with a token control field, usually of one octet.

The required control information could have been coded into the link level control field, but instead is placed directly ahead of the link field. There are three reasons for this. First, it provides adherence to the ISO/OSI model's sense of encapsulation. This says that a given layer must not modify or require the use of any data in a higher layer for its own proper operation. Observing this requirement saves software development and redevelopment as users switch between X.25 and LANs. Second, the ability to send "piggyback" tokens requires separate access and link control fields; link information can go to one station while control is (optionally) passed to another. This is an efficiency enhancement that allows a reduction in the bandwidth used for network management. Third, special format access frames can be sent. Since the access control field may be thought of as defining the rest of the frame (for example, an opcode), very short access frames can be transferred and evaluated without modifying link control programs.

While both access methods are conceptually simple, there are several implementation challenges in the token scheme. These include network initialization, building and maintaining the logical ring (online addition/removal of stations), and the resolution of fault recovery conditions. Centralized and fully distributed are two categories of solutions for these tasks.

The centralized scheme uses an administrative station to watch for and resolve unusual network conditions. Removing this chore from the bulk of the stations simplifies their processing requirements and thus their cost. This administrative method is also generally more expeditious than distributed schemes, since the latter require delays in their distributed algorithms and require all stations to rediscover their part of the network configuration each time.

In the distributed scheme, reliability and ease of configuration are achieved; the network configures itself each time it is initialized. The best of both worlds, speed and reliability, are achieved in the hybrid system. Here the distributed algorithms are retained as backup in case of an administrative failure.

The token protocol makes no assumptions about, or "improper" use of, the transmission medium or transmission rate. Any collisions are treated simply as manifestations of noise and are consistently handled as exceptions. No expected collisions mean no confusion as to cause, resulting in improved maintainability and serviceability. The use of strictly "inband" signaling allows true media independence. Radio frequency, infrared, CATV, baseband coaxial, fiber, and other broadcast media are usable with no change in the access algorithm or any sacrifice of efficiency. This flexibility will be useful as data rates and distances grow and as new transmission technologies are developed.

Inband signaling also means that existing components and technology can be used. This gives network

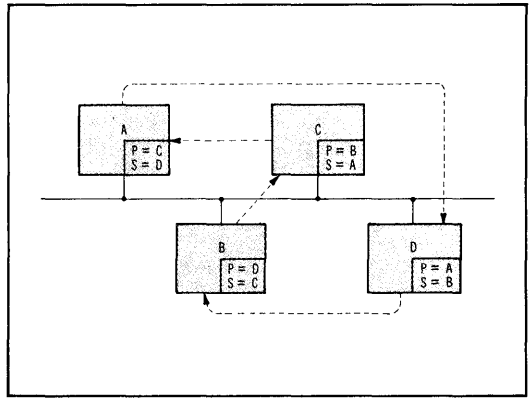

Fig 4 Model of logical ring. Each station has sufficient intelligence to receive and validate tokens from its predecessor (P) and send tokens to its successor (S). Physical ordering of stations is not relevant. Dashed lines indicate control flow

implementors the option to capitalize on established production efficiencies and low costs such as are represented by CATV components. From the use of existing broadcast technologies follows the applicability of existing regulations and trained cable installers.

Token protocol's insensitivity to transmission speed is another important factor. It is unreasonable to assume that all network users need the same arbitrary data rate, such as 10M bits/s. Users with lesser requirements should be able to scale systems to their needs and budgets. Some CSMA implementations have minimum frame size restrictions that are directly based on the data rate and the physical length of the medium, to say nothing of the cost of multi-megabit hardware. There is no reason that a few cathode ray tube terminals cannot be linked together with inexpensive twisted pair cable, using the same token protocol and controllers as those used in applications with higher speed requirements.
tions with higher speed requirements.

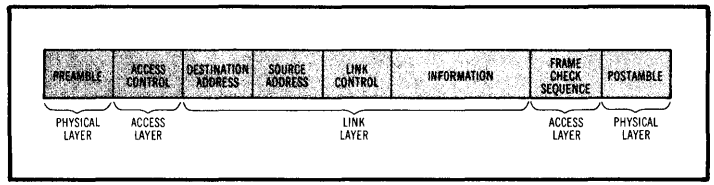

Fig 5 Token frame. Separate access control field allows option of passing control to one station while sending link level information to another

Depending on the application, networks must either be fair (where all stations have equal access to the medium), or include some priority mechanism. The token access method supports both conditions by being generically fair, but also allows tuning of network and station parameters if desired. Features such as sending "n" frames while holding the token are easily supported. This allows prioritization of stations where some may be allowed to transmit more than others before giving up the token. The network may be set to guarantee access to all nodes within strict time boundaries, as required in control applications.

If tokens solve all LAN problems, why is there any controversy? The answer to this lies in the real and perceived complexities of the token access scheme.

Is token complexity worth it?

Complexity considerations must be evaluated on two fronts: technical (Can it be implemented reliably?) and economic (Is any additional incurred cost justified?). Intensive efforts by individual companies and standards groups have yielded some commercial offerings and several technical proposals. The token access method has been reviewed and evaluated by academicians, network implementors, and users. With the systems, models, and documents available today, it can safely be said that the token scheme is implementable.

The LSI developer is challenged to deliver this complex protocol at low cost. With such an LSI controller, a day can be envisioned when users need be as little concerned about low level network protocols as they are today with bit locations and formats on floppy discs. Efforts in protocol design are nonrecurring, but the benefit of a sophisticated, forward-looking design course will manifest itself more and more as network requirements grow.

Algorithm details and standardization

In the general token scheme just described, detailed algorithms vary depending on the system and design requirements. Choice of an initialization algorithm, for example, depends heavily on the address range allowed in the network: a 48-bit address range uses a different station sort scheme than does an 8-bit range. Work on the distillation of these tradeoffs is underway by standards committees and commercial organizations.

Standardization is key to volume manufacture of token controllers and to the interconnectivity of multivendor equipment. Standardization also advances the development of network diagnostic equipment and tools.

Summary and conclusions

A new science requires fresh consideration of engineering challenges. The needs of users and the progress of implementation technology, especially LSI, can be projected. There is no reason to accept any scheme simply because it exists, as proposed in References 5 and 6. Professional skill and judgment must be used in selecting all elements of any system, especially one as new and with such potential impact as the local area network.

The general token access scheme enjoys current commercial use, generality, and expandability that make it a truly useful standard. Investment costs in up-front complexity will be continually reduced with further LSI developments and as network uses proliferate.

References

1. J. M. McQuillan, "Local Network Architectures," *Computer Design,* May 1979, pp 18-26
2. H. Zimmerman, "OSI Reference Model—The ISO Model of Architecture for Open Systems Interconnection," *IEEE Transactions on Communications,* Apr 1980, pp 425-431
3. J. M. Kryskow and C. K. Miller, "Local Area Networks Overview—Part 2: Standards Activities," *Computer Design,* Mar 1981, pp 12-20
4. The Ethernet, A Local Area Network—Data Link Layer and Physical Layer Specifications, Version 1.0, Sept 30, 1980, Digital Equipment Corp, Intel Corp, Xerox Corp
5. P. Franson, "It's time to get on the Ethernet bus," *Electronic Business,* Oct 1980, p 6
6. L. J. Curran, "Seconding an Ethernet motion," *Mini-Micro Systems,* Nov 1980, p 73

Please rate the value of this article to you by circling the appropriate number in the "Editorial Score Box" on the Inquiry Card.

High 713 *Average 714* *Low 715*

LOCAL NETWORKS

Token passing cashes in with controller chip

Token-passing protocols can upgrade a data-communication system, especially if a dedicated controller relieves the host from token-processing tasks.

Mark Stieglitz, Network Products Manager
Western Digital Corp., Communications Division
2445 McCabe Way, Irvine, Calif. 92714

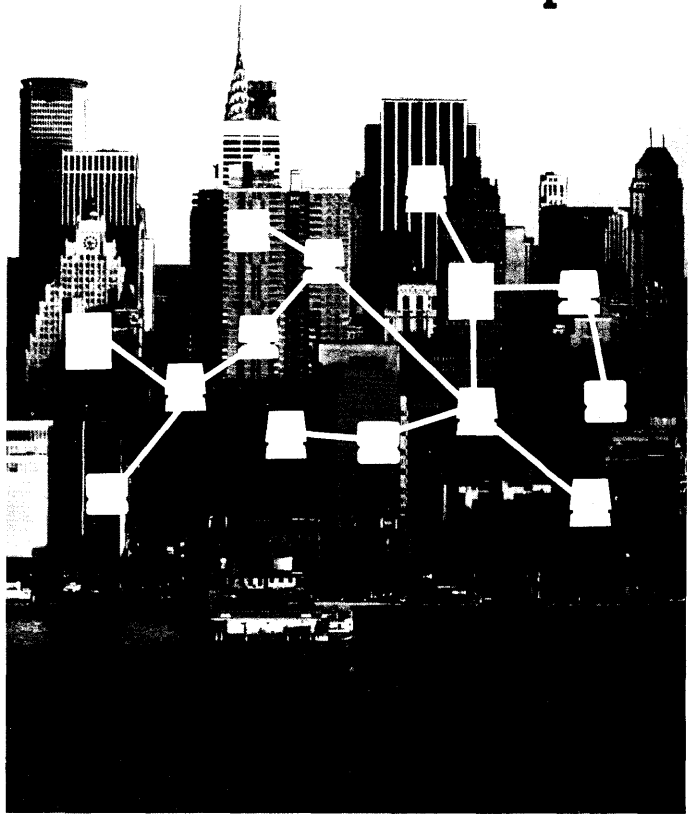

Designers can now implement a distributed-access token-passing systems without worrying about the complex details involved in the communications protocol. Those are taken care of by one LSI chip, called the token-access controller.

Token passing is one method of sharing a communications path. It enjoys the benefits of distributed-access systems while eliminating the drawbacks of schemes employing carrier-sense multiple access with collision detection (CSMA/CD). Until recently, however, token-passing techniques had little currency because of their need for complex controllers. This need relegated tokens primarily to proprietary uses. (For a complete review of local networks, including token-passing techniques, see "Broad Standards, Many Implementations Are on the Way," ELECTRONIC DESIGN, Sept. 30, p. 87.)

The introduction of Ethernet in 1980 marked the beginning of commercial local networks using distributed-access techniques. A distributed system does not rely on a single device for polling. Early versions of Ethernet were designed to use simple controllers because most of the development work began in the 1970s, before the LSI era. Although CSMA/CD offered a simplified access protocol, users incurred cost, performance, and flexibility penalties. Progress in LSI technology has now given designers the benefits of complexity—increased efficiency and enhanced flexibility—but without high cost.

The availability of the WD2840 token-passing controller chip brings token-access communications capability to a range of critical-process applications that were previously unsuited to the token method.

Reprinted from ELECTRONIC DESIGN - Oct. 14, 1982

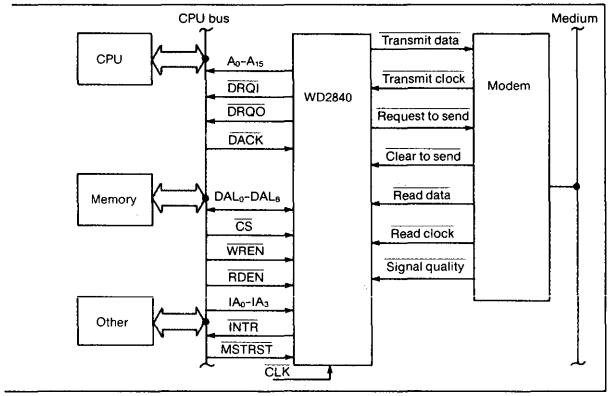

1. Two interfaces connect the WD2840 token-access controller chip to a local-area network. The network side (medium) is electrically conventional, whereas the host side (CPU) combines both a control and status register and a direct-memory-access interface.

