ARCNET

June 22, 1982

John A. Murphy

PRELIMINARY

TABLE OF CONTENTS

CHAPTER 1. ARCNET PROTOCOL

1.1 LINE PROTOCOL

The line idles in a spacing (logic Ø) condition. A transmission starts with an ALERT BURST consisting of six unit intervals of mark (logic 1). Eight bit characters are then sent with each character preceeded by two unit intervals of mark and one unit interval of space. Five types of transmission are sent:

INVITATIONS TO TRANSMIT
: An ALERT BURST followed by three characters; an EOT (End Of Transmission) and two (repeated) DID (Destination IDentification) characters. Used to pass control of the line from one adaptor to another.

FREE BUFFER ENQUIRIES
: An ALERT BURST followed by three characters; an ENQ (ENQuiry) and two (repeated) DID (Destination IDentification) characters. Used to ask an adaptor if it is able to accept a packet.

PACKETS
: An ALERT BURST followed by from 8 to 26Ø characters; a SOH (Start Of Header), a SID (Source IDentification), two (repeated) DID's (Destination IDentification), a COUNT, from 1 to 253 data characters, and two CRC (Cyclic Redundancy Check) characters. Used to move data between adaptors.

ACKNOWLEDGEMENTS
: An ALERT BURST followed by one character; an ACK (ACKnowledgement). Used to acknowledge PACKETS and as an affirmative response to FREE BUFFER ENQUIRES.

NEGATIVE ACKNOWLEDGEMENTS
: An ALERT BURST followed by one character; a NAK (Negative AcKnowledgement). Used as a negative response to FREE BUFFER ENQUIREs.

The receiver validates all incoming transmissions by checking for:

At least one mark and exactly one space preceeding each character;

An EOT, ENQ, SOH, ACK, or NAK following the ALERT BURST;

Proper CRC (packets only);

Proper number of characters (3, 8 to 260, or 1);

At least nine spaces following the last character.

1.2 LINK CONTROL

Each RIM in a system has a unique ID (IDentification) from 1 to 0377 selected by jumpers. (ID 0 may not be assigned to any RIM since destination 0 is used to indicate a BROADCAST to all RIMs!). System operation is based on an INVITATION TO TRANSMIT being passed around the system with each adaptor passing it to NID (Next ID), the RIM with the next higher ID in the system. When a RIM receives an INVITATION TO TRANSMIT containing its ID it assumes control of the line. Which RIMs are in the system is determined during SYSTEM RECONFIGURATION.

When a RIM is first turned on, or has not received an INVITATION TO TRANSMIT for approximately 840 ms., it causes a SYSTEM RECONFIGURATION by sending a RECONFIGURE BURST consisting of eight marks and one space repeated 765 times. The purpose of this burst is to terminate all activity on the line. It is longer than any other type of transmission and will therefore interfere with the next INVITATION TO TRANSMIT and keep any RIM from seeing it and assuming control of the line. It also provides line activity so that the RIM sending the INVITATION TO TRANSMIT releases control of the line. Thus the RIM that had control releases it and no other RIM picks it up. When any RIM sees idle line for 78.2 us. it knows the system is being reconfigured and initializes its NID to its own ID. It then starts a time-out equal to 146 us. times 255 minus its own ID. If this time-out expires with no line activity the RIM starts sending INVITATIONs TO TRANSMIT. (Note that this time-out will expire only in the RIM with the highest ID in the system.)

After sending an INVITATION TO TRANSMIT the RIM waits for activity on the line (the RIM receiving the invitation sending a FREE BUFFER ENQUIRY, PACKET, or INVITATION TO TRANSMIT or any RIM sending a RECONFIGURATION BURST). If there is no activity for 74.7 us. the RIM increments NID and tries again. If it hears any activity before the time-out expires it releases control of the

line. During SYSTEM RECONFIGURATION INVITATIONs TO TRANSMIT will be sent to all 256 possible ID's. Each RIM, however, will have saved NID, the ID of the RIM that assumed control from it. From then until the next SYSTEM RECONFIGURATION (which will occur only when a new RIM is powered up or when a RIM gets dropped from the system due to line errors causing it to miss an INVITATION TO TRANSMIT), control is passed directly from RIM to RIM with no wasted INVITATIONs TO TRANSMIT sent to ID's not in the system.

The time required to do a SYSTEM RECONFIGURATION depends on the number of RIMs in the system and the propagation delays between them, but will be in the range of 24 to 61 ms.

1.3 DATA EXCHANGE

When a RIM receives an INVITATION TO TRANSMIT it checks to see if it has a packet to send, i.e. if TA (Transmitter Available) is false. If not it sends an INVITATION TO TRANSMIT to NID. Otherwise it tests byte 001 in the transmit buffer, the DID (Destination IDentifier). If this byte is 000 the packet is a BROADCAST and the RIM sends the packet. Otherwise it sends a FREE BUFFER ENQUIRY to the DID RIM and waits up to 75.1 us. for a response. If the response to the FREE BUFFER ENQUIRY is an ACK it sends the packet. If after sending the COUNT it finds that the COUNT has been set to 000 it truncates the rest of the packet (insuring that no RIM will receive it), sets TA, and sends an INVITATION TO TRANSMIT to NID. If the response to the FREE BUFFER ENQUIRY is a NAK it sends an INVITATION TO TRANSMIT to NID and will send another FREE BUFFER ENQUIRY the next time it receives an INVITATION TO TRANSMIT. If the RIM times-out waiting for a response to the FREE BUFFER ENQUIRY it sets TA and sends an INVITATION TO TRANSMIT to NID.

After sending a packet the RIM waits up to 75.1 us. for a response. If it receives an ACK it sets TMA and then TA and sends an INVITATION TO TRANSMIT to NID. If it times out waiting for an ACK (packets are never NAK'ed) it just sets TA and sends an INVITATION TO TRANSMIT to NID.

When a RIM receives a FREE BUFFER ENQUIRY it tests RI. If RI is true it sends a NAK. Otherwise it sends an ACK.

When a RIM receives an SOH (indicating the start of a PACKET) it writes the SID into the receive buffer and then checks the first DID. If this byte is 000 (indicating a BROADCAST) the RIM tests byte 001 in the receive buffer for a 000 (reception of BROADCASTs enabled). If reception of BROADCASTs is enabled, or if

the first DID is the RIM's own ID, the RIM writes the second DID, and the rest of the PACKET into the receive buffer. Otherwise it ignores the rest of the PACKET. If after being written into the receive buffer the PACKET fails either the CRC or length validation phases, the RIM ignores it. Otherwise it tests byte 001 in the receive buffer, the DID. If this byte is 000 the packet is a BROADCAST and the RIM simply sets RI. If this byte is the RIM's own ID the RIM sends an ACK before setting RI. If the DID is neither 000 nor the RIMs own ID the RIM ignores the packet.

1.4 PERFORMANCE FACTORS

The most interesting aspect of performance in a local network is usually the amount of time a node may have to wait before being able to send a message. In a token passing scheme this waiting time is bounded by the time it takes the token to make the rounds of each and every node. There are several possible sequences of events that can occur when a node receives the token; two of these make up the vast majority of cases, the others having to do with error conditions, messages sent to non-existent nodes, etc. The two cases of interest are a simple token pass and a message followed by a token pass. A simple token pass takes between 28 and 59 microseconds, depending on the propagation delay between the two nodes. (In even the largest real world systems, the propagation delay between nodes, averaged over a complete token trip, is nearly zero, resulting in an average time per token pass of something very close to 28 microseconds.) A message followed by a token pass takes between 141 and 296 microseconds, depending on propagation delays, plus 4.4 microsecond per byte of data. Thus the time required for the token to make a complete trip around the network is approximately 28 us. per node plus 113 us. per message plus 4.4 us. per byte.

In a system, for example, consisting of 10 nodes and sending messages no longer than 100 bytes, a complete token trip will be at least 280 us. (when no messages are sent) and no longer than 5810 us. (when all 10 nodes send a 100 byte message). If only a single node is sending messages, it can send one every 813 us.; a rate of 1230 messages, or 123000 bytes, per second. If all 10 nodes are sending messages, each one can send one every 5810 us.; a rate of 172 messages per second per node.

At the other end of the size spectrum, the original ARCNET system in the Datapoint Research and Development Department now has anywhere from 150 to over 200 nodes active at any given time. It supports two totally independent operating systems and a wide variety of uses including program loading, word processing, print

spooling, program development, electronic mail, etc. The traffic load rarely falls below 400 messages per second, yet less than 2% of the nodes send a message on the average token trip. The time required for a token trip, therefore, stays very close to the no traffic value; peaks of three times the no traffic value are extremely rare.

CHAPTER 2. ARCNET HARDWARE

Processors interface to ARCNET via special purpose communications adaptors called RIMs (Resource Interface Modules). Interaction between the processor and the RIM is on a message by message, as opposed to a byte by byte, basis. The RIM contains four 256 byte message buffers and the arbitration logic required to share them between the processor, the RIM transmitter, and the RIM receiver. To send a message, the processor writes the message into a RIM buffer and issues a transmit command; the RIM sets a status bit (and can cause an interrupt) when the message has been sent. To receive a message, the processor assigns a RIM buffer to the RIM receiver; the RIM sets a status bit (and can cause an interrupt) when a message has been received and is available in the buffer.

The RIM IC is an NMOS LSI implementation of the major portions of the RIM.

2.1 ANATOMY OF A RIM

A complete RIM consists of four basic sections:

* CONTROLLER - the RIM IC interfaced to 1k bytes of RAM buffer
* CLOCK - the clock oscillator and drivers
* LINK INTERFACE - the circuitry required to interface the CONTROLLER to the serial data link (baseband coax, broadband coax, fiber optic cable, etc.)
* PROCESSOR INTERFACE - the address decoders, bus drivers, etc. required to interface the CONTROLLER to the system bus of the host processor.

2.1.1 CONTROLLER

The CONTROLLER section consists of the RIM IC, an eight bit node ID switch, a shift register, a 1k x 8 static RAM, and a gated latch. Control signals from the RIM IC periodically load the switch into the shift register and then shift (most significant bit first) the node ID into the RIM IC. The 8 line Address-Data bus and 2 line PAGE bus provide the address and data path to the PROCESSOR INTERFACE. The Address-Data bus is multiplexed between address and data under control of the RIM IC.

2.1.1.1 BUFFER MEMORY AND CONTROL

The buffer memory is composed of 1K of static RAM. The memory is segmented into four pages, 256 bytes each. It is a dual port memory, accessed directly by the RIM IC, or indirectly by the processor under control of the RIM IC. The RIM IC's bus control logic provides the control strobes for both access modes.