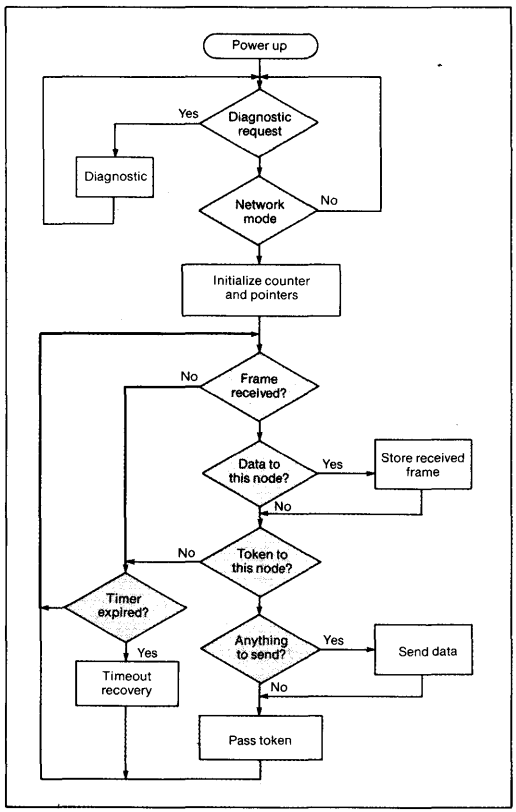

2. After initialization, the token-access controller idles in its "watching" loop (bold lines), waiting for a data frame or token addressed to it. The upper portion of the flow chart shows the steps in the initialization procedure.

What's more, the controller sufficiently reduces system communication costs, encouraging its use with very inexpensive end products.

The controller is designed to connect distributed intelligent devices over a shared broadcast medium —usually coaxial cable, free-air radio, or twisted-pair party line. Shared by all stations through the use of a token-passing protocol, the broadcast medium enables each attached device to hear everything on the network. A station is a microprocessor-based device that incorporates the controller. Complementing the token protocol's efficiency is a high-level, software-friendly DMA (direct memory access) interface contained within the controller. In conjunction with conventional hardware and serial communications interfaces, the DMA interface and token protocol make the WD2840 simple to use.

The primary purpose of a token-access controller is to free the designer from data communications concerns. Once the chip is initialized, for example, the host microprocessor need never bother with the protocol; it merely processes frames addressed to it —the controller filters out all others—and generates messages to send later. In fact, it is the controller that sends messages when a token is received. This decoupling of the functions between the network and any user processing simplifies programming and system timing considerations.

Tasks that affect network performance—such as processing tokens and generating acknowledgments —are performed inside the controller. Thus designers can use any type of host processor in a station (Fig. 1). The circuit interfaces with systems in which a processor is busy with a specific application.

Illustrating that point is the microprocessor found in a CRT terminal. Its duties are to scan the keyboard and perform a limited amount of editing. That leaves enough processing power remaining to drive the controller, which handles data flow only for one specific terminal. If the processor falls behind momentarily, just one terminal is affected; all others in the network continue to operate at full speed. Thus a network of controller chips is not slowed by its weakest link.

Token passing in a system environment

When a token-access controller receives a data frame addressed either to itself or to all broadcast stations—there are 254 stations in a system—it transfers the frame to the host's memory via a DMA operation. If the frame is invalid, the close coupling of protocol handling and DMA operations allows the chip to manage its own housekeeping. For example, if errors are detected through a CRC (cyclic redundancy code) or signal-quality check, the memory space is automatically reclaimed.

On the successful reception of a data frame, the controller sets an interrupt and checks to determine whether an acknowledgment was requested by the sending station. If so, it sends the acknowledgment, adding the receiver status, as well. A typical message might be "Received successfully" or "Encountered DMA problem on this end. Please retry." While the controller evaluates each frame as it looks for its own data message, it simultaneously checks for tokens passed to it. This combined token-and-data frame—in which each part can be directed to a different station—is referred to as piggybacking, a feature that increases system efficiency, since most of the overhead associated with conventional token-passing is eliminated (Fig. 2).

Receipt of a token allows the WD2840 to transmit

A three-controller architecture

A single NMOS LSI chip, the WD2840 token-access controller, comprises three major elements: a fast serial communications subsystem, a two-channel DMA controller, and a microprocessor with internal ROM and RAM.

The device's three preprogrammed microcontrollers handle media access and host memory-management functions. This type of architecture facilitates internal parallel processing: for example, prefetching a new buffer address while transmitting or receiving data. Although the token-passing protocol is essentially a half-duplex scheme, separate receiving and transmitting subsystems permit loopback testing.

The primary microcontroller has the capabilities and instruction set of a conventional 8-bit microprocessor, including subroutines, bit manipulation, conditional branching, and arithmetic operations. This part, whose chief task is to implement the token algorithms and maintain the host memory chain, has its firmware located in the internal 1-kbyte ROM. Repetitive and simple operations such as DMA fetching and storing are controlled by the receiving and transmitting microcontrollers.

Two internal timers keep the network independent of the host microprocessor's timing. The first timer is set for a relatively short duration to limit the time it must wait for the required response from a transmission. The second timer has a longer and less critical duration that restricts the period that a network can normally be "idle." Limits on idle time are useful for initialization and some error-recovery operations.

The serial-to-parallel and parallel-to-serial converter block includes standard 16-bit cyclic redundancy code (CRC) checking and generation, along with the framing logic specified by the HDLC (High-level Data Link Control) protocol. Also, the receiver contains an input FIFO buffer to speed internal processing and relieve DMA latency constraints.

Hardware interfacing is designed for flexibility. The DMA interface, for example, uses the familiar DMA request signals as outputs—one for input requested, one for output. Together with a grant signal that both notifies the token-access controller that the bus is available and optionally enables the address drivers, it permits synchronization with slower external memories.

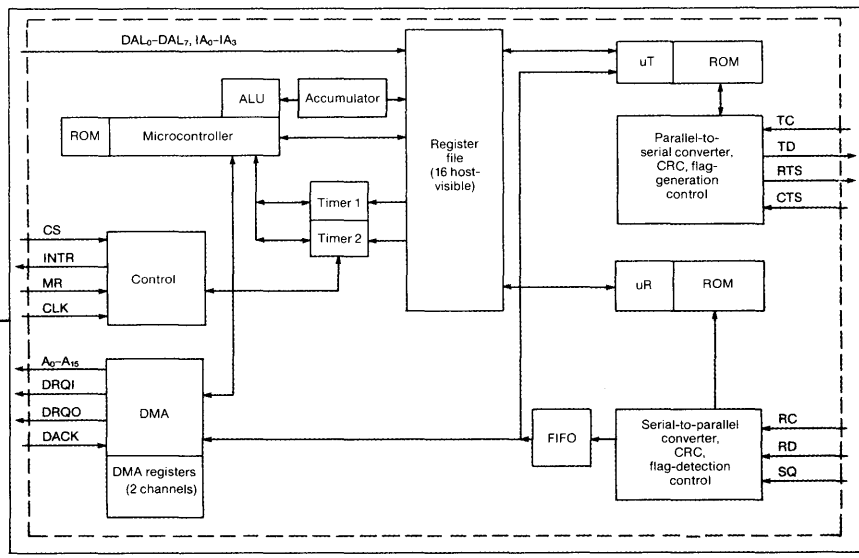

messages that its host has queued. Before transmitting, however, the device checks its internal hold-off register, AHOLT, to determine whether it should defer use of the token on the current cycle. With this optional deference capability, the system designer can bias the intrinsically fair token protocol in favor of selected stations. Those that have more access opportunities have effectively higher priorities.

Key to a token system is the visibility that each station has to network loading conditions. The less often a token is received in a given time, the greater the load on the network. This indication of load data is available to the host as an optional interrupt/token received, and the host can scale down its data over time. Since the host knows the importance of data it sends, it can defer (or set higher delay values in the priority registers) data transmissions of lesser importance to a later time.

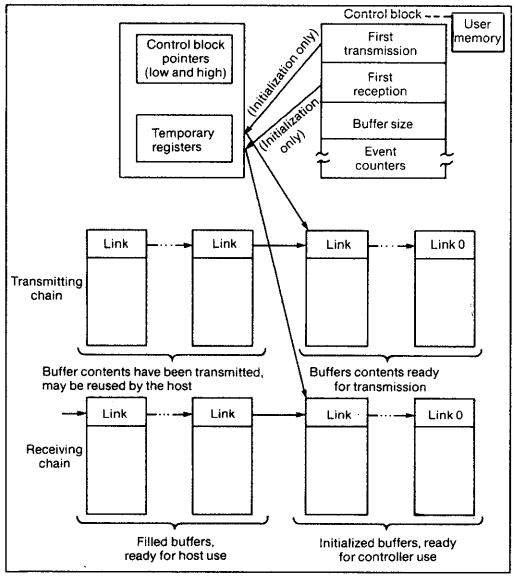

3. The chaining technique allows either the host processor or the controller chip to vary the number of buffers in a system. In chaining, buffers are linked so that data frames may span multiple buffers, making memory operations more efficient and simplifies the host's memory allocation tasks.

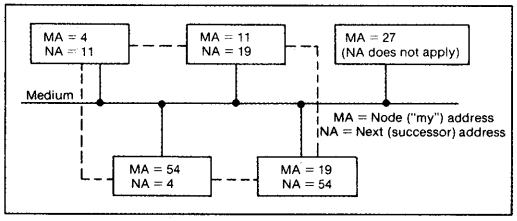

4. In a logical ring, a token can pass from station 4 (MA-4) to 11, to 19, to 54, and back to 4. The physical order of the ring is irrelevant, since token passing is based only on station addresses.

If transmission proceeds, messages are sent automatically to their appropriate destination addresses, with acknowledgment requests optionally encoded into each frame's header. Such transmissions continue for each frame queued until either a preprogrammed limit or the end of a transmission chain is reached (preprogramming is an optional priority feature). When transmission is complete, the controller passes the token to the next station. After frames having an acknowledge option are transmitted, the transmitting controller awaits a response from the intended receiver. Responses can be positive (indicating that the frame was received correctly), negative, or nonexistent. In the last situation, the receiving station either received the frame incorrectly or was out of service. The waiting period is controlled by an interval timer. If a time out occurs from a no-response condition, the WD2840 automatically retries the transmission.

Automatic retransmissions overcome most network noise glitches quickly and automatically. In this case, the host makes no decisions. If the retry is unsuccessful, the frame is tagged and an interrupt is generated that allows the host to decide the disposition of the frame. To avoid holding up the network, the tagged frame is passed over in the transmission chain and the next frame's transmission is attempted.

A universal device

The WD2840 interfaces with any conventional, general-purpose microprocessor or minicomputer bus. Equally as important as that capability are the chip's contributions to host-speed independence (decoupling) and system efficiency. To achieve these goals, the device combines I/O register programming, interrupts, and dual DMA interfaces for data, exception reporting, and extended control (see "A Three-Controller Architecture").

The memory interface is a self-contained subsystem that consists of two sets of 16-bit registers, byte counters, and DMA control logic. Backing this up is internal intelligence that interprets and manipulates the high-level buffer control structures.

Fetching and storing user data destined for or received from the network are the most important functions of the DMA system. However, these are not simple tasks because of the speed decoupling

needed between the real-time controller and the non-real-time host processes. The WD2840 solves the problem using an open-ended FIFO (first-in, first-out) method of buffer chaining (Fig. 3). In chaining, either the controller or host adds or removes buffers as they consume or generate data, but neither need be concerned about the state of the other.

Though residents of the host's buffer memory, chains are visible to the token-access chip. They are constructed by the host from linked buffers prior to the controller's initialization. Linking the buffers maximizes use of the memory by permitting frames to span multiple buffers, an advantage when most frames are short but long frames must occasionally be accommodated. When receiving a frame, the chip fills the receiving buffer pointed to by an internal register until the frame is complete. If the buffer is filled and the frame is not complete, the device automatically reads the link field of the filled buffer to find the next one available and continues receiving. Of particular importance when interfacing with an existing operating system, this automatic link handling of variable-sized buffers simplifies the host's task of allocating memory.

The transmitting and receiving chains' linked buffers are maintained cooperatively through the use of control fields located in the first buffer of each frame. This header information includes frame status, destination address, and actual frame length in bytes. The length, which determines how many buffers the associated frame spans, is written by the host in the transmitting chain. Each device has its own status bytes and can only read the status of the other device, thus preventing deadlocks.

The control field written by the host is called the frame-control byte and determines what options to put into a frame. An example of a per-frame option is the wait-for-acknowledgment command bit. This bit is tested while the frame is transmitted; if set, it causes the controller to await a response after a frame transmission is complete. The control byte written by the controller is called the frame-status byte. It indicates receiving or transmitting status, including the received "negative acknowledgment,"

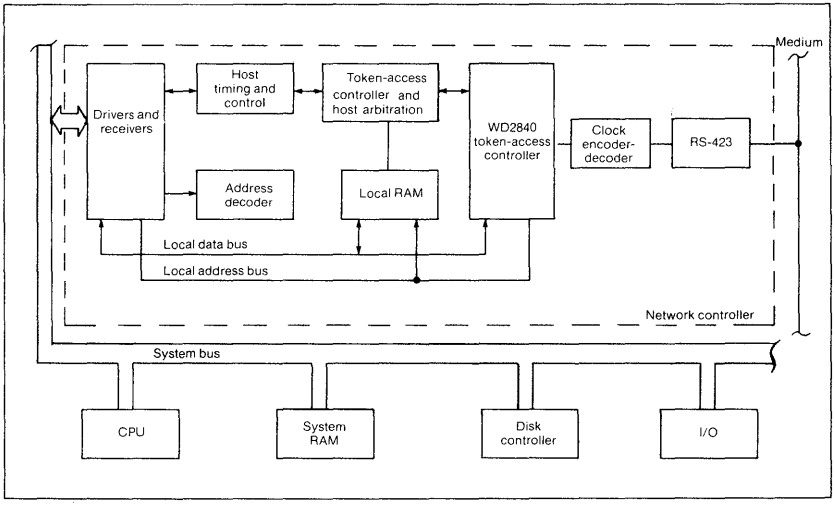

5. A high-performance word-processing application needs a two-ported memory interface between the token-access controller and the local network. This minimizes memory access latency time when the disk controller "hogs" the DMA bus for several consecutive cycles.

if any has been received.

Also contained in the logical memory interface is a series of 8-bit event counters. Located in the host's visible memory in the last section of the WD2840's control block (Fig. 3), they tabulate all noncritical but important network conditions. One such condition is the detection of a transmission error; by the time the host learns of the error, the controller has retransmitted. However, the chip increments the appropriate counter, since the event is valuable to the host for diagnostic purposes.

There are 16 control registers visible to the host (see table). The register file stores initial timer values and such fixed station parameters as station addresses and transmission limits. Included is a pointer to the initial DMA control block, which is used both when initializing the controller and when network exceptions are encountered. The register file also contains locations not visible to the host that are used by the internal controller for scratchpad functions, including the 16-bit pointers to the active receiving and transmitting buffers.

The controller's network (serial) interface accommodates standard modems and clock encoders. The device accepts non-return-to-zero (NRZ) data and permits the transmitting and receiving clocks to be stopped immediately upon completion of a frame if desired. Ordinarily this type of operation is not allowed, as most conventional devices use the clocks to clear the internal shift registers and perform other functions. A pair of request-to-send and clear-to-send signals is available for externally generating preambles for any type of medium.

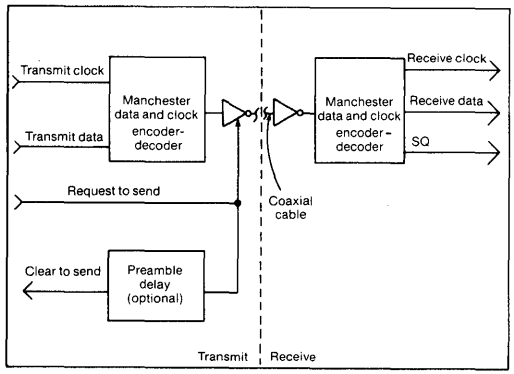

6. One of the simplest network interfaces is the Manchester encoder-decoder, which operates through an RS-442 bus transceiver. A nonvalid Manchester output signal connects to the controller's signal-quality pin.

Locating errors

Furnishing the network interface with novel features, the signal-quality input warns the controller that the frame in progress contains an error due to a media, modem, or clock-recovery fault. An example of such a error is the detection of a missed clock transition by a Manchester decoder (see "Manchester II Transfers Data with Integrity, Speed," ELECTRONIC DESIGN, March 19, 1981, p. 233). Coupled with appropriate external logic (optional), this input enhances data integrity beyond that offered by the frame-check sequence included with each transmission.

A common problem in token systems involves the actions necessary if a station is removed or fails; that is, what if the logical ring is broken? Prompt correction of such a condition affects the entire network and is therefore handled autonomously by the controller chip (Fig. 4).

The primary responsibility for assuring that a token arrives at its intended destination rests with the controller chip that sent it. Similarly, if the token does not arrive, the sending device also must retrieve it. To accomplish this, the sending controller sets its short internal timer to indicate the longest time that it must wait for a receiving station to use the token. If the token pass is unsuccessful—time expires on the internal timer without a successful pass—the chip begins error-recovery procedures.

Since the chief cause of these dropped tokens is a station improperly leaving the network, the WD2840 initiates a station scan. First it polls the address space for an active station to which it can pass the token. When one is found, the chip then updates its successor register so that the poll need not be repeated on the next cycle. After the logical ring is restored, the host is informed by means of a new-successor interrupt.

The WD2840 incorporates several types of diagnostics that are necessary in conventional intelligent subsystems and LSI systems. Network-level tests provide confidence and maintainability for the distributed token-access system and its associated medium and modems. System tests validate local interfaces, such as external RAM and the interrupt

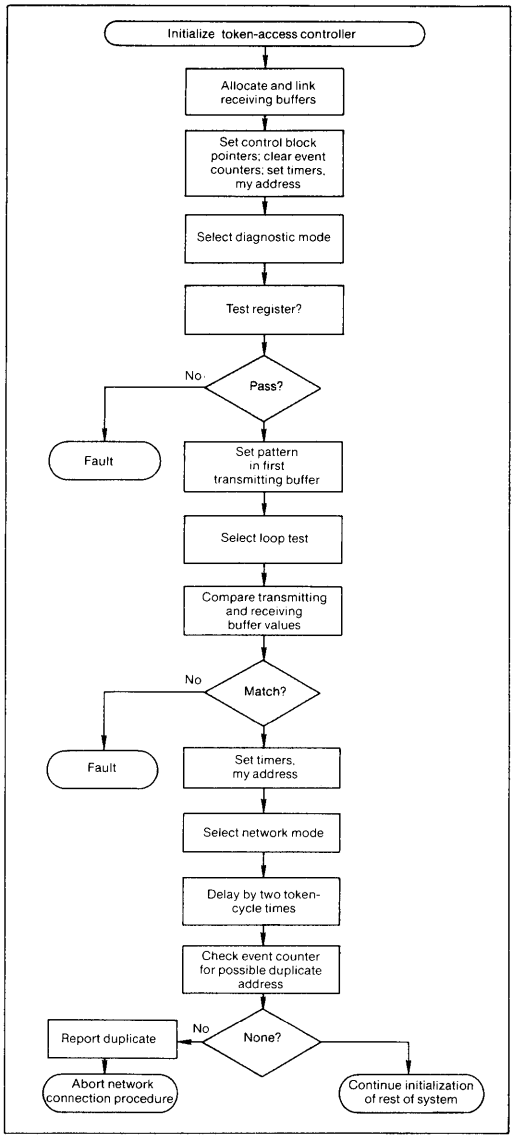

7. Initializing the token-access controller is necessary before a host processor can access a local network. The procedure includes buffer allocation, building a linked receiving chain, and performing diagnostic testing.

subsystem, and internal tests validate the chip itself, including the internal controllers and register file. All three diagnostics are used cooperatively.

As a new station powers up—or attaches itself to the network for any reason—it tests itself before transmitting on the network, thus ensuring that any faults do not disrupt the operating network. Self-testing is initiated by the host, which interprets the results. The host must be involved because an LSI device cannot always find and report its own failure.

Network testing occurs continuously. For example, a token-access controller watches for frames transmitted on the network by another station having the chip's own address as its destination. Source stations having the same address are prohibited from a token network. If that occurs, there is usually a hardware failure in the station or a misconfiguration (setting DIP switches incorrectly).

Local networks in word processing

A common application of a distributed network is multiple-user word processing. A typical system consists of a combined file server and print station connected to several remote CRT terminals. Each node contains its own processor and controller, and all communicate via the network. The CRTs are essentially "dumb" terminals having hardware modifications and internal firmware extensions that permit network use. The most critical station is the file server.

CRT terminals can be added or removed from the network while it operates. Either the control chip in each station or the host software configures the network. The choice depends on the speed with which new stations must be admitted and the tolerance of the network to access delays. This application tolerates delays of about 100 ms, so a simple token-access–control polling method is used.

The file server has a two-ported memory interface that minimizes memory access latency resulting from the disk interface's so-called hogging mode. Controllers that retain the DMA bus for several consecutive cycles are in the hogging mode. The two-ported memory is physically located on the network interface module and appears logically in the host processor's address space (Fig. 5).

The host interface matches the timing of the normal microcomputer bus to that of the controller and its local memory. This includes the memory-mapping logic of the host's operating system. Arbitration logic controls access to the local RAM and ensures that simultaneous memory requests by the host and controller do not end in improper memory operations or timing deadlocks. The logic is designed for FIFO-type command priority, with ties awarded to the controller.

The media interface here is implemented in its simplest form: a Manchester encoder-decoder chip and an RS-449 (three-state) bus transceiver (Fig. 6). That device handles its own preamble generation and detection at the start of each frame, so that the clear-to-send pin of the controller is tied directly to its request-to-send pin. Even more, the Manchester part provides a nonvalid Manchester output that is tied to the controller's signal-quality input.

Simple modems of this type are suitable for operation over moderate distances—about 1 km at a 1-Mbit/s transmission rate using twisted-pair cable. Other commercially available modems use more elaborate techniques to increase message reliability and distance or for other types of media.

The software interface with the host's file manager has three phases: initialization, file transmission, and command reception. The controller's initialization and network maintenance routines are included in the operating system. The software receives incoming frames from the WD2840, checks for proper frame sequencing, and builds messages that are compatible with normal operating-system file requests.