During a transmit sequence, the RIM IC reads data from the transmit buffer, one of the four 256 byte pages of the buffer memory. The particular page used is selected by a two bit field in the transmit command written into the RIM IC command register. During a receive sequence, the RIM IC writes data in the receive buffer, also one of the four 256 byte pages of the buffer memory. The particular page used is selected by a two bit field in the receive command written into the RIM IC command register.

When the processor is accessing the RAM, the RAM address and the data are exchanged directly with the processor interface. The RIM IC only generates the control signals.

A suitable memory is the MOSTEK MK48Ø1AN-4 RAM. The RAM's chip select pin (CS*) is grounded and RIM IC address output A8 is not used. An external latch latches the lower 8 bits of the address.

2.1.1.2 BUS ARBITRATION AND CONTROL LOGIC

Since the buffer is a two port memory, arbitration logic is required to govern access to the memory. The Address-Data bus, AD7-ADØ, provides a path for three separate types of data transfer:

1. RIM IC access to buffer memory
2. Processor access to buffer memory

3. Processor access to RIM IC control registers.

The RIM IC requests occur when the transmitter or receiver needs to read or write the buffer. The processor request can result from either a buffer memory access request or from a control/status I/O command.

The two types of processor data transfers cannot occur simultaneously, so arbitration between these is unnecessary. RIM IC memory requests are asynchronous to both types of processor requests, however, so arbitration is required.

The bus arbitration logic allocates bus cycles for the three separate functions. In the event of simultaneous requests from the processor and the RIM IC, the processor has priority. At the end of a processor cycle, the RIM IC will be granted use of the bus if it is waiting. All RIM IC memory cycles must be completed within 2.2 microseconds.

2.1.2 CLOCK

The CLOCK section provides two 5 MHz clock signals. One of these is free running and is applied to the CLK inputs of the RIM IC and the ID shift register. The other clock is synchronized to the received data under control of the DSYNC* output of the RIM IC and applied to the CA input of RIM IC. Sychronization is accomplished by sampling DSYNC* on the rising edge of CA and stopping the clock when DSYNC* is found to be low. This will occur during the ALERT BURST and the stop element of each received chacacter. CA is left high while the clock is stopped. The clock is restarted when the received signal from the LINK INTERFACE, the RX input to the RIM IC, goes low as the start element of a character is received. When the clock is restarted CA should go low 1ØØ ns. (nominally) after the leading edge of the start element, so that the rising edge of CA occurs in the center of the start element.

2.1.3 LINK INTERFACE

The LINK INTERFACE section of the RIM interfaces the serial data link to the CONTROLLER section via the TX* and RX pins of the RIM IC. The exact implementation of this section is dependent on the type of data link (twisted pair, coax, etc.) and signaling scheme (NRZ, Manchester, etc.) used. The version described here is the STANDARD ARCNET INTERFACE; a baseband system using dipulse

signaling on RG-62 coax and allowing runs of up to 2000 feet.

2.1.3.1 TRANSMITTER

The serial output of the transmitter in the RIM IC is TX*. This pin remains high for spaces and goes low for 200 ns for marks.

The STANDARD ARCNET INTERFACE samples TX* on the trailing edge of CA and generates a 200 nanosecond wide dipulse for each mark.

2.1.3.2 RECEIVER

The serial input of the receiver in the RIM IC is RX. The data format at this pin is NRZ with spaces low and mark high.

In the STANDARD ARCNET INTERFACE the incoming signal on the coax is passed through a matched filter and line receiver. The output of the line receiver is then converted to NRZ. A low to high transition at the filter output sets a flip-flop. The flip-flop is sampled at 50 ns intervals, and when a one is detected, the receiver resets the flip-flop and generates a 400 ns pulse on RX.

2.1.4 PROCESSOR INTERFACE

The processor sees the RIM as two seperate sets of contiguous eight bit registers, or memory locations: the two location set of control/status in the RIM IC, and the 1024 location set in the buffer memory. The PROCESSOR INTERFACE section of the RIM provides the interface necessary to allow the host processor to access these two areas. Depending on the requirements of the particular processor and the whims of the hardware designer and/or system programer, these areas may appear as I/O registers, locations in memory space, locations in I/O space, or some combination of these.

The interface is asynchronous, with timing controlled by AS and DWR*. Processor access cycle requests begin on the trailing edge of AS if either IOREQ* or MREQ* is asserted. The RIM IC responds by asserting its WAIT and REQ outputs. (WAIT will return to its normal state near the end of the cycle; REQ is the output of a tranparent latch gated by AS, and will therefore remain true until AS goes high with neither IOREQ* nor MREQ* asserted.) When

the Address-Data bus arbitration logic grants the processor a cycle, the RIM IC asserts ADIE* and AIE* to gate address information from the processor onto the Address-Data bus. (The signal ADIE* is used to gate either address or data onto the Address-Data bus, while AIE* is used to gate the higher order address bits.) On read cycles, determined by R/W* being high at the trailing edge of AS, the RIM IC latches the address information, returns ADIE* to its normal state, and places read data on the Address-Data bus. On write cycles, determined by R/W* being low at the trailing edge of AS, the RIM IC latches the address information, and asserts ILE* to gate processor data onto the Address-Data bus. The input DWR* may be used, if neccessary, to cause the RIM to wait until the processor data is ready. Some function of the REQ, WAIT, and AIE* outputs, depending on the particular host processor and system bus structure, can be used to indicate the completion of the cycle.

The arbitration and bus control logic is designed so that the 2.2 microsecond limit will always be met, providing the processor does not introduce excessive delays on the DWR* (delayed write) control signal. To assure valid RIM operation the delay between the fall of AS and the assertion of DWR* must not exceed 950 ns. To assure the maximum I/O transfer rate (minimum number of processor wait cycles) the delay between the fall of AS and the assertion of DWR* should not exceed 350 ns. If it can be determined that WRITE data will always be valid in time (see timing specifications in Appendix A), then DWR* should be grounded.

2.1.4.1 CONTROL/STATUS

The low order address bit determines the function; the higher order address bits and control signals are decoded to produce a chip select signal (IOREQ*) which is latched internal to the chip by AS. The interface functions include:

- WRITE RIM COMMAND - issues a general purpose RIM control command. This command provides functions such as enable transmitter, disable receiver, clear selected status bits, etc.

- READ STATUS - reads the 8-bit RIM status latch.

- WRITE STATUS MASK - loads a mask register which enables interrupts from selected status bits.

2.1.4.2 BUFFER ACCESS

The low order ten address bits address the buffer; the higher order address bits and control signals are decoded to identify a processor access request command to the RAM buffer (MREQ*) which is latched internal to the chip by AS.

The buffer address is latched in the input latches. There are two input latches, HIGH ADDR LATCH for P1-PØ (system bus address lines A9-A8) and ADDR/DATA IN for AD7-ADØ. Separate enable signals are employed to allow the HIGH latch to supply A9-A8 continuously while the ADDR/DATA IN latch is disabled (e.g., during processor read operations).

During processor data transfers the RIM IC acts as a memory control unit, interpreting MREQ*, R/W*, and DWR*, and generating control strobes for the buffer memory and the processor interface. The RIM IC will assert WAIT for as long as necessary. This halts the processor, allowing the RIM IC to coordinate the processor with its own memory cycle timing. The signal WAIT can also be used to latch the data on processor read cycles.

2.2 POWER UP AND INITIALIZATION

The RIM IC has the following power-up requirements:

1. When power is first applied, POR* must be active for at least 1ØØ ms to allow the on-chip bias generator to charge the substrate. The bias generator is independent of CA and CLK.

2. CLK must run for at least ten cycles before POR* is negated. This permits initialization of the bus control logic.

3. While POR is asserted, CA may be either running or held high. If CA is running, POR* may be negated asynchronously with respect to CA. If CA is held high, POR* may be released before CA starts to run.

While POR* is asserted, DSYNC* = 1. The RIM's controller will start operation approximately four CA cycles after POR* is negated.

2.3 RIM IC PINOUT AND PIN DESCRIPTION

2.3.1 PINOUT

PIN CONFIGURATION

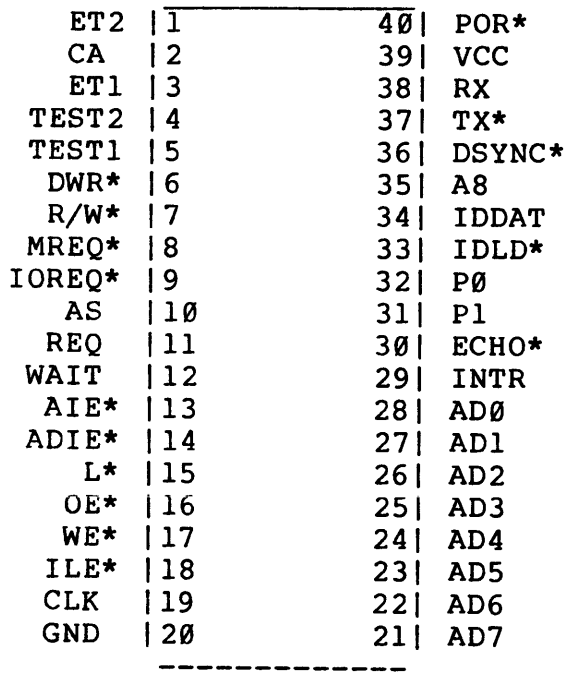

2.3.2 PIN DESCRIPTION

SYMBOL	FUNCTION
P1-P0 (Output, tri-state)	PAGE bus: The most significant two bits of the buffer memory address, tri-stated except during RIM access cycles to the RAM buffer. These pins are system bus address bits 9-8 in a 1k buffer system, and bits 10-9 in a 2k buffer system.

A8 (Output, tri-state)	ADDRESS: This pin is used as an eleventh address bit in 2k buffer systems. Tri-stated except during RIM access cycles to the RAM buffer. This pin is unused in a 1k buffer system, and system address bit 8 in a 2k buffer system.
AD7-ADØ (Input/Output, tri-state)	Address-Data bus: The lower 8 bits of the buffer memory address and the 8 bit data path in and out of the RIM chip. ADØ is also used for I/O command decoding of processor control/status commands to the RIM.
IOREQ* (Input)	I/O REQUEST: A low level on IOREQ* indicates that the processor is requesting use of the RIM bus to transfer control/status information. IOREQ* is sampled by the RIM on the falling edge of AS.
MREQ* (Input)	MEMORY REQUEST: A low level on MREQ* indicates that the processor is requesting use of the RIM bus to transfer data to or from the RAM buffer memory. MREQ* is sampled by the RIM on the falling edge of AS.
R/W* (Input)	READ/WRITE: A high level on R/W* indicates that the processor access cycle to the RIM or to the RAM will be a read cycle. R/W* is sampled by the RIM on the falling edge of AS.
AS (Input)	ADDRESS STROBE: Used by the RIM to sample the status of IOREQ*, MREQ*, and R/W*. RIM bus arbitration is initiated on the falling edge of AS if either IOREQ* or MREQ* is asserted.
REQ (Output)	REQUEST: Acknowledgement from the RIM that a processor I/O cycle or memory cycle request has been sampled.