Before the operating system can access the network, the controller must be initialized. That process consists of allocating buffers, building an initial linked receive chain, and performing self-diagnostics. The flow chart shown in Fig. 7 gives the sequence of events. After diagnostics are complete, the host's initialization routine clears the event counters and writes the address of the chip's control block into the latter's internal registers. It then stores the proper values in the controller's registers (station addresses, priority values, timer settings) and puts the chip into the network mode.

A driver removes the incoming frames from the receiving chain and then ties them into the operating system in response to a controller interrupt. These messages are then separated into network management and information request groups. Network management frames serve primarily for the orderly addition and removal of stations. The file server periodically polls the network address space to allow new members in, but all stations process station-removal requests as they occur. There are many ways to maintain a network, each having a tradeoff between simplicity and timeliness. For this application, new stations need not be added very rapidly, allowing for greater simplicity.

Information requests always include a sequence number added by the controller's driver. These numbers are used in conjunction with the chip's automatic acknowledgment and thus ensure data integrity: the controller makes certain that no requests are lost, and the driver filters possible duplicates. The resultant messages are then removed from the controller's receiving chain and reformatted (including blocking if needed) before being passed to the filer. As a background task, the receiver driver initializes and attaches any free buffers returned by the file manager to the head of the chain for future use. Moreover, the operating system can add new buffers to the pool as the load increases.

Transmission is initiated after the filer obtains previously requested data. The information is passed to the network data, which then formats it into controller-compatible buffers, adds the correct sequence number and destination address, and finally attaches it to the transmitting chain. As its background task, the transmitting driver periodically checks the chain for buffers that can be returned to the transmission pool. This is an option that can be performed whenever a frame-transmission interrupt occurs.□

A summary of register files

Register	Name	Function
0	CR0	Control register 0
1	CR1	Control register 1
2	SR0	Status register 0
3	IR0	Interrupt register
4	SR1	Status register 1
5	SR2	Status register 2
6	CTR0	Temp counter
7	NA	Next address
8	TA	Acknowledge timer
9	TD	Network dead timer
A	CPBH	Control block pointer (most significant byte)
B	CPBL	Control block pointer (least significant byte)
C	NAR	Next address, request
D	AHOLT	Access hold-off limit
E	TXLT	Transmit limit
F	MA	My address

Token-access controller minimizes network complexity

MARK STIEGLITZ, Western Digital Corp.

Users can benefit from the increased speed this transmission method provides

In a data-communications network, contention among stations trying to get through to the central computer is inevitable. One of the more effective procedures to eliminate this contention is a form of distributed polling known as token passing. Despite its effectiveness, however, token passing has not been very popular with system integrators. Most network architects have been intimidated by the complexity of the algorithms required to set up the station linkages and to recover from network exception conditions, and have settled for less complex control methods. A new LSI token-access controller (TAC) residing in each station of the network minimizes this complexity for network designers.

A token is a message granting a polled station the temporary but exclusive right to transmit on the medium, a right the station must then relinquish to the next designated station. This method has been historically used on sequential media on which access sequence is implied by the physical interconnection, but tokens can also be used on broadcast media such as baseband coaxial or CATV systems by assigning unique addresses to each station or node (MA: "my address") and passing transmission rights between them (Fig. 1).

The simplicity and non-reliance on quirks of a medium make token methods superior for use on a wide array of applications. Relatively simple (from the data-movement viewpoint) applications such as file transfer to the complex time-critical applications of factory automation are supported with the same access

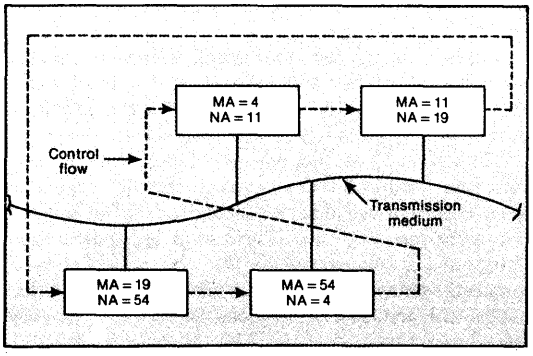

Fig. 1. Access control flow. *When a station (MA: "my address") has transmitted its data, it sends the transmission rights—the token—to the station identified in the next address (NA) register. Station numbers are in ascending order but need not be sequential for network efficiency.*

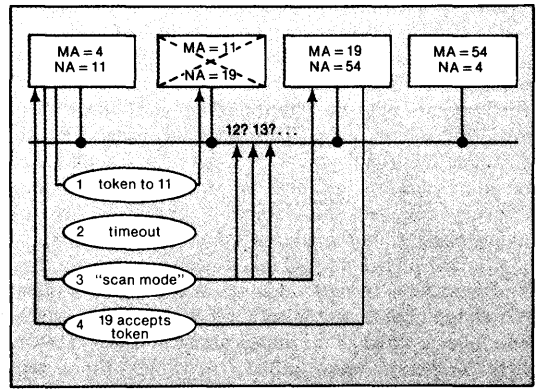

Fig. 2. Station dropout. *Station 4 attempts to pass the token to 11, which has dropped out of the network. Station 4 then "times out" and scans for another station to which it can pass the token, finding 19.*

TAC must handle three main exception conditions: network initialization and recovery from failed nodes, addition of stations to the access ring while the network is in use and recovery from an error situation in which two or more tokens have been generated on the network.

protocol. Also, data rates optimized for the application, not mandated by the network implementation, are possible with the same LSI network controllers.

TAC's tasks

TAC must handle three main exception conditions: network initialization and recovery from failed nodes, addition of stations to the access ring while the network is in use and recovery from an error situation in which two or more tokens have been generated on the network.

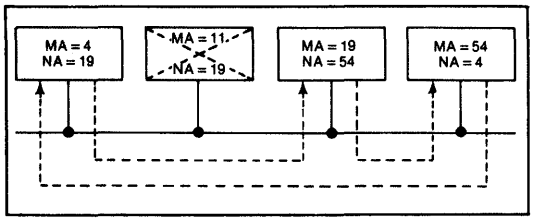

Fig. 3. Station patched out. *Station 11 is logically removed from the network when station 4 changes its next address (NA) register to 19. Thus, station 11 no longer consumes network time.*

Initialization and recovery

Initialization is setting up the network's token linkages and determining the correct values for registers in all TACs wanting to be part of the access ring (desiring INRING status). Failed-node recovery refers to the network restart when a token is lost or damaged.

Token loss results from exception or expected conditions. Error cases may be a product of noise hits on the transmission medium corrupting the token message, or simply a controller failure. A token loss usually occurs with the intentional removal of a station from the access ring. The ratio of noise hits to controller failures depends on a network's application and administration, but both are recovered identically. Recovery requires the detecting station to set its linkage register to the address of an active station.

Initialization is a form of failed-node recovery in that all of the access linkage registers of the network must be updated. Before initialization it is not known which node follows which. Two timers assist in these cases: a fairly long-time value called TD, which times out network inactivity and a shorter timer called TA, which is the maximum turnaround time required for a response (token or data) to be sent by the receiving station of a previous message. These timers are user settable and depend greatly on a network-transmission rate, and to a lesser extent, on an application. The timers work together and are the key to solving the initialization and failed-node challenges.

Two manifestations of a failed node can occur. One happens when a token holder tries to pass the token to the next station in the ring. If the next station does not respond to the token, the token-passing station soon knows because it knows how long it should take for a node to pass the token or to send a data message (time TA). In this case, the node that tries to pass the token has primary responsibility to recover. It does so by entering a scan mode from an access level, and polling the network for another successor.

Assume that station 11 (Fig. 2) is removed from the network and station 4 is attempting to pass the token to it. Station 4 will time out because 11 does not respond to the token within time TA and will attempt recovery by passing the token to station 12. Station 12 will not respond because it also is not present, which will cause

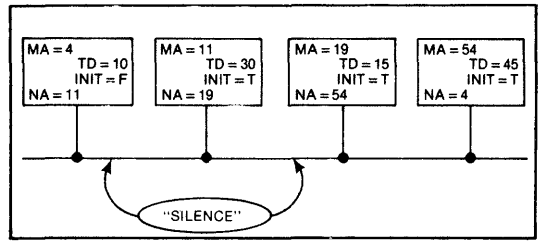

Fig. 4. With no token, there is no transmission. *All stations detect this and start internal timers. When one expires and has the proper control bit enabled, it restarts the network. TD is the station-inactivity timer. INIT is a switch that, when false (F), tells TAC not to attempt recovery regardless of TD. When INIT is true (T), TAC attempts to initialize the network after there has been no activity for duration of TD.*

station 4 to try 13. This "polling" continues by station 4 until it finally gets to 19. Station 19 will respond, causing station 4 to update its next address register (NA) to 19, bypassing station 11. The next time station 4 gets the token, it immediately passes it to 19 (after sending any messages).

At this stage, station 11 is logically removed from the network, or "patched out" (Fig. 3). If station 11 wants to get back into the network later, the standard station-adding procedures must be followed.

This station-by-station access polling consumes network time (each poll takes TA time), and may appear to be an inefficient use of network bandwidth. But this is a rare error-recovery case. Further, it is handled completely and autonomously by the TAC, which at least bounds the delay. The host μp is not burdened with this critical task and, as a result, does not slow the recovery procedure.

The second failed-node manifestation occurs when a station holding the token itself fails before it has a chance to pass the token to another. If station 4 has the token (Fig. 4) and dies before passing it, no activity

Additional stations can be added to an operating token network at any time. The distributed method does not rely on a specific station. Thus, there are no problems or efforts spent selecting the administrator.

occurs on the network and no station has the short timer (TA) running. All stations, however, have the timer TD, or network dead timer, running. The station whose TD timer expires first takes recovery responsibility.

To simplify network administration, not all stations must be able to reinitialize the network. The first station whose network timer TD times out tests the control bit "INIT," saying, in effect, "when timer TD expires, should I claim the token?" If INIT is false, the station waits, as does every other station, for a station's TD to expire that has INIT true.

One station on the network's timer TD that has the ability to initialize will eventually expire. That station will claim the token and send its messages, or send the token to its successor station as directed by its NA register. Thus, if station 19 happens to have the shorter timer TD and has its initialize enable bit set (Fig. 3), station 19 assumes the token, sends whatever messages it had queued and sends the token to station 54.

On receiving the token, station 54 (Fig. 4) sends its messages and tries to pass the token to station 4. If station 4 has recovered from its problem (its failure caused this recovery condition), it receives the token and transmits with it.

If station 4 still does not transmit, station 54 has primary recovery responsibility (54 has started its timer TA) and will enter the scan method. The scan starts at station 5 and searches until it finds the next available on-line station (11 in this example).

The power-up initialization case behaves in the same manner. As stations come up, they wait for a message or for their timer TD to expire.

Host software is responsible for setting up the next address register before enabling the transmitter in the TAC. This is set to the station address plus one (which will in effect cause a polling by that station) or, if it has some prior knowledge of what the network configuration looks like, it sets NA to reflect the correct address of the successor. Host software is also responsible for setting the time-out values in the recovery timers (TA and TD). The value for TA should be consistent among all stations of a network, but TD is not critical, and thus may vary greatly because it is used only in exception situations.

For example, there can be half a dozen stations on the network that are intended to recover from catastrophic conditions such as loss of token. These stations can all have substantially different time values TD so that if a couple of them are not on-line at a time, one will come up and reinitialize the network.

Additional stations can be added to an operating token network at any time. If a supervisory communication path can be assumed, a candidate station requests of the administrator that it be admitted to the access ring. This approach is not unlike the method pay-TV companies use to enable new subscribers' decoder boxes. When installed, a service representative of the cable-TV company telephones (the supervisory communications method) the central site, which then sends the properly addressed enabling signal over the network. While this method is efficient from the network viewpoint (the infrequent control messages are handled "out of band"), such duplicate communications schemes do not usually exist.