WAIT (Output)	WAIT: The RIM asserts WAIT at the start of processor access cycles to indicate that it is not ready to transfer data. WAIT is negated when the RIM is ready for the processor to complete the cycle.
DWR* (Input)	DELAYED WRITE: During processor write cycles the processor uses DWR* to signal to the RIM that there is valid data on the data bus. On processor write cycles, if DWR* does not meet setup (logic Ø) to the start of state 4, then the RIM bus control logic enters an internal wait state until DWR* is asserted. DWR* may be asserted asynchronously. DWR* is not used on read cycles. If the processor will always meet setup for WRITE data, DWR* should simply be grounded.
INTR (Output)	INTERRUPT: INTR is asserted by the RIM to indicate that an enabled interrupt condition has occured. INTR is negated when the interrupting status condition is reset, or when the corresponding mask bit is set to zero.
ILE* (Output)	INTERFACE LATCH ENABLE: An active low signal which gates the processor output data onto the Address-Data bus during processor write cycles.
ADIE* (Output)	ADDRESS DATA INPUT ENABLE: An active low signal which enables the processor to drive the Address-Data bus.
AIE* (Output)	ADDRESS INPUT ENABLE: An acitve low signal which enables the processor to drive A8 and P1-PØ.
L* (Output)	LATCH: An active low signal which latches the lower eight bits of the Address-Data bus into an eight bit address latch external to the RAM.
WE* (Output)	WRITE ENABLE: The trailing edge of WRITE ENABLE signals the RAM to execute a write operation.

OE* (Output)	OUTPUT ENABLE: An active low signal which enables the RAM output data onto the Address-Data bus.
IDLD* (Output)	ID LOAD: An active low signal which synchronously loads the setting of the ID switches into the external ID shift register.
IDDAT (Input)	ID DATA IN: An input which accepts the serialized ID switch setting. The ID is shifted into the RIM, high true, most significant bit first.
ET2-ET1 (Input)	EXTENDED TIMEOUT FUNCTION: See notes 1 and 4.
TX* (Output)	TRANSMIT DATA: Serial transmit data output to the LINK transmitter.
RX (Input)	RECEIVE: Serial receive data input from the LINK receiver.
TEST1-TEST2 (Input)	TEST PINS 1 and 2: See notes 2 and 3.
ECHO* (Input)	ECHO DIAGNOSTIC ENABLE: See notes 2 and 4.
CLK (Input)	CLOCK: A continuous 5MHz clock used for timing of the RIM bus cycles and for bus arbitration, serial ID input, and the RECON TIMER. The chip requires a high level drive on the CLK input.
CA (Input)	CA: A 5 MHz clock used to control the operation of the RIM sequencer. CA is periodically halted in the high state. The chip requires a high level drive on the CA input.
DSYNC* (Output)	DELAYED SYNC: The RIM asserts DSYNC* to cause the external clock generator logic to halt the CA clock.

POR* (Input)	POWER ON RESET: Sets the program counter to zero and initializes various internal control flags and status bits. Sets the POR status bit, causing INTR to be asserted.
VCC (Input)	+5 Volt Supply.
GND (Input)	Ground.

NOTES:

1	These pins are used to configure the RIM IC for special applications.
2	These pins are used only for testing.
3	These pins should be grounded for normal operation.
4	These pins should be tied high for normal operation.

CHAPTER 3. TRANSMISSION MEDIA AND NETWORK TOPOLOGY

The ARCNET protocol imposes very few restrictions on the transmission medium. Basically, the receiver at any node must be able to hear the transmitter at any other node, any two nodes must be connected by a single path, and the path between any two nodes must have a propagation delay of 31 microseconds or less.

The STANDARD ARCNET INTERFACE is a baseband system using dipulse signalling on RG62 coax. The usual disadvantage of a baseband system is the difficulty of tapping the line for distribution. This problem is circumvented through the use of "hubs." Each node connects, through a length of coax, to a port on a hub; the coax is properly terminated at each end and no taps of any kind are used. In a sense, the hub may be thought of as an amplifier and a number of ideal taps all mounted in the same box! These virtual taps are ideal in the sense that they have no insertion loss and no tap loss, and yet provide total suppression of reflections, even from unterminated lines. Each of the ports of a hub may be connected to an ARCNET node, to another hub, to an unterminated length of coax, or to nothing at all.

Each port of a hub consists of a line driver and receiver identical to those in the RIMs. The outputs of the receivers are OR'ed together. In its idle state the hub has all of the receivers enabled. As soon as one of the receivers receives anything, the hub latches into a state in which the output of that receiver is fed to the inputs of the transmitters for all of the other ports; all of the other receivers, meanwhile, are disabled. The hub remains in this state until the transmission it is repeating is finished, and then drops back into its idle state.

The determination of when a transmission is finished is based on time. There are never more than nine consecutive spacing elements in a transmission (the start element and eight zeros); a dipulse is received at least once every ten unit intervals, or 4.Ø microseconds. The RIMs have a turnaround time of something greater than 12 microseconds; there is at least 12 microseconds between the end of the last data element of one transmission and the start of the alert burst of the next. Were it not for the reflection problem, the hub could drop back into its idle state when the receiver had not heard anything for some period of time between 4.Ø and 12.Ø microseconds.

In order to provide protection against reflections from

unterminated lines, the hub should not fall into its idle state until the reflections cease. The STANDARD ARCNET INTERFACE limits individual runs of coax to 2000 feet of RG62 coax, which has a velocity of propagation of 84%. Thus the reflection from an unterminated, or shorted, cable will return in less than 4.9 microseconds. Changing the 4.0 microsecond limit mentioned above to 4.9 microseconds will therefore make the hub ignore reflections.

In summary, then, an N port hub is a device with N+1 stable states: the idle state, in which all N receivers are enabled, and N active states, one corresponding to each of the N receivers driving N-1 transmitters. The transition from the idle state to an active state is caused by the reception of a dipulse by one of the receivers; the transition back to the idle state is caused by 8.5 (+/- 40%) microseconds of silence.

CHAPTER 4. SOFTWARE INTERFACE TO THE RIM

4.1 PROCESSOR INTERFACE - CONTROL/STATUS

The processor exchanges control and status information with the RIM by reading and writing two registers in the RIM IC. The high order bits of the address are decoded outside the chip to produce a chip select signal. The low order bit specifies the address of a specific I/O register. The register functions are:

ADØ	READ FUNCTION	WRITE FUNCTION
Ø	Status	Mask
1	(reserved)	RIM Command

A brief explanation of the I/O functions is given below.

4.1.1 READ STATUS

Execution of a read from RIM address Ø latches the current RIM status bits into the status latch and places the latched status on the data bus. The status register format is as follows:

BIT	STATUS	INDICATION
Ø	TA	Transmitter Available. Indictes that RIM transmitter is available for use; any previous ENABLE TRANSMIT command has completed.
1	TMA	Transmitted Message Acknowledged. Indicates that the message sent by the last ENABLE TRANSMIT command was acknowledged by the receiving RIM.
2	RECON	Reconfiguration Flag. Indicates that a SYSTEM RECONFIGURATION has occured since the last time this bit was reset by a CLEAR FLAGS command. (The occurance of a SYSTEM RECONFIGURATION is determined by the line being idle for at least 78.2 us.)

3	TEST	Test Flag. This bit is intended for test and diagnostic purposes. It will be zero under any normal operating conditions.
4	POR	Power on Reset. Indicates that the RIM IC has undergone a power-on-reset since the last time this bit was reset by a CLEAR FLAGS command.
5	ETS1	Extended Timeout Status 1. Indicates the state of the ET1 input pin of the RIM IC. This pin is used to configure the RIM IC for special applications and will be high (the status bit will be one) under any normal operating conditions.
6	ETS2	Extended Timeout Status 2. Indicates the state of the ET2 input pin of the RIM IC. This pin is used to configure the RIM IC for special applications and will be high (the status bit will be one) under any normal operating conditions.
7	RI	Receiver Inhibited. Indicates that the RIM receiver is inhibited; no messages will be received until an ENABLE RECEIVE command is issued. If an ENABLE RECEIVE command has been issued since the last power on reset, RI indicates that a message has been received and is available in the buffer.

4.1.2 WRITE MASK

The RIM is capable of generating an interrupt signal to the processor when selected status bits become true. A write to the MASK register specifies which status bits can generate an interrupt. A one bit in a position of the mask register enables an interrupt from the corresponding status bit. Status bits 1 (TMA), 5 (ETS1), and 6 (ETS2) will never cause interrupts --- they have no mask bit. Status bit 4 (POR) is non-maskable and will always cause an interrupt. The four maskable status bits are anded with their respective mask bits, and the results, along with the POR status bit, are OR'D to produce the processor interrupt request signal.

The interrupt request signal is negated when the interrupting status bit is reset, or when the corresponding bit in the mask register is set to zero.

4.1.3 WRITE RIM COMMAND

The processor issues commands to the RIM by writing a command byte to RIM I/O address 1, the RIM command register. The commands are:

COMMAND BYTE	FUNCTION
1 ØØ ØØØ ØØ1	DISABLE TRANSMITTER - Causes transmitter to cancel an uncompleted transmit command. The next time the RIM is polled TA will be set. Note that failure of TA to come on within about 36Ø ms. after issuing this command is an indication that there are no other RIMs connected to this one.
1 ØØ ØØØ Ø1Ø	DISABLE RECEIVER - Causes receiver to cancel any uncompleted receive command. The next time the RIM is polled RI will be set. Note that failure of RI to come on within about 36Ø ms. after issuing this command is an indication that there are no other RIMs connected to this one.
1 ØØ Ønn Ø11	ENABLE TRANSMIT from page nn - Clears TA and TMA status flags and commands the RIM to begin a tranmit sequence, using buffer page nn, the next time it is polled. TA is set upon completion of the transmit sequence. TMA will have been set by this time if the RIM has received an acknowledgement from the destination RIM. (Note that this is strictly a hardware level acknowledgement which which is sent by the receiving RIM before its host processor is even aware of the arrival of the message. Note also that the acknowledgement may get lost due to line errors, etc., so that TMA not being set is not a guarantee that the message was not received.) This command should not be issued unless TA is true!
1 bØ Ønn 1ØØ	ENABLE RECEIVE to page nn - Clears RI status flag and allows the RIM to receive messages into buffer page nn. If bit "b" is a one the RIM will accept broadcasts as well as messages addressed to it. If "b" is zero only messages addressed to this particular RIM will be accepted. RI is set upon successful reception of a message.