A more acceptable solution is to allow the control communications to share the data bandwidth. To avoid data collisions and retain the prized asset of a token system—determinism—new stations are added on a controlled-polling basis. To accomplish this, the TAC requires the host to initiate the test for a new station. Although, in this case, host interaction is required to expand the network, that interaction doesn't set back the goal of autonomous TAC network control in that adding stations is not a real-time requirement. The time to add a new station is not critical to the performance of the rest of the network.

There are three primary methods by which a station can be added to a network. The first is a distributed method, in which each station in the network can poll for new stations in the gap between its address and the next address (between MA and NA). Second is a centralized method, in which an individual station designated by the network architect can interrogate the entire address space seeking a new station desiring INRING. The third—central seam—is a simpler (from the host point of view), centralized method in which a station can send a global frame causing all the on-line TACs to reset their next address register. This causes each TAC to poll its address space at its next token-pass attempt. Each method has advantages and disadvantages.

The distributed method does not rely on a specific station. Thus, there are no problems or efforts spent selecting the administrator, nor is there any concern about backup administrators. In the distributive method, each station has the same responsibility to allow new access members as other stations. This method is the most host intensive and requires each station to maintain a timer (that can be configuration set as to its value) as to how often it should poll its gap for new stations.

For example, assume the timer in each station is 5 sec. and that station 4's timer has expired (Fig. 5). The host attached to station 4 notes that the next address register (NA in the TAC) is set to 11, which indicates that a new station might be added to the network as station number 5, 6, 7, 8, 9 or 10.

A token is a message granting a polled station the temporary but exclusive right to transmit on the medium, a right the station must then relinquish to the next designated station.

The host queues a frame into the TAC transmit chain, polling station 5. This frame will be sent by 4 with an acknowledgement requested from 5. If 5 is present it responds; otherwise, the TAC aborts its attempt after time TA. The TAC marks the result on the frame in the host memory space and proceeds with other tasks.

After this exchange, the host, at its leisure, checks

IMPLEMENTING TAC

The TAC is a single-chip NMOS LSI device that performs all real-time communication tasks in a μp-based system. The assumed existence of a μp allows some less critical, non-network performance-affecting tasks to be performed outside the TAC, such as flow control and adding new stations.

Removing these functions from the device results in:

- Less processing power, which, in turn, makes the chip smaller and less expensive;
- No processing burdens, enabling the TAC to respond faster to network conditions, thus improving efficiency;
- Saved firmware space and processing power, which can be used for internal diagnostics and a more sophisticated host interface—a chained frame buffer scheme, which is a trade-off in favor of system efficiency.

To meet the network requirements and to include the other features expected in LSI, such as internal validation, a three-processor design was used consisting of a primary microcontroller, a receiver and a transmitter.

The primary microcontroller performs all token-algorithm support such as network initialization and error recovery, manages host interrupts and coordinates internal and system diagnostics. It also evaluates the host commands and arms and supervises the receiver and transmitter microcontrollers.

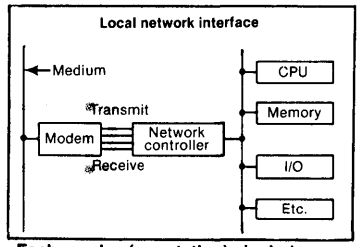

Each node (or station) includes a modem, *a network controller (TAC: token access controller) and appropriate hardware as required by the application ("host").*

The receiver does minor frame (group of bytes) filtering and frame validation and, independently of and simultaneously with the primary controller, performs DMA operations storing incoming data. The transmitter sends data via its DMA interface when allowed by the main controller, that is, when a token is received.

The register file is used by the host to set memory pointers, network address registers, long-term parameters such as frame transmit limits (allowing users to select exhaustive or non-exhaustive transmission) and the conventional command, status and interrupt indications. The TAC's primary interface to the host is its DMA system. Data to and from the network and options selectable on a frame-by-frame basis are read in this manner.

The half-duplex network interface has standard RTS/CTS handshakes. Another feature of the receiver is a signal-quality input that allows errors that are easily detected by the modem (such as a missing clock detected in a Manchester decoder or low carrier in a broadband system) to be signaled to the TAC. The use of these low-level checks further enhances the basic frame integrity beyond that of the CRC.

Messages are sent between stations on the network in frames. The frame structures are similar to the industry standard HDLC; delimiters are unique flag patterns with zero insertion used for data transparency. In addition to adding the required control fields to support the token protocol, the TAC recognizes three basic frame types: a short token pass frame, a short frame conveying only acknowledgement and control information and variable-length frames holding user information and optional network control information.

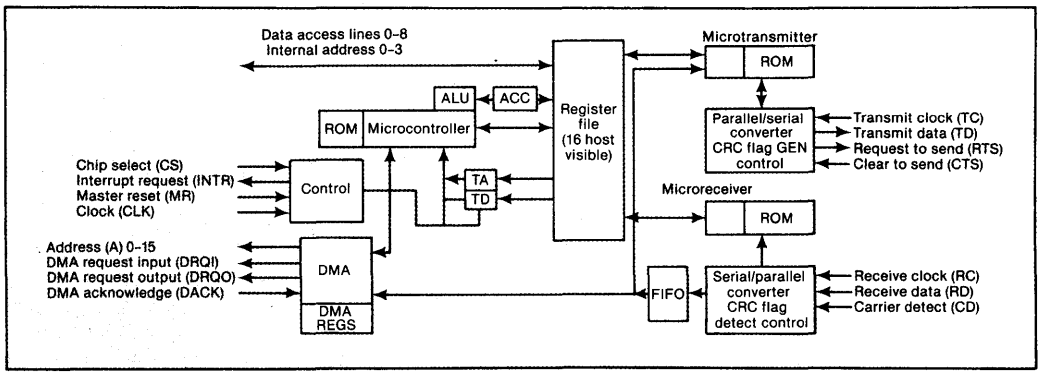

TAC includes three processors implemented in a single LSI device *that interfaces with the host through a register file and a DMA subsystem.*

In the centralized station addition method, a single station can poll the entire address space, seeking a new station that desires INRING.

the transmit status of the frame. The host sees that the frame acknowledgement timed out, meaning that station 5 has not been added to the network, or that station 5 is on the network and whether the request INRING is set in the network code field. In either case, the host takes appropriate action. If the desired INRING bit is set, station 4 changes its NA register to 5, allowing its next token to be passed to 5. This action puts station 5 in the ring.

Depending on an application's sophistication, a control message can be sent to station 5. That message says, "Your successor is X." In this case, X = 11, so that 5 is not forced to poll for its successor. In any case, 4 updates its next address register to 5 and does not need to go through this distributive polling cycle again

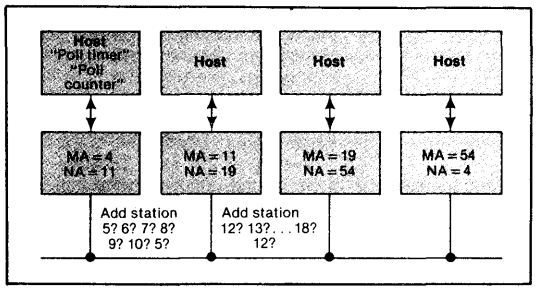

Fig. 5. Distributed polling. *Each host polls the gap in its address space for the possible addition of new stations. The host internal poll timer and poll counter set the polling rate and range as desired.*

because there is no gap between 5's address and the next address; there is no possibility that a new station can be inserted between addresses and 5. If 5 didn't respond to 4's poll, station 4 updates its poll counter so that the next time that the poll timer times out, station 6 will be tried.

If node 6 responds, its desired INRING bit is tested as above. If 6 does not respond, the host will queue a poll to station 7 the next time its poll timer expires. This continues until the host completes 10, when the cycle goes back to 5 and repeats. In this example, with a gap of 6 stations (between 4 and 11), and with a 5-sec. clock, a new node can be added within 30 sec.

In the centralized station-addition method, a single station can poll the entire address space, seeking a new station that desires INRING. One reason for centralizing this function might be the more careful control that can be placed in a network. There can also be optimizations. For example, the central polling station can keep track of the stations that already exist and, therefore, bypass some address ranges. A polling station may know the network will never have more than, say, 75 stations. In the example of Fig. 6, when station 4 starts polling, it polls only to address 75 before resetting to zero. This works like the distributed method except that a single station does all the work.

When the polling station determines that a station has been added, it must place the new station in the access ring. For example, station 4 is the centralized station doing all the polling (Fig. 6), and it discovers that station 27 has recently been added. Station 4 knows this because station 27 now responds to a first-time poll, and because its status bit is set, indicating that it wants to be added to the ring. (Some stations may be receive only, never desiring the right to initiate transmissions.) Station 4 sends a high-level message to the software in station 19, telling it to change its next address register to 27. This message

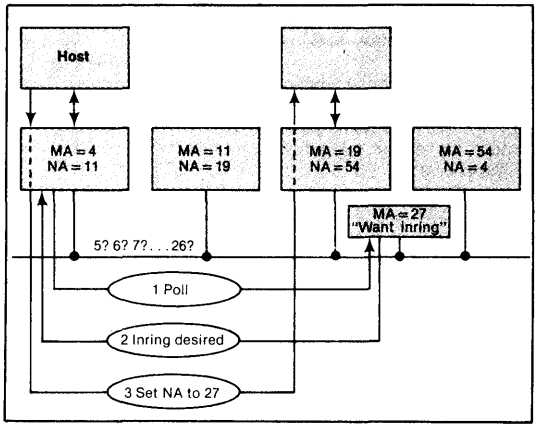

Fig. 6. Central polling. *A single station—in this case, station 4—dubbed "the administrator," can be charged with all polling tasks. This simplifies the software in the other stations and centralizes network control.*

can also prompt station 19 to tell 27 its next address register should be 54. This gets confusing, but it is all done with high-level software. These tasks are not real time and are quite efficient from the network point of view.

Station 4, the administrator, need not create and maintain a table of active stations on the network because the poll response returns three pieces of information. As node 4 polls the stations on the network, it finds out (a) that the polled station does not respond at all, as it would if it polled station 12 in Fig. 6; (b) that the station is already part of the network and is already in the ring or is receive only, as it would if station 4 happened to poll station 11 or 19; and (c) whether the station is attached to the network, is alive and wants to be in the ring, as is the case with a poll to 27. These indications are conveyed by a combination of status bits sent back by the acknowledge frame. This acknowledge frame and status information are transferred at a TAC device level, so a host is not concerned with whether its station wants to be in the ring. The host simply sets up the proper bits in the control

To simplify network administration, not all stations must be able to reinitialize the network.

registers; the bits are relayed automatically by the TAC. Thus, with a simple algorithm, an administrative station can poll the entire network address range and know the network's exact membership and status.

Central scan

Central scan is the simplest method of adding stations to a network. It involves sending a global frame to all stations on the network, which forces each to update its own next address register to its station address plus one (NA = MA + 1). Assume station 4 is the centralized station and sends the scan command frame (Fig. 7). Station 11, upon receiving it, automatically sets its next address register to 12 (the TAC does this; the host is not involved but is notified of the situation). Also, station 19 sets its next address register to 20, and station 54 sets its NA register to 55.

The result of this is a round of polling at the TAC level. Station 11, on completing its use of the token, tries to send it to 12. The token to station 12 times out because 12 is not present. Station 11 reclaims the token trying to send it to 13 and so on, causing 11 to poll for station addition. The drawback of this is the huge time disruption incurred by the simultaneous polling.