1 ØØ Ørp 11Ø	CLEAR FLAGS - Conditionally clears Power-On-Reset and RECON status bits. If p=1, clears Power-On-Reset status flag. If r=1, clears system RECON status flag.

4.2 PROGRAMMING

To transmit a message the processor selects a buffer and writes into it in the following format:

ADR	CONTENTS	FUNCTION
Ø	xxx	Unused (RIM sends local ID regardless of this value)
1	DID	Destination IDentifier
2	CP	Continuation Pointer (buffer address of first data byte)
...		
CP	data	First data byte
...		
255	data	Last data byte (always at end of buffer)

The processor then waits for TA (status bit Ø) to be true and gives the transmit command. At the completion of its transmit sequence the adaptor conditionally sets TMA (status bit 1) and then sets TA.

If the DID is non-zero (the PACKET is not a BROADCAST) the RIM will wait for a free buffer (RI false) at the destination (without locking out other users) before sending the packet. Note that if the host processor at the destination is not servicing its RIM the adaptor at the source will never find a free buffer and never set TA. There must, therfore, be a software time-out on TA. When this timer times out the processor should disable the TRANSMITTER to force the RIM to abandon the transmission. Note that if the disable TRANSMITTER command does not cause TA to return true within 36Ø ms. it is an indication that there are no other RIMs in the system.

To enable the receiver the processor selects a buffer to use, waits for RI (status bit 7) to be true, and gives the receive command . When a packet addressed to the local adaptor or sent as a BROADCAST (if reception of BROADCASTs was enabled) is completely and correctly received the adaptor sets RI. The processor selects the buffer used by the receiver and reads the buffer. The buffer contents are as follows:

ADR	CONTENTS	FUNCTION
Ø	SID	Source IDentifier
1	DID	Destination IDentifier (local ID or Ø)
2	CP	Continuation Pointer (buffer address of first data byte)
...		
CP	data	First data byte
...		
255	data	Last data byte (always at end of buffer)

The format of the buffers is shown below.

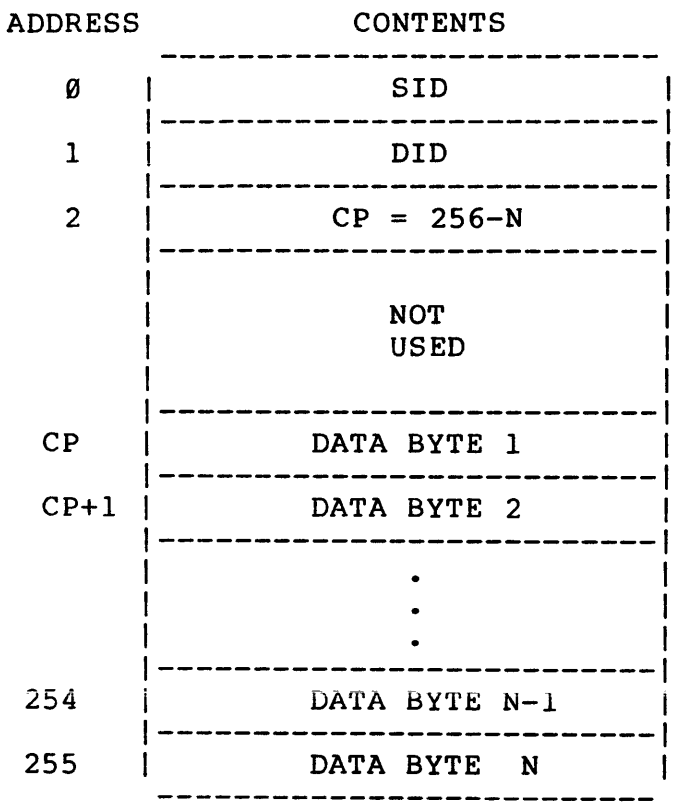

N - DATA LENGTH
SID - SOURCE ID (NOT USED FOR TRANSMIT BUFFERS)
DID - DESTINATION ID (Ø FOR BROADCASTS)

CHAPTER 5. DIAGNOSTIC FEATURES

The RIM IC incorporates several diagnostic features to simplify the task of verifying correct operation and/or identifying incorrect operation. These features are applicable to RIM IC at three levels - chip test, board test, and unit test.

5.1 CHIP TEST DIAGNOSTICS

The RIM IC input pins TEST1 and TEST2 are used, in combination with input pins ET1 and ET2, to select from among a number of internal diagnostic functions. These functions are designed to simplify testing the RIM IC in a dedicated test system and normally cannot be used once the chip is incorporated into a PC board. The tests require intricate synchronized manipulation of the four test control pins, the two clocks, POR*, etc. and are typically run on automated test equiptment by the chip manufacturer.

5.2 BOARD TEST DIAGNOSTICS

The board test diagnostic functions allow testing of the RIM IC at the PC board level. These functions typically require direct access to the RIM circuitry, but do not require that the chip be removed from the PC board.

5.2.1 ECHO TEST

The ECHO TEST is designed to execute a partial test on the RIM IC and its buffer without using the local processor. It is particularly useful in single board systems for identifying whether the RIM is functional when the condition of the processor is unknown.

Grounding the ECHO* pin on the RIM IC effectively "locks" both RI and TA false: the reveiver is always ready to accept a message, and the transmitter will send a message every time it gets the token. If the receiver and transmitter are both dealing with the same page of the buffer, as they will be following a power-on-reset, then the RIM will continuously re-transmit, or echo, any message it receives.

The ECHO TEST uses a known good RIM and a processor (or equivalent test fixture) to transmit a message to the RIM under test and then receive one of the echos. Since the echos will have the same DID as the original, and since both the RIM under test and the known good RIM must be able to receive messages containing this DID, either the test message must be a broadcast or both RIMs must have the same ID. The test sequence is as follows:

1. Through RIM commands, disable the transmitter and the receiver on the known good RIM.

2. Ground the ECHO* pin on the RIM under test. Momentarily assert POR* to set the transmit and receive page registers to zero, and to set the broadcast reception enabled flag.

3. On the known good RIM, load the transmit buffer with a packet for the RIM under test, and enable the transmitter. When TA becomes true, enable the receiver.

4. When RI becomes true, check the contents of the receive buffer for correct data.

Note that this test should never be performed in an operating system! If the test message is a broadcast, the continuous stream of (echoed) broadcasts from the RIM under test will undoubtedly be a source of consternation to other nodes in the network! And while the trick of using the same ID for the known good RIM and the RIM under test works fine in a universe containing only two RIMs, duplicate RIM IDs are definately not allowed in a normal system!

5.3 UNIT TEST DIAGNOSTIC FEATURES

The unit test features allow partial testing of the RIM through processor commands.

5.3.1 RIM ID

At POR the RIM IC microcode will execute two write cycles to the buffer. The data stored is

Address	Data
0000	0321 (octal)
0001	RIM ID

The constant Ø321 may be read by the processor to determine (with reasonable probability) that the test function is working. The processor may then read address ØØØ1 to determine the RIM ID. This is a fairly powerful test in that it verifies at least partial operation of:

* ID logic (on-chip and off-chip)

* Major portions of the on-chip logic

* RIM-Write-RAM and Processor-Read-RAM cycles.

5.3.2 TEST FLAG

The TEST FLAG is controlled via a special command to the RIM IC:

1 ØØ ØØt 111	LOAD TEST FLAG - Sets or clears the TEST FLAG depending on whether t is one or zero, respectively.

The TEST FLAG provides a tool for checking the processor-to-RIM interface. Since the TEST FLAG is also a status bit, the effect of the RIM command may be verified by reading the status register. The interrupt structure may also be checked by enabling and disabling interrupts with the WRITE MASK command, while setting and resetting the TEST FLAG.

Due to the sharing of some internal functions of the RIM IC, the TEST FLAG should not be set unless status bit TA is true.

CHAPTER 6. NON-STANDARD OPERATION OF THE RIM IC

The RIM IC is capable of being configured for two different types of "non-standard" operation. The maximum allowed propagation delay may be increased to provide for operation over large distances, and messages longer than 253 bytes may be sent. The use of these features results in a non-compatible system and should therefore be implemented only after careful consideration!

6.1 EXTENDED DISTANCE OPERATION

The Extended Timeout feature allows the RIM IC to operate over greater distances than standard RIM's. DC levels at the ET2 and ET1 inputs control the maximum distance over which the RIM can operate. The RIM's distance limitations are a function of internal timeouts which are part of the RIM protocol. The two types of internal timeouts are the Response timeout and the RECON timeout. The Response timeout determines how long the RIM waits for a response after sending a message, and thus also determines the maximum propagation delay allowed between any two nodes. The RECON timeout occurs when a RIM fails to receive an invitation to transmit from another RIM in the expected amount of time.

Altering the values of ET2 and ET1 cause the timeouts to change as follows:

ET2	ET1	MAX PROP DELAY (microseconds)	RECON TIMEOUT (seconds)	TIME TO RECONFIGURE (milliseconds)
1	1	31	Ø.84	24 to 61
1	Ø	131	1.68	77 to 223
Ø	1	271	1.68	149 to 439
Ø	Ø	549	1.68	291 to 6Ø3

Note that all RIM's on the network must use the same timeouts, and that the ARCNET standard is the 31 microsecond (ET1 = ET2 = VCC) setting.

6.2 EXTENDED LENGTH MESSAGE OPERATION

The extnded length message feature allows the RIM IC to handle messages longer than the standard 253 bytes. Use of this feature involves both hardware and software considerations.

6.2.1 HARDWARE CONSIDERATIONS

RIMs supporting the extended length message feature require a 2k buffer, organized as four 512 byte pages, instead of the standard 1k buffer. RIM IC outputs PØ and Pl select the buffer page, as in the standard configuration, but are connected to system address lines A9 and AlØ instead of A8 and A9. RIM IC output A8, which is unused in the standard configuration, prodives the eleventh bit of address for the buffer.

6.2.2 SOFTWARE CONSIDERATIONS

Use of the extended length message feature is controlled via a special command to the RIM IC:

1 ØØ ØØc 1Ø1 — DEFINE CONFIGURATION - The Long Packet Enable flag, an internal flag in the RIM IC, is set to the value of c. When this flag is set, a continuation pointer (contents of address 2 in a buffer page) value of zero indicates that the actual continuation pointer value is 256 plus the contents of address 3.

Messages of length 1 to 253 are sent in the usual manner, with the last byte of the message located at buffer address 255 and the contents of address 2 equal to the buffer address of the first byte of the message.

Messages of length 257 to 5Ø8 are loaded into the buffer with the last byte at address 511, the contents of address 2 set to zero, and the contents of address 3 equal to the buffer address of the first byte of the message.

Note that messages of length 254 to 256 must be padded out to a length of at least 257 bytes in order to be handled.

6.2.3 COMPATIBILITY CONSIDERATIONS

RIMs equipped and configured for extended length message operation can coexist in the same system with standard RIMs. The DEFINE CONFIGURATION command merely informs the RIM IC of the existence of the 2k buffer (as opposed to the standard 1k) and thus need only be issued at initialization time. Operation with standard length messages proceeds in the normal fashion.

If an extended length message is sent to a RIM that does not have its Long Packet Enable flag set, the receiver will ignore it. The transmitting RIM will set TA, but not TMA.

If an attempt is made to send an extended length message when the DEFINE CONFIGURATION command has not been used to set the Long Packet Enable flag, the packet will not be sent and TA will stay off until a DISABLE TRANSMITTER command is issued. The host software sees this exactly as it sees the case of a message to a RIM that never makes a receive buffer available.

The final compatibility consideration involves compatability with existing Datapoint ARCNET systems using RIMs implemented without the RIM IC. To insure compatibility with such systems, messages longer than 32Ø bytes should not be sent.

The format of the extended length buffers is shown below.

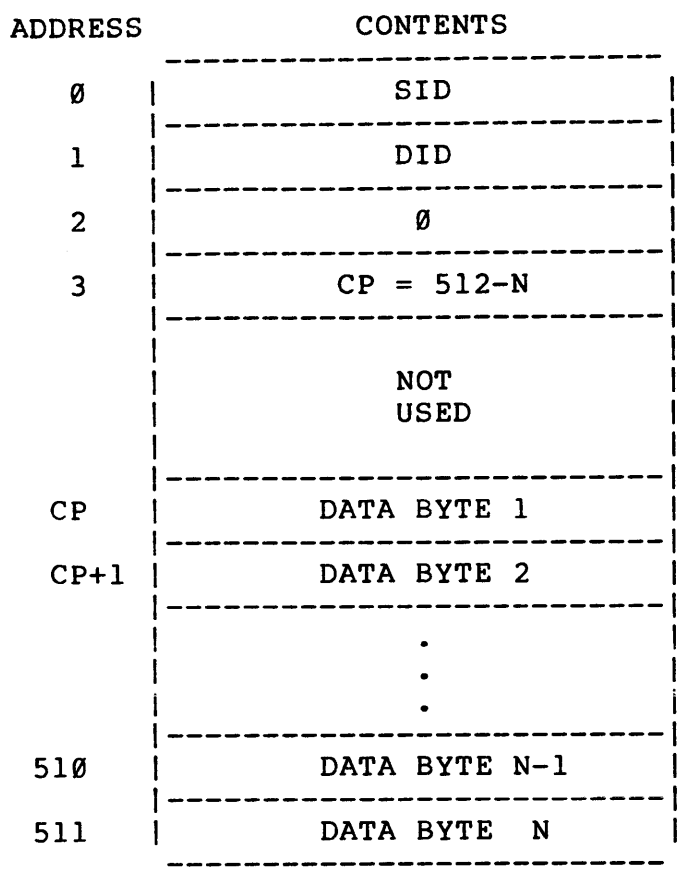

N - DATA LENGTH
SID - SOURCE ID (NOT USED FOR TRANSMIT BUFFERS)
DID - DESTINATION ID (Ø FOR BROADCASTS)

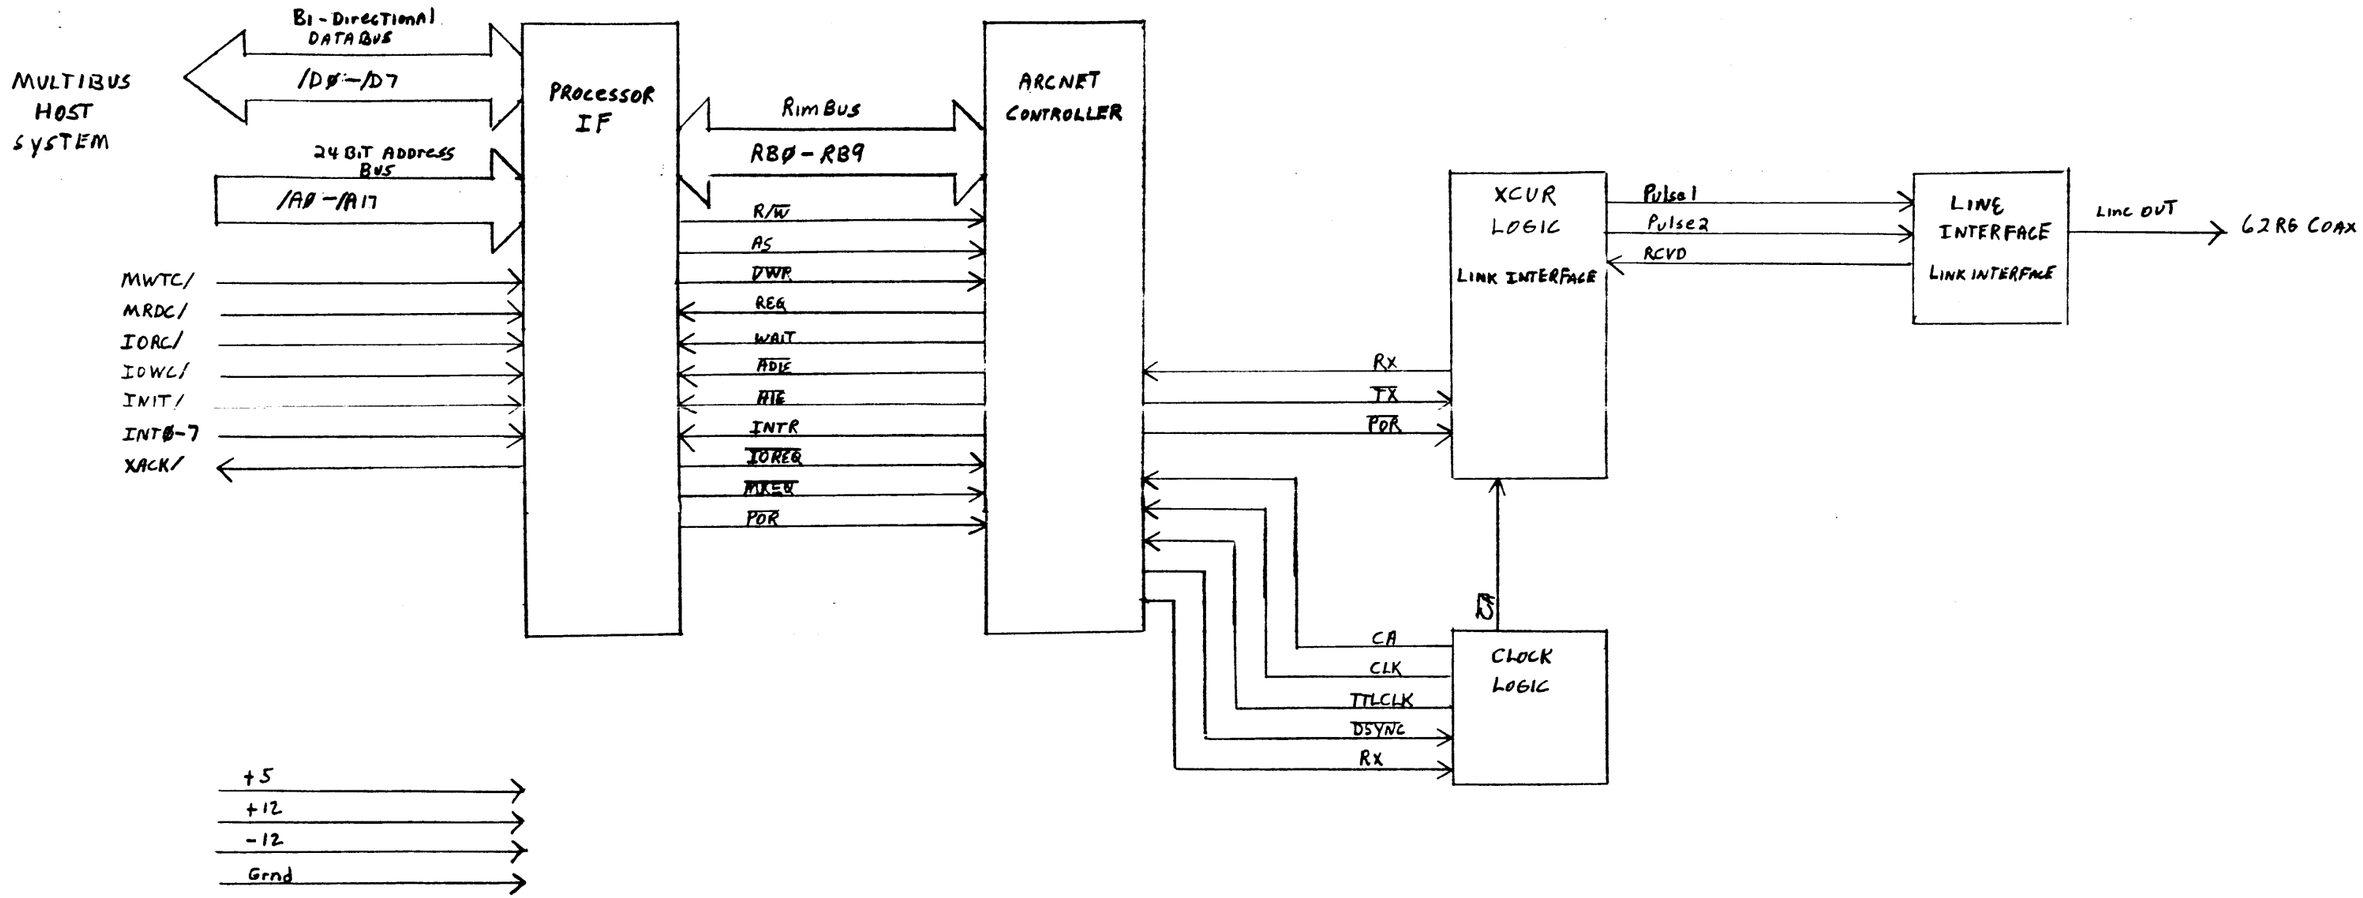
MULTIBUS HOST SYSTEM
Bi-Directional DATA BUS
/D0 - /D7
24 BIT Address Bus
/A0 - /A17
MWTC/
MRDC/
IORC/
IOWC/
INIT/
INT0-7
XACK/
PROCESSOR IF
RimBus
RB0 - RB9
R/W
AS
DWR
REG
WAIT
ADE
ALE
INTR
IOREQ
MREQ
POR
ARCNET CONTROLLER
Rx
TX
POR
XCUR LOGIC
LINK INTERFACE
Pulse1
Pulse2
RCVD
LINE INTERFACE
LINK INTERFACE
Linc OUT
62 RG COAX
CA
CLK
TTLCLK
DSYNC
Rx
CLOCK LOGIC
+5
+12
-12
Grnd