It is not required that station 4 send this scan control frame to all stations at the same time. If it is known that station 11 exists in the network and that a station may be trying to add into the network after station 11 in the address space, a command can be sent to 11 telling it to set its next address register to 11 + 1. Now 11 will go through scanning station 12, 13, 14 ... again without intervention from station 11's host software. This directed scanning has the effect of smoothing the polling disturbance over a greater time.

The trade-off of all these methods is the software complexity distribution. If a TAC user assumes more responsibility, providing more intelligence distributed in the software, the system can be more sophisticated in handling new stations. If a user wants the TAC to handle this task itself, saving host software development, he pays only slightly in inefficiency. TAC gives the user an option.

Recovery from multiple tokens

Multiple tokens are not allowed on a token bus because their presence causes a breakdown of the orderly nature of the protocol. Their presence can only be the result of a combination of exception and hardware failure conditions but, once present, must be handled immediately.

The primary defense against multiple tokens is prevention. The control algorithms and the frame formats have been designed to minimize multiple tokens. For example, the TAC can refuse to allow a

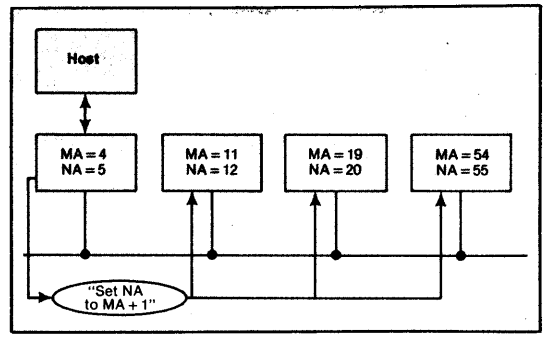

Fig. 7. Central scan request. *A special command can be sent by any station causing all attached TACs to set their NA register to the address of the next possible node. This causes each TAC to poll without the help of the host.*

piggyback token (a single frame containing both the token and a user-information field) with the data-acknowledge option. If this were allowed, conditions could result in which the data was negative acknowledged by its receiver and retransmitted, but the token arrived successfully at its destination—in this case, twice—creating two tokens.

Duplicate tokens, or at least network confusion, can result from more than one station having the same network address. Unless the stations are receive only, their simultaneous responses to data frames and tokens will probably result in their response not being accepted. While conceptually simple to prevent, address duplication can be the result of hardware failure (a bad DIP switch), operator error or configuration error (if a device is moved from one network to another).

Because the access controller must monitor the network for messages addressed to itself anyway, it is simple to check for messages sent by a station with its address (most frames contain both a source and a destination address). Part of a host's attachment algorithm would normally check this counter in the TAC before allowing it to transmit anything, thereby catching most of these duplicate station faults before they have a chance to affect the network.

A token access controller can also detect duplicate tokens by knowing that, when it has the token, no other station can transmit. This ability is supported in the TAC by incorporating separate receive, transmit and control sub-controllers. This allows the receiver to monitor the medium while the primary controller is, for example, searching the host's memory for a frame to be sent. If another token exists or is suspected, the TAC drops its token, allowing the other to circulate. If there is no other token, the network is left in a no token state and is easily restarted with the aid of recovery timer TD. ■

Mark Stieglitz is manager, local networks, Western Digital Corp., Newport Beach, Calif., and chairman of the IEEE committee working to set a token-passing protocol standard.

Quality/Reliability To Leading Edge Technology

QUALITY PROGRAM DESCRIPTION

The Quality Organization shown in Figure 2 assures compliance to design control, quality and reliability specifications, pursuant to corporate policy.

CORPORATE QUALITY POLICY

It is the policy of Western Digital Corporation that every employee be committed to quality excellence in producing products/processes which conform to acceptable requirements. The total quality program is managed and monitored by the quality assurance organization. Quality assurance is chartered to review marketing product requirements, qualify hardware and software designs, certify manufacturing operations and monitor performance/control conformance to product specifications.

Primary responsibility for execution for the quality program rests with functional organizations to design, produce and market high quality and high reliability products specified to our customers.

LSI QUALITY ASSURANCE PROGRAM HIGHLIGHTS

- LSI manufacturing assurance provisions are derived in part from MIL-M-38510 and MIL-STD-883B as applied to high grade commercial components.
- All process raw materials used in the Mask/Wafer fabrication and assembly operations are monitored by Material Assurance.
- Material Assurance maintains a thorough control of incoming material and has developed unique "use/stress tests" (look ahead sample build acceptance) which critical material must pass before acceptance.
- The Product Assurance Department continuously monitors the internal and external manufacturing flow (shown in Figure 1) and issues process control reports displaying detailed data and trends for the associated areas.
  - — Document control is an integral part of Product Assurance. All specifications are issued and controlled by this activity.
  - — The Western Digital Malaysian assembly operation uses specifications and quality control provisions controlled by Document Control. Indicators of Malaysia quality are reviewed weekly.
  - — Purchased FAB and assembly operations are individually qualified and are certified against standard specifications during vendor qualification and monitored against reliability criteria.
  - — Defect control within the process assures the highest levels of built-in reliability.
- Quality audits and gates are located throughout the manufacturing process in order to assure a stable process and thus, a quality product to our customers. Figure 1 illustrates the manufacturing/screening/inspection flow diagram and identifies the steps as they relate to the production of LSI devices.
- Testing assures quality margins through 100% testing by manufacturing and, in addition, all products must pass a specified AQL sample test performed by QA at maximum operating temperature as follows:

Outgoing Quality Levels

SUBGROUPS	INSPECTION LEVEL
Subgroup 1 — Final 100% Electrical Audit @ Max °C	0.5 AQL*
Subgroup 2 — Visual (Marking, Lead Integrity, Package, Verify customer shipper)	1.0 AQL
Subgroup 3 — Shipping Visual Audit	1.0 AQL

*The double sampling techniques used allow considerably better AQL's in most all cases.

- LSI devices are 100% tested on industry standard test systems like that shown below. Quality outgoing testing (auditing) is done on the Fairchild Sentry Series 20 where possible to allow better correlation with customers.

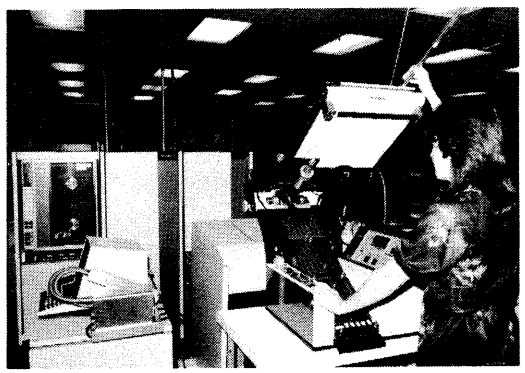

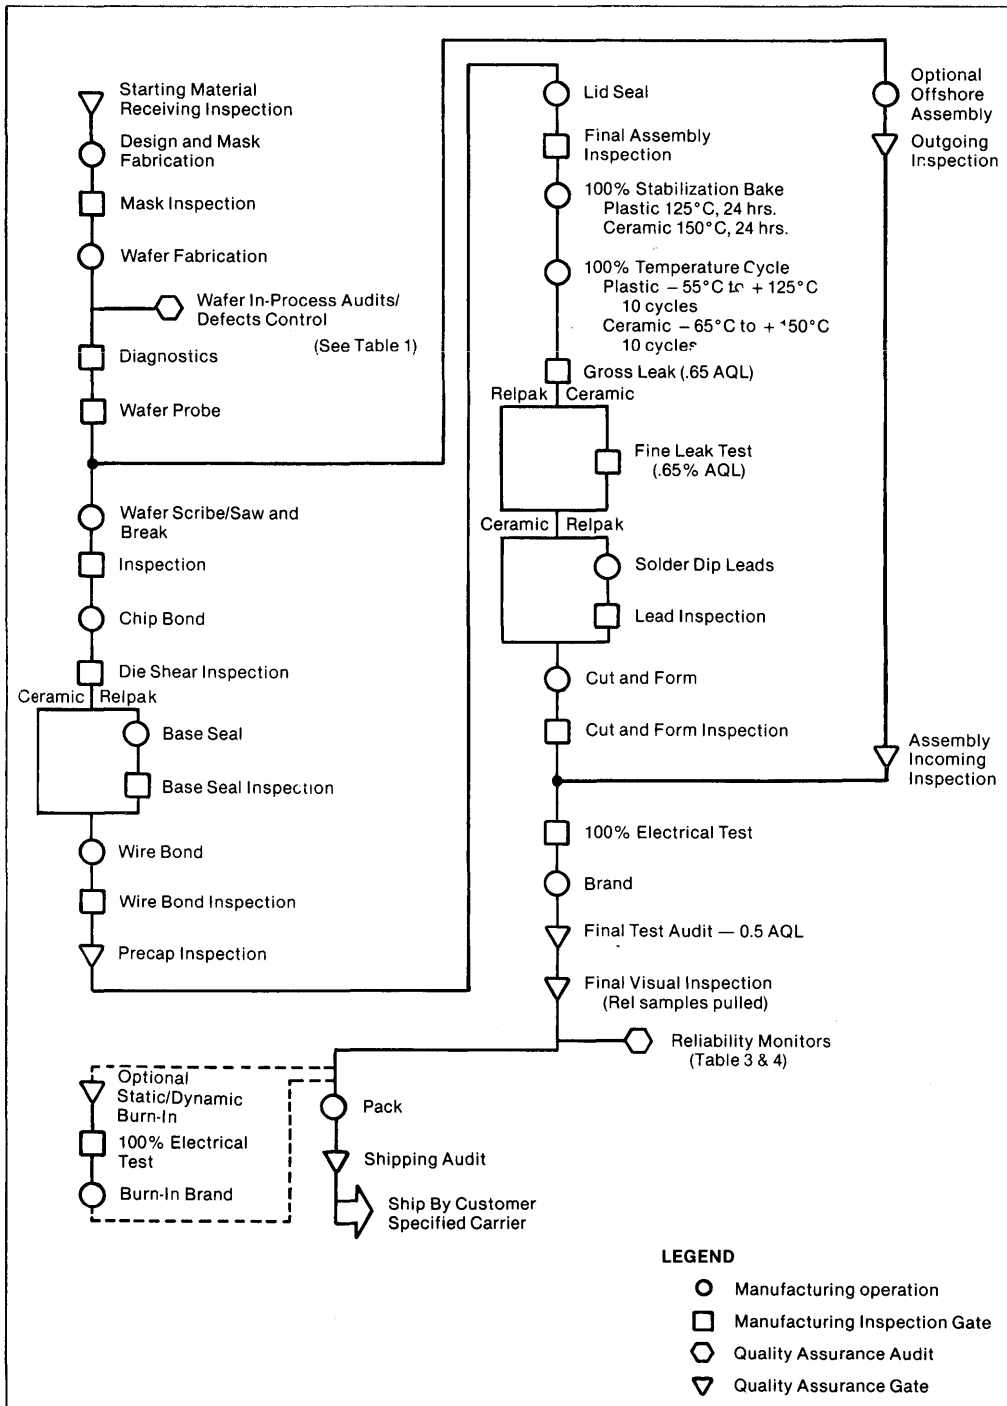

Figure 1 LSI PRODUCTION FLOW

Reliability Means Lasting Value

• DESIGNING FOR RELIABILITY

The production release procedure for an LSI device is designed to assure maximum reliability with a Quality checklist for:

- ☑ Test program qualifications
- ☑ Characterization report
- ☑ Field test (Beta Test) report
- ☑ Reliability Lifetest Qualifications
- ☑ Infrared Thermal Analysis
- ☑ Static Protection

All new devices and major process changes must pass reliability qualification before incorporation into production using the criteria defined in Tables 2-4. The infrared microscope shown on the right assures optimum burn-in temperatures and margins of safety. The dynamic burn-in system shown on the right is one of two custom designed systems which assure protective device isolation during burn-in.