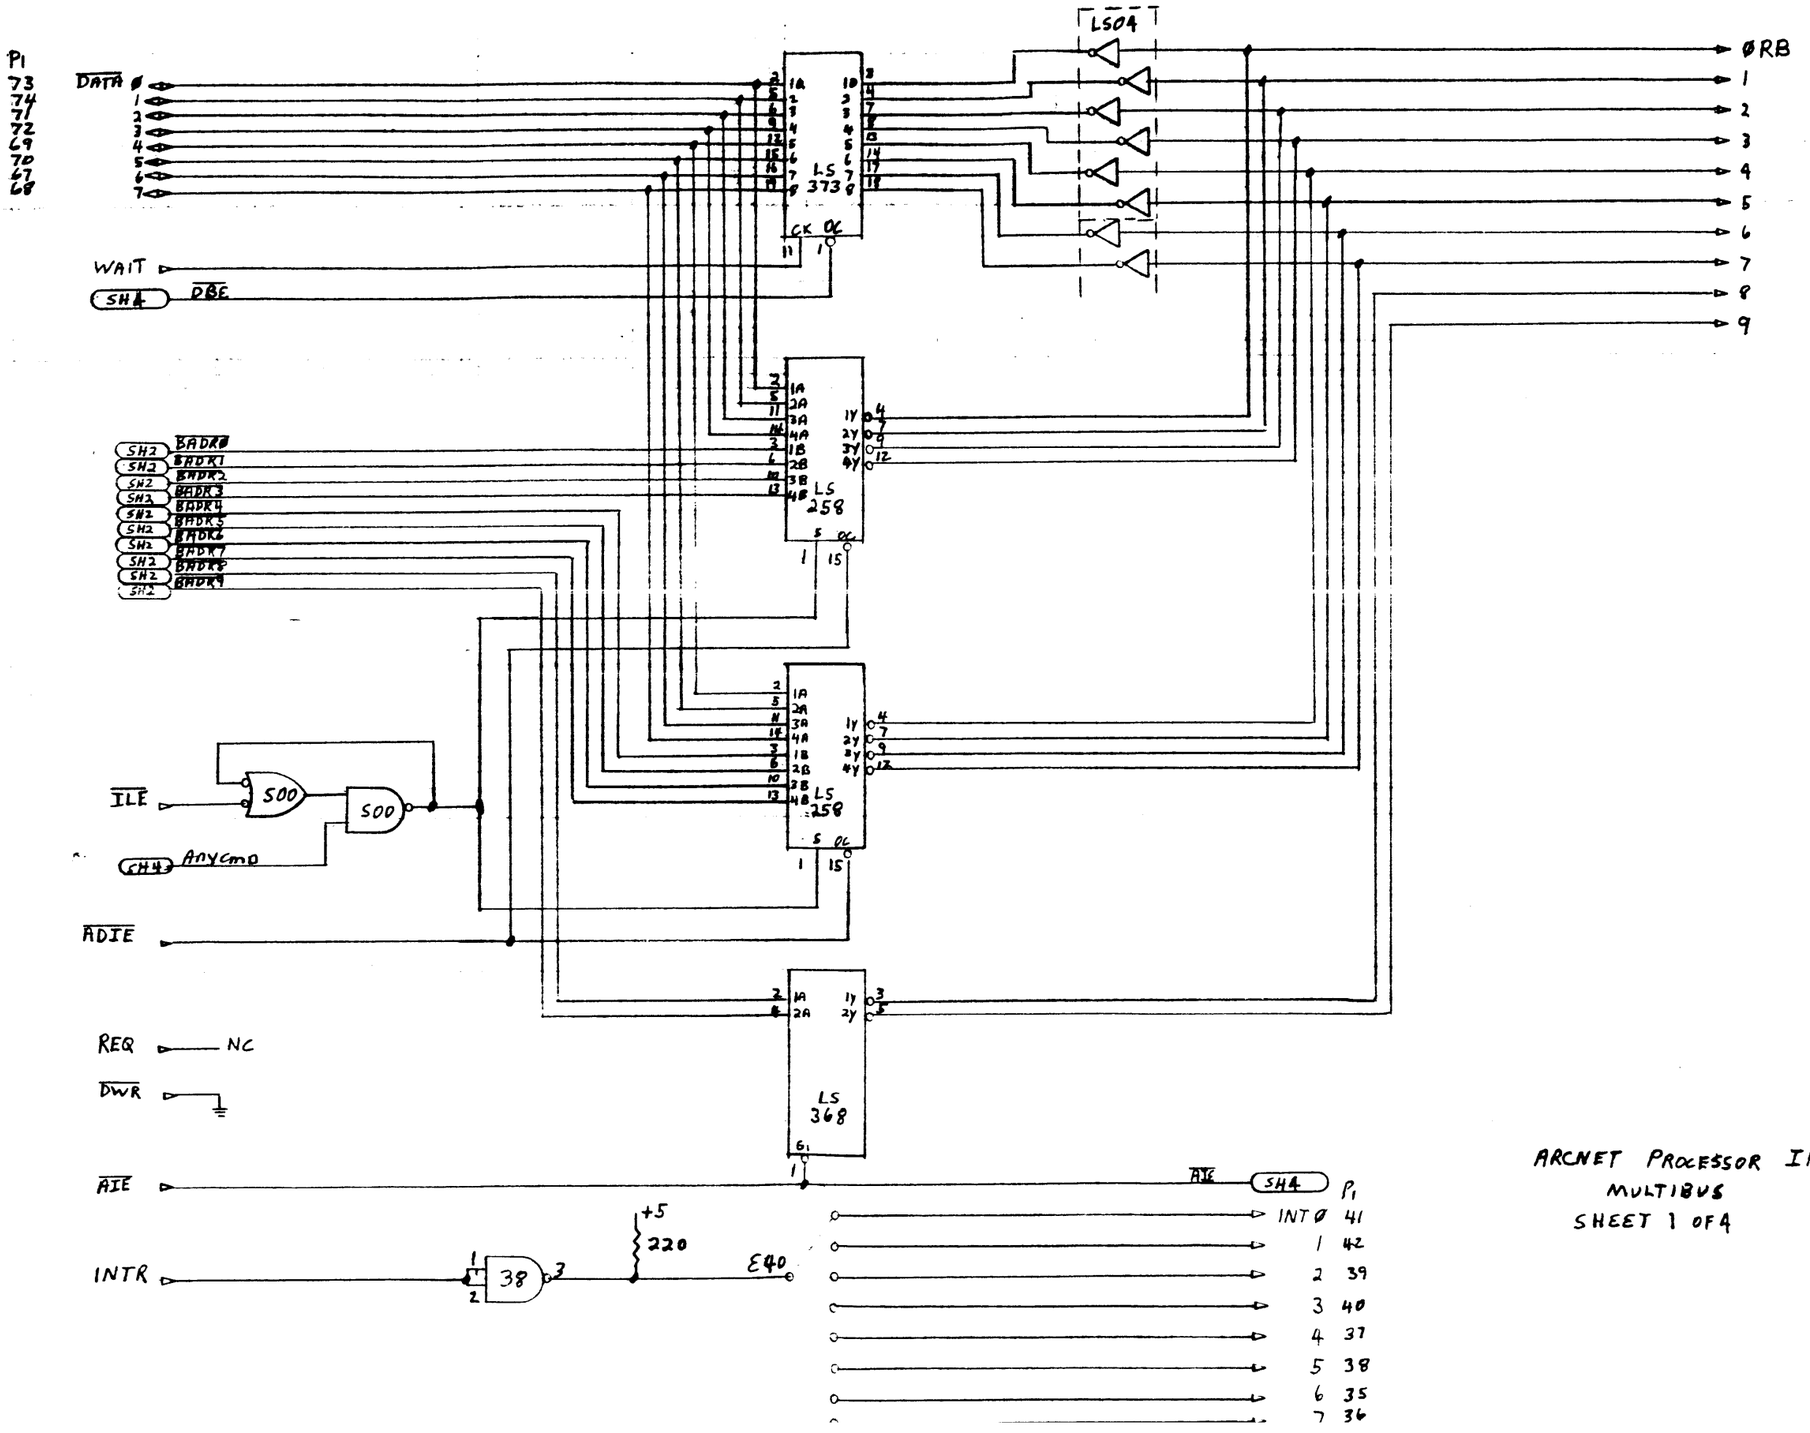

ARCNET PROCESSOR IF
MULTIBUS
SHEET 1 OF 4

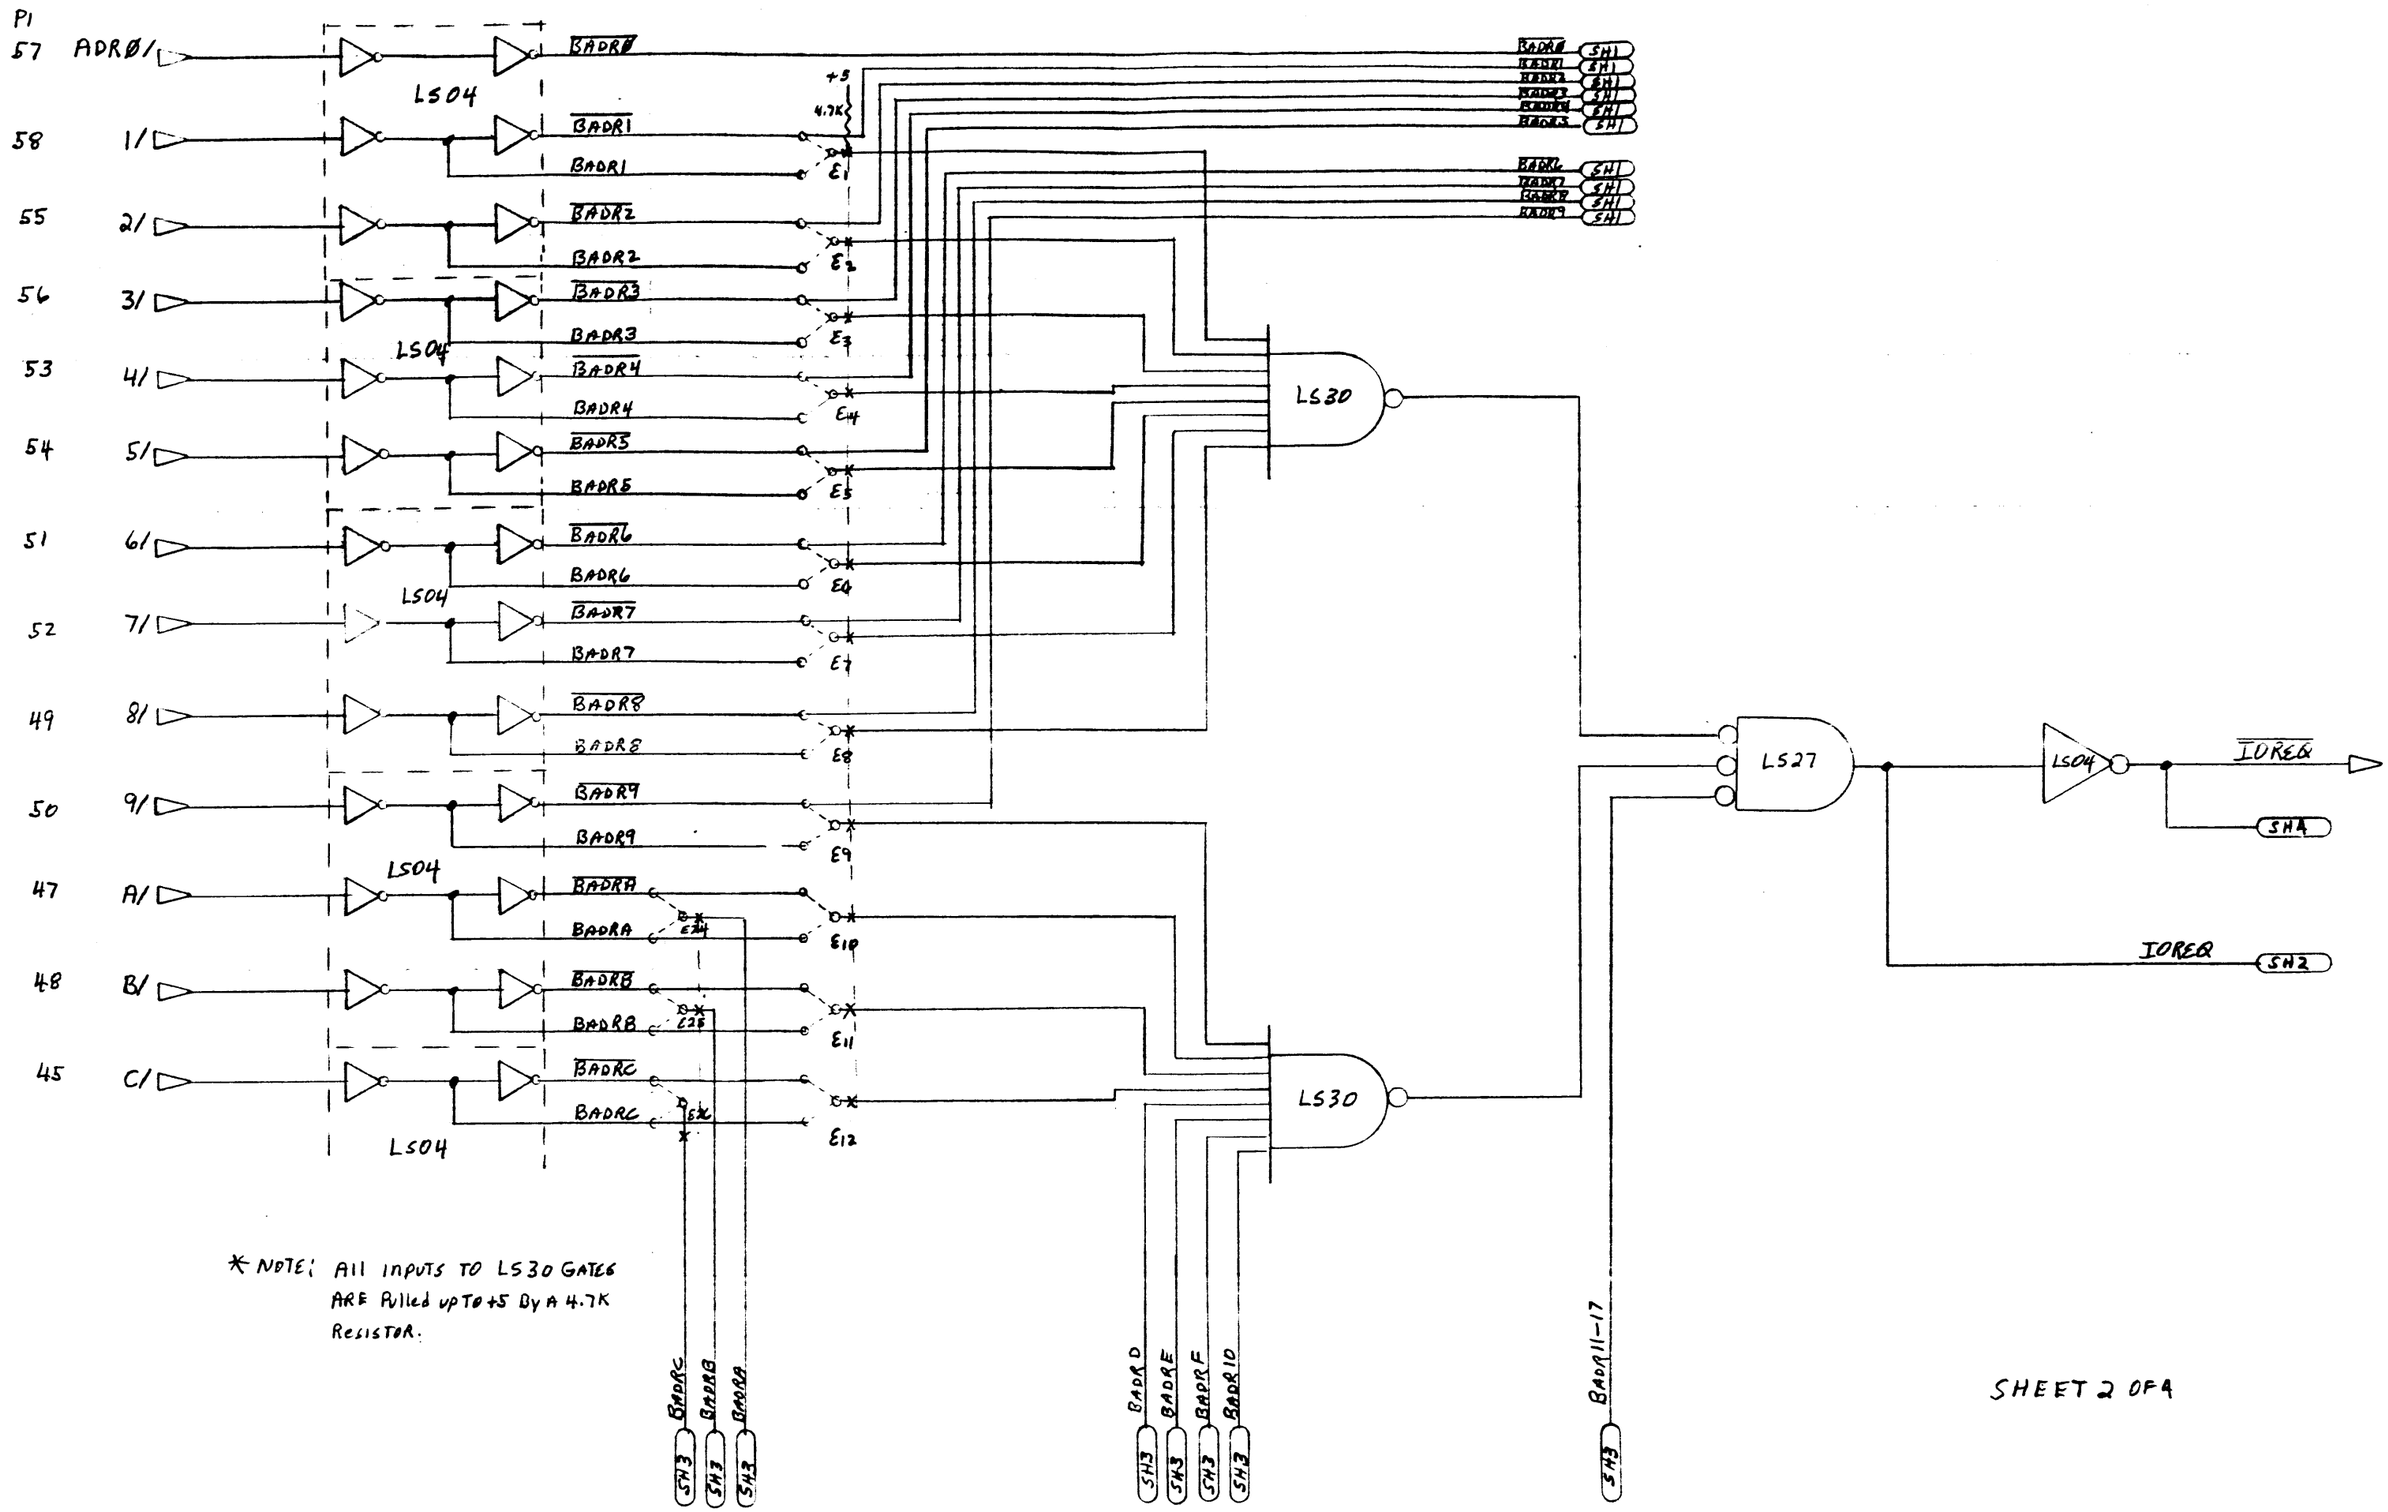

* NOTE: All inputs to LS30 gates are pulled up to +5 by a 4.7K resistor.