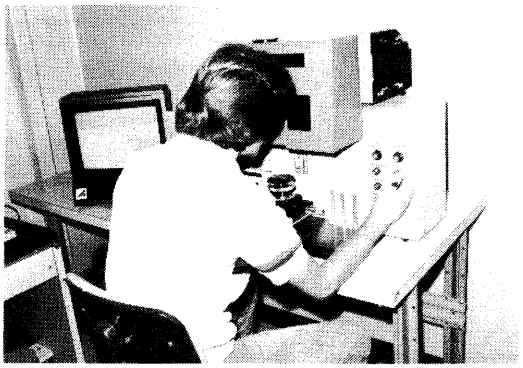

• MAINTAINING RELIABILITY IN PRODUCTION

Process defects control are defined to continually measure built-in reliability, as measured by the following criteria:

TABLE 1

PROCESS RELIABILITY CONTROL	METHOD	CONDITION	SAMPLE*
Subgroup 1 — Defects Control			
a. Oxide Integrity	Non-destructive bubble test	Pinhole defect density	5 wafers
b. Polysilicon Integrity	SEM Analysis	Visual	5 wafers
Subgroup 2 — Electro-Migration Control			
Metal Step Coverage	MIL-STD-883 Method 2018	SEM Analysis	5 wafers
Subgroup 3 — Defect Density	Critical layers Field Gate Contact Metal	Visual of Photo defects (Defects/in^2)	8 wafers each layer
Subgroup 4 — Passivation/Insulation Integrity	MIL-STD-883 Method 2021	Visual of Pinhole defect density	Final Silox 5 wafers Intermediate 5 wafers

*Inspection intervals are defined by the in-line process control data reviewed on a lot-by-lot basis.

• PROGRAMS TO ASSURE OPTIMUM RELIABILITY

Improved levels of reliability are available under custom reliability programs using static and dynamic burn-in to further improve reliability. These programs focus on MOS failure mechanisms as follows:

FAILURE MECHANISMS IN MOS

FAILURE MECHANISM	EFFECT ON DEVICE	ESTIMATED ACTIVATION ENERGY	SCREENING METHOD
Slow Trapping	Wearout	1.0 eV	Static Burn-In
Contamination	Wearout/ Infant	1.4 eV	Static Burn-In
Surface Charge	Wearout	0.5-1.0 eV	Static Burn-In
Polarization	Wearout	1.0 eV	Static Burn-In
Electromigration	Wearout	1.0 eV	Dynamic Burn-In
Microcracks	Random	—	100% Temp. Cycling
Contacts	Wearout/ Infant	—	Dynamic Burn-In
Oxide Defects	Infant/ Random	0.3 eV	Dynamic Burn-In at max. voltage
Electron Injection	Wearout	—	Low Temp. Voltage Operating Life

Temperature Acceleration of Failure

The Arrhenius Plot defines a failure rate proportional to exp(−Ea/kt) where Ea is the activation energy for the failure mechanism. The figure on the right indicates that lower activation energy failures are **not** effectively accelerated by temperature alone; hense, maximum voltage operation is selectively applied to optimize the burn-in process.

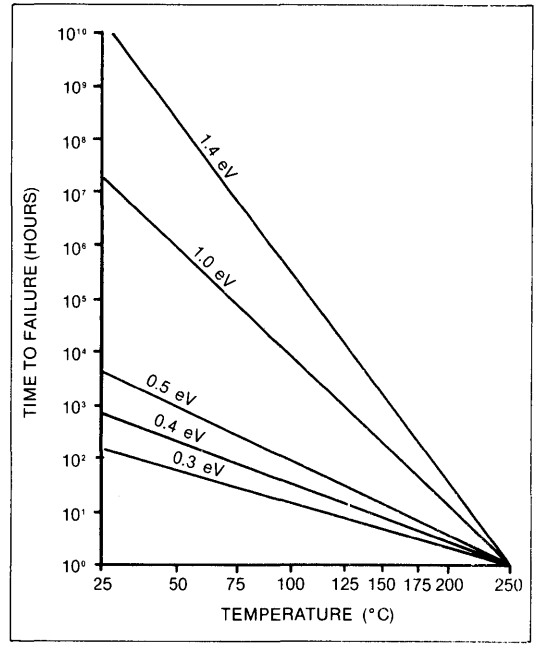

Static Burn-In (125°C — 48 hours or 160 hours)

Provided on a sample basis for process monitor/control of 0.5 eV — 1.0 eV failure mechanisms. 100% static burn-in may be specified at an additional cost. However, static burn-in is considered only partially effective for internal LSI gates at logic "O" levels.

Dynamic Burn-In (Pattern test/125°C — 8 hours to 160 hours)

Accelerated functional dynamic operating life effectively controls internal MOS gate defects buried from external pin access. The input pattern is optionally pseudo-random or fixed pattern programmable to simulate 1000-3000 hours of field operation at maximum operating voltage(s).

High-Rel "K" Testing Program

General conformance to MIL-STD-883B method 5004.4, Class B with static Burn-In (Dynamic Burn-In may be specified as an option).

LSI RELIABILITY STANDARDS

TABLE 2 STANDARD RELIABILITY LEVELS

TEST	METHOD	CONDITION	FAILURE
Infant Mortality (see note)	Static Burn-In	125°C — 160 hrs.	<0.5%
Long Term Failure Rate	Dynamic Life Test	125°C — 1000 hrs.	<.05%/1000 hrs. @ 55°C 60% Confidence

*NOTE: Devices failing the infant mortality target remain on burn-in until acceptable failure rates are obtained.

TABLE 3 GROUP A DEVICE RELIABILITY MONITORS

TEST	METHOD	CONDITIONS	LTPD
Subgroup 1			
a. Internal Visual			15
b. Thermal Shock	1011	Test Failure Used (cond. B or C)	
c. Bond Strength	2011	Test Failures (cond. B)	
d. Die Shear Strength	2019	Test Failures	
Subgroup 2			
a. Seal — Gross Leak		Fluorocarbon detection 10^{-3} atm/cc/sec	15
b. Seal — Fine Leak	1014	Test Condition A	
Subgroup 3			
a. Rotating Steady State Life Test	1005	Static 160 hr. Burn-In 125°C plus 125°C Lifetest — 1000 hrs.	5
b. Electrical Parameters	—	Final electrical @ 25°C (with data @ 70°C)	

TABLE 4 GROUP B PACKAGE RELIABILITY MONITORS

TEST	METHOD	CONDITIONS	LTPD
Subgroup 1			
a. Thermal Shock	1011	Test Condition B or C	15
b. Temperature Cycling	1010	Test Condition B or C	
c. Seal — Gross Leak	—	Fluorocarbon detection 10^{-3} atm/cc/sec	
d. Seal — Fine Leak (ceramic)	1014	Test Condition A	
e. Electrical Parameters	—	Electrical at max -C	
f. 85/85 Moisture Resistance (plastic only)	—	85% RH/85°C for 1000 hours PDA = 10%	
g. Electrical Parameters	—	Final electrical @ 25°C	
Subgroup 2			
a. High Temp. Storage	1008	Test Condition B or C	
b. Mechanical Shock	2002	Test Condition B	15
c. Seal — Gross Leak	—	Fluorocarbon detection 10^{-3} atm/cc/sec	
d. Seal — Fine Leak (ceramic)	1014	Test Condition A	
e. Electrical Parameters	—	Final electrical @ 25°C/max. C	
Subgroup 3			
a. Lead Integrity	2004	Test Condition B2 (Lead Fatigue)	15
b. Seal — Gross Leak	—	Fluorocarbon detection 10^{-3} atm/cc/sec	
c. Seal — Fine Leak (ceramic)	1014	Test Condition A	

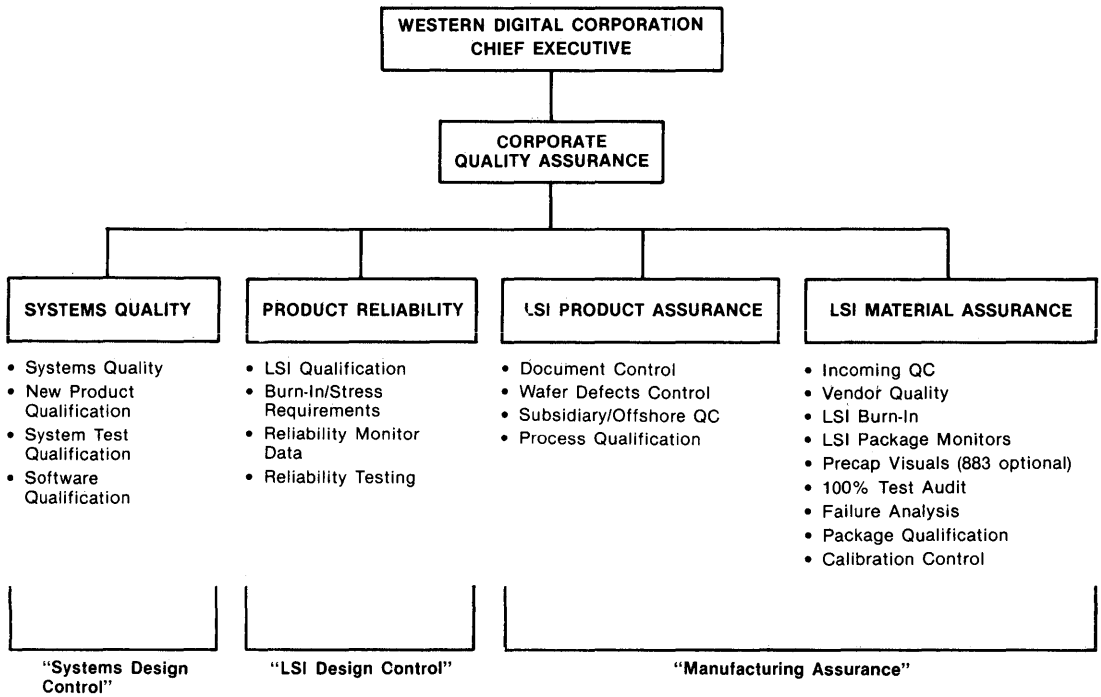

Figure 2 QUALITY ORGANIZATION

Printed in U.S.A.

Announcing Burn-In Program Availability/Warranties

Western Digital now supports customer burn-in requirements for both static and dynamic burn-in under the strict control of the QA-Reliability organization.

This burn-in provides high performance 125°C static and dynamic burn-in for 8-160 hours to eliminate infant mortality and improve reliability. This process is executed using custom modified 32Bit AEHR test commercial burn-in equipment which provide monitored fixed pattern or pseudorandom burn-in with power supply and resistor device pin isolation.

LSI dynamic burn-in is verified in all cases by the design engineer for proper functioning. LSI Chip sets are also individually burned-in with dynamic equivalency to assure high performance bundled reliability.

The warranty on the program will optionally provide certificate of compliance to standard or custom designed burn-in programs and guarantee <.05%/Khrs failure rate.

CAUTION

Using outside burn-in methods not certified as acceptable by Western Digital may result in voided warranty, due to mishandling, junction temperature stress, or electrical damage. Further, since most burn-in houses do not support testing, catastrophic system condition can result in substantial damage before a problem is identified.

One consistent problem experienced with outside LSI burn-in houses can cause reliability problems; namely, parallelling totem pole MOS outputs, where the output states are not predictable, can cause a single (or a few) device(s) to sink all the current from the other devices on the burn-in tray — electromigration or current zaps are both possible.

Western Digital burn-in diagrams, dated after 1/1/82, must be used exactly as shown and will be provided upon request.

SEE YOUR LOCAL REPRESENTATIVE FOR COSTS AND ORDERING INFORMATION ON THIS NEW PROGRAM.

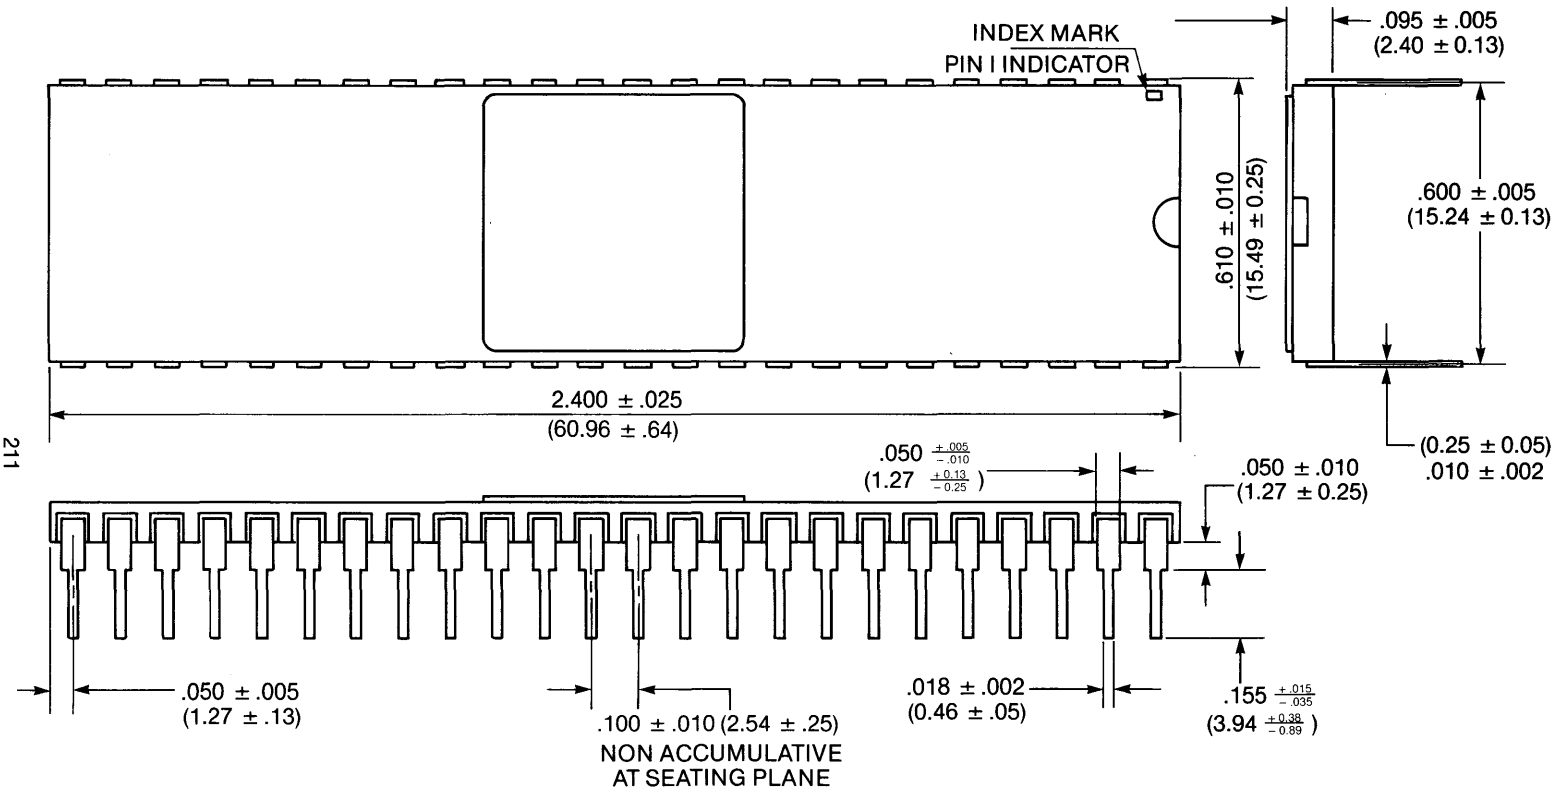

48 LEAD PACKAGE
CERAMIC SIDE BRAZE

NOTE:
( ) DENOTES METRIC VALUE IN MM

Component Products
Terms and Conditions

1. **ACCEPTANCE:** Unless otherwise provided, it is agreed that sales are made on the terms, conditions and warranties contained herein and that to the extent of any conflict, the same take precedence over any terms or conditions which may appear on Buyer's order form. Seller shall not be bound by Buyer's terms and conditions unless expressly agreed to in writing. In the absence of written acceptance of these terms, acceptance of or payment for any of the articles covered hereby shall constitute an acceptance of these terms and conditions.

2. **F.O.B. POINT:** All sales are made F.O.B. point of shipment. Seller's title passes to Buyer and Seller's liability as to delivery ceases upon making delivery of articles purchased hereunder to carrier at shipping point in good condition; the carrier acting as Buyer's agent. All claims for damages must be filed with the carrier. Unless specific instructions from Buyer specify which method of shipment is to be used, the Seller will exercise his own discretion.

3. **DELIVERY:** Shipping dates are approximate only. Seller shall not be liable for any loss or expense (consequential or otherwise) incurred by Buyer if Seller fails to meet the specified delivery schedule because of unavoidable production or other delays. Seller may deliver the articles in installments, Seller shall not be liable for any delay in delivery or for non-delivery, in whole or in part, caused by the occurrence of any contingency beyond the control either of Seller or Seller's suppliers, including, by way of illustration but not limitation, war (whether an actual declaration thereof is made or not), sabotage, insurrection, riot or other act of civil disobedience, act of a public enemy, failure or delay in transportation, act of any government or any agency or subdivision thereof, judicial action, labor dispute, accident, fire, explosion, flood, storm or other act of God, shortage of labor, fuel, raw material or machinery or technical failure where Seller has exercised ordinary care in the prevention thereof. If any contingency occurs, Seller may allocate production and deliveries among Seller's customers.

4. **TERMS AND METHODS OF PAYMENT:** Where seller has extended credit to Buyer, terms of payment shall be net thirty (30) days from date of invoice. The amount of credit or terms of payment may be changed or credit withdrawn by Seller at any time. If the articles are delivered in installments, Buyer shall pay for each installment in accordance with the terms hereof. Payment shall be made for the articles without regard to whether Buyer has made or may make any inspection of the articles. If shipments are delayed by Buyer, payments are due from the date when Seller is prepared to make shipments. Articles held for Buyer are at Buyer's sole risk and expense.

5. **TAXES:** All prices are exclusive of all federal, state and local excise, sales, use, and similar taxes. Such taxes; when applicable to this sale or to the articles sold, will appear as separate additional items on the invoice unless Seller receives a properly executed exemption certificate from Buyer prior to shipment.

6. **PATENTS:** The Buyer shall hold the Seller harmless against any expense or loss resulting from infringement of patents or trademarks arising from compliance with Buyer's designs or specifications or instructions. The sale of products by the Seller does not convey any license, by implication, estoppel, or otherwise, under patent claims covering combinations of said products with other devices or elements. Except as otherwise provided in the preceding paragraph, the Seller shall defend any suit or proceeding brought against the Buyer so far as based on a claim that any product, or any part thereof, furnished under this contract constitutes an infringement of any patent of the United States, if notified promptly in writing and given authority, information and assistance (at the Seller's expense) for the defense of same, and the Seller shall pay all damages and costs awarded therein against the Buyer. In case said product, or any part thereof, is in such suit held to constitute infringement and the use of said product or part is enjoined, the Seller, shall at its own expense, either procure for the Buyer the right to continue using said product or part, or replace same with non-infringing product, or modify it so it becomes non-infringing, or remove said product and refund the purchase price and the transportation and installation costs thereof. The foregoing states the entire liability of the Seller for patent infringement by the said products of any part thereof.

7. **ASSIGNMENT:** The Buyer shall not assign his order or any interest therein or any rights thereunder without the prior written consent of Seller.

8. **WARRANTY:** Seller warrants articles of its manufacture against defective materials or workmanship for a period of one year from date on which Seller delivers said articles. The liability of Seller under this warranty is limited at Seller's option, solely to repair, replacement with equivalent articles, or an appropriate credit adjustment not to exceed the original sales price of articles returned to the Seller provided that (a) Seller is promptly notified in writing by Buyer upon discovery of defects, (b) the defective article is returned to Seller, transportation charges prepaid by Buyer, and (c) Seller's examination of such article disclosed to its satisfaction that defects were not caused by negligence, misuse, improper installation, accident, or unauthorized repair or alteration by the Buyer. In the case of equipment articles, this warranty does not include mechanical parts failing from normal usage nor does it cover limited life electrical components which deteriorate with age. In the case of accessories, not manufactured by Seller, but which are furnished with the Seller's equipment, Seller's liability is limited to whatever warranty is extended by the manufacturers thereof and transferable to the Buyer. This Warranty is expressed in lieu of all other Warranties, expressed or implied, including the implied Warranty of fitness for a particlar purpose, and of all other obligations or liabilities on the Seller's part, and it neither assumes nor authorizes any other person to assume for the Seller any other liabilities. This Warranty should not be confused with or construed to imply free preventative or remedial maintenance, calibration or other service required for normal operation of the equipment articles. These Warranty provisions do not extend the original Warranty period of any article which has either been repaired or replaced by Seller. In no event will Seller be liable for any incidental or consequential damages.

9. **TERMINATION:** Buyer may terminate this contract in whole or from time to time in part upon 60 days written notice to Seller. In such event Buyer shall be liable for termination charges which shall include a price adjustment based on the quantity of articles actually delivered, and all costs, direct and indirect, incurred and committed for this contract together with a reasonable allowance for pro-rated expenses and profits. Any termination or back off in scheduling will not be allowed on shipments scheduled for the month in which the request is made and for the month following.

10. **GOVERNMENT CONTRACTS:** If the articles to be furnished under this contract are to be used in the performance of a Government contract or subcontract and a Government contract number shall appear on Buyer's purchase order, those clauses of the applicable Government procurement regulation which are mandatorily required by Federal Statute to be included in Government subcontracts shall be incorporated herein by reference.

11. **ORIGIN OF ARTICLES:** Seller engages in off-shore production, assembly and/or processing and makes no warranty or representation, expressed or implied, that the articles delivered hereunder are United States articles or of U.S. origin for the purpose of any statute, law, rule, regulation or case thereunder. If Buyer ships the articles hereunder out of the U.S. for assembly, then at Buyer's request in writing, Seller shall provide information applicable to identification of any articles not of U.S. origin.

Corita Kent, the cover artist, is an American whose work presents an optimistic, yet philosophical view of the world we live in. A former Catholic nun and teacher, Corita now devotes her life and energies to her artwork and the "human needs she feels transcend national and religious barriers." A true "citizen of the world," Corita's philosophy positions her "on the positive side of hope." Her depiction of the Western Digital mission . . . "Making the leading edge work for you" . . . dramatizes the spectrum of solutions we provide our customers.

Printed in U.S.A.

WESTERN DIGITAL
CORPORATION

2445 McCabe Way, Irvine, CA 92714
(714) 557-3550 • TWX 910-595-1139