SHEET 2 OF 4

JUMPER CHART	
E1 – E23	I/O ADDRESS SELECT
E24–E27	MEMORY ADDRESS SELECT
E38, E39	MEMORY / I/O MODE SELECT
E40	INTERRUPT BIT SELECT

SHEET 3 OF 4

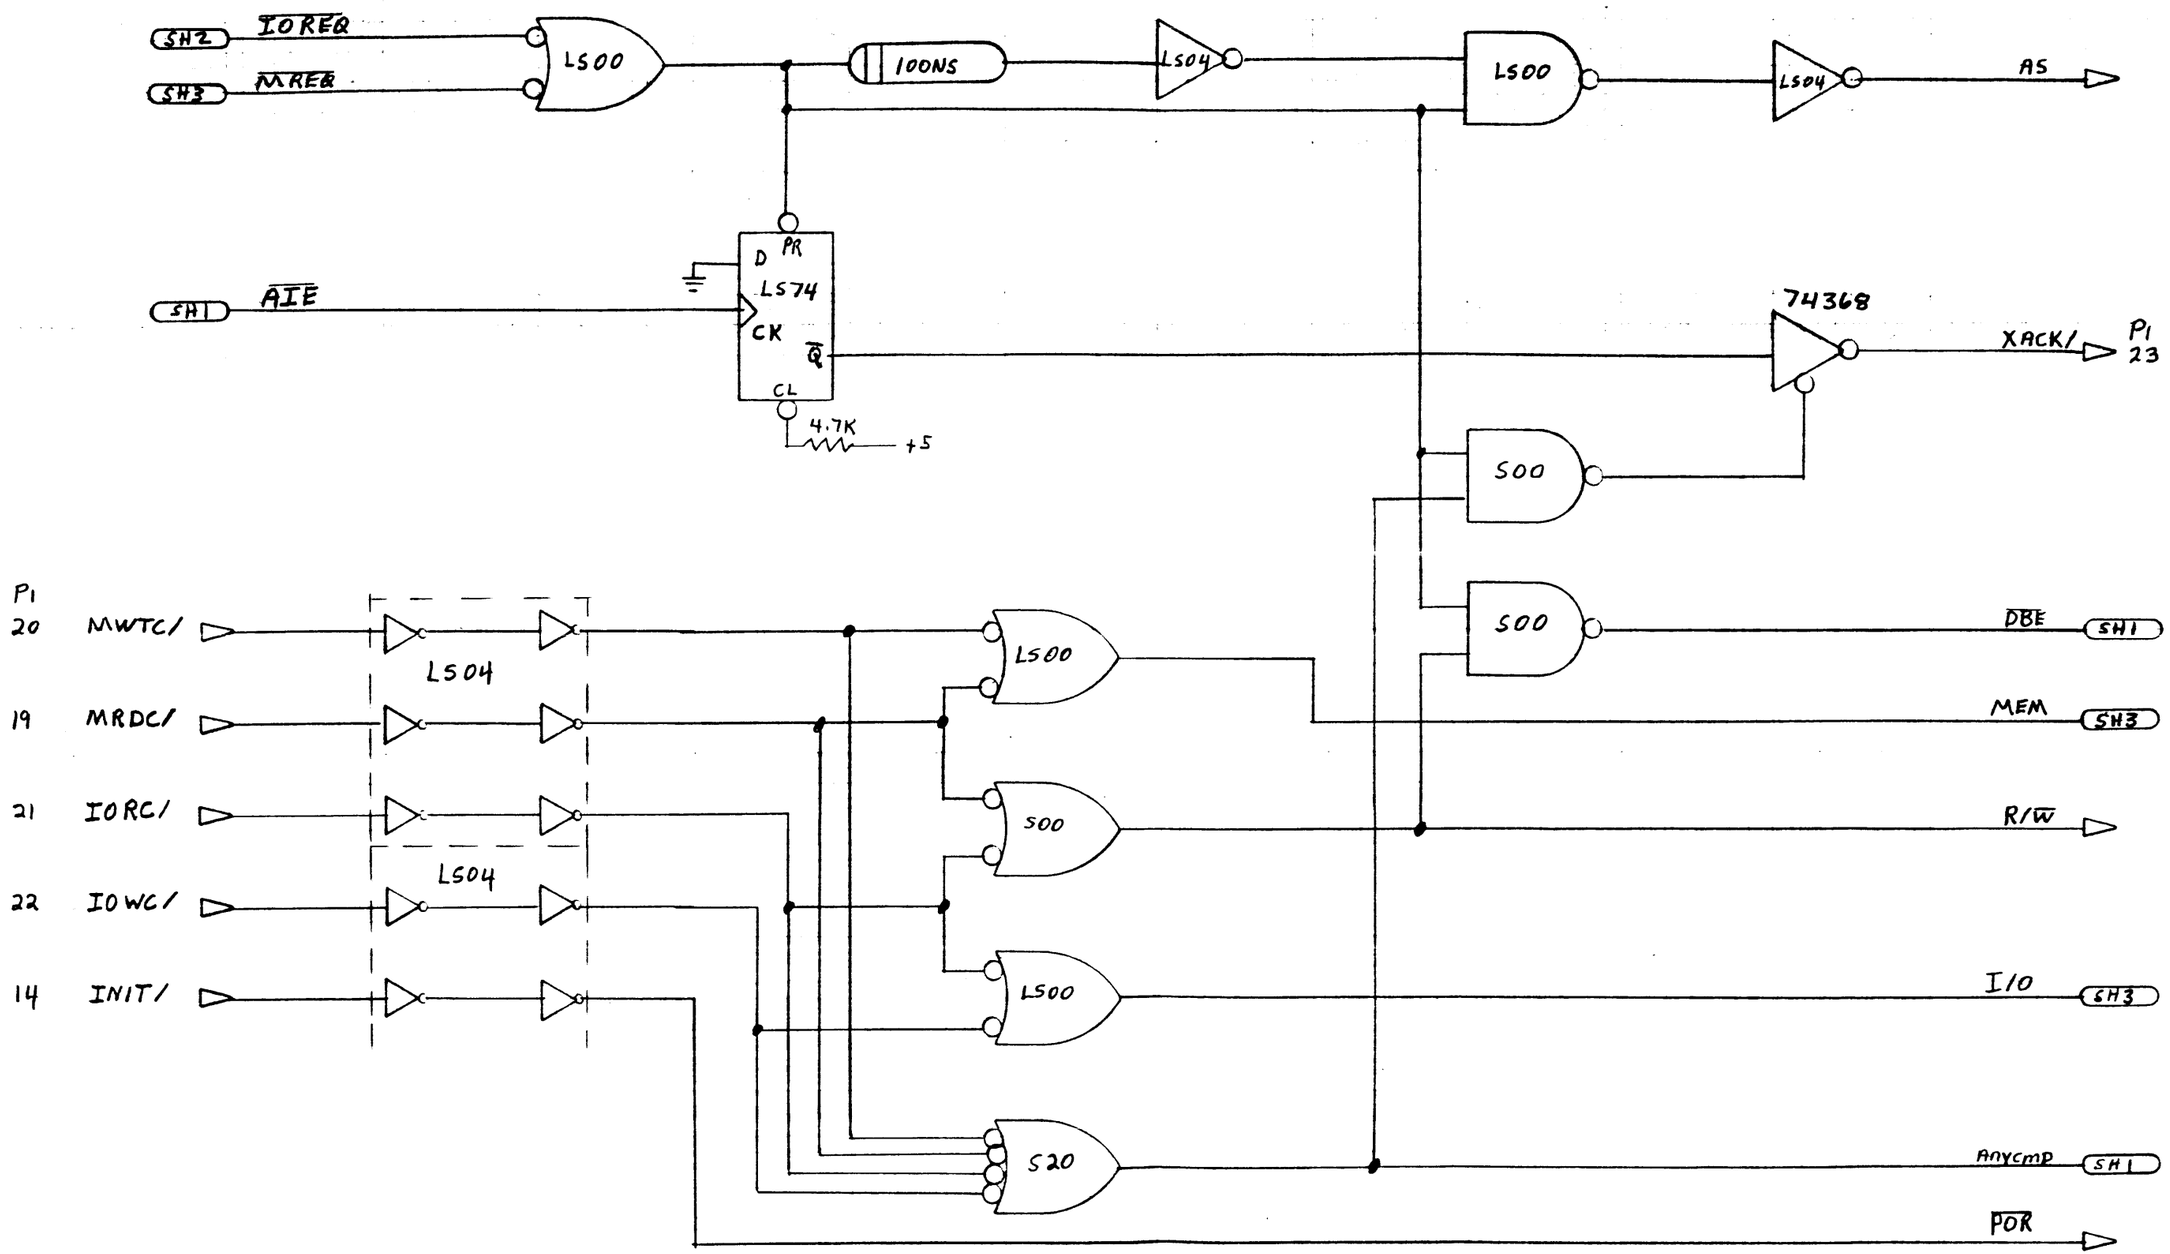
SH2
IOREQ
SH3
MREQ
LS00
100NS
LS04
LS00
LS04
AS
SH1
AIE
D
PR
LS74
CK
Q
CL
4.7K
+5
74368
XACK/
P1
23
S00
P1
20
MWTC/
19
MRDC/
21
IORC/
22
IOWC/
14
INIT/
LS04
LS04
LS00
S00
S00
LS00
S20
DBE
SH1
MEM
SH3
R/W
I/O
SH3
ANYCMP
SH1
POR

SHEET 4 OF 4

4.7K

+5

RBQ 8 7 6 5 4 3 2 1 RB0

AIE

ILE

ADTE

100Ω

NC

RBQ,10

13

AIE

47Ω

AS 10

IOREQ 8

MREQ 9

4 TST2

5 TST1

PDR 48

R/W 7

DWR 6

CLK 18

CA 2

RX 38

1K +5

ECHO 39

1 ET2

3 ET1

34 IDOUT

47Ω

TTLCLK

ILE 18

ADTE 14

WE 17

OE 16

L 15

RB7 21

22

23

24

25

26

27

RB0 28

100Ω

SMC COM 9026

TX 37

DSync 36

INTR 29

WAIT 32

REQ 11

IDLD 33

+5 4.7K

19 23 24

L A8 A9

21 WE

20 OE

18 CS

1K×8

480-4

IO 0

LS 373

11 1

OC

1D 2D 3D 4D 5D 6D 7D 8D

1Q 2Q 3Q 4Q 5Q 6Q 7Q 8Q

TX

DSYNC

INTR

WAIT

REQ

1 +5

1K

16 mSB 2 3 4 5 6 7 8 9 LSB

LS166

SH/LD CLK SER A B C D E F G H

15 7 13

9 1 6

+5 4.7K

RIM ID DIP SWITCH "ON" = 0

ARCNET CONTROLLER LOGIC

ARCNET LINK INTERFACE
XCVR LOGIC A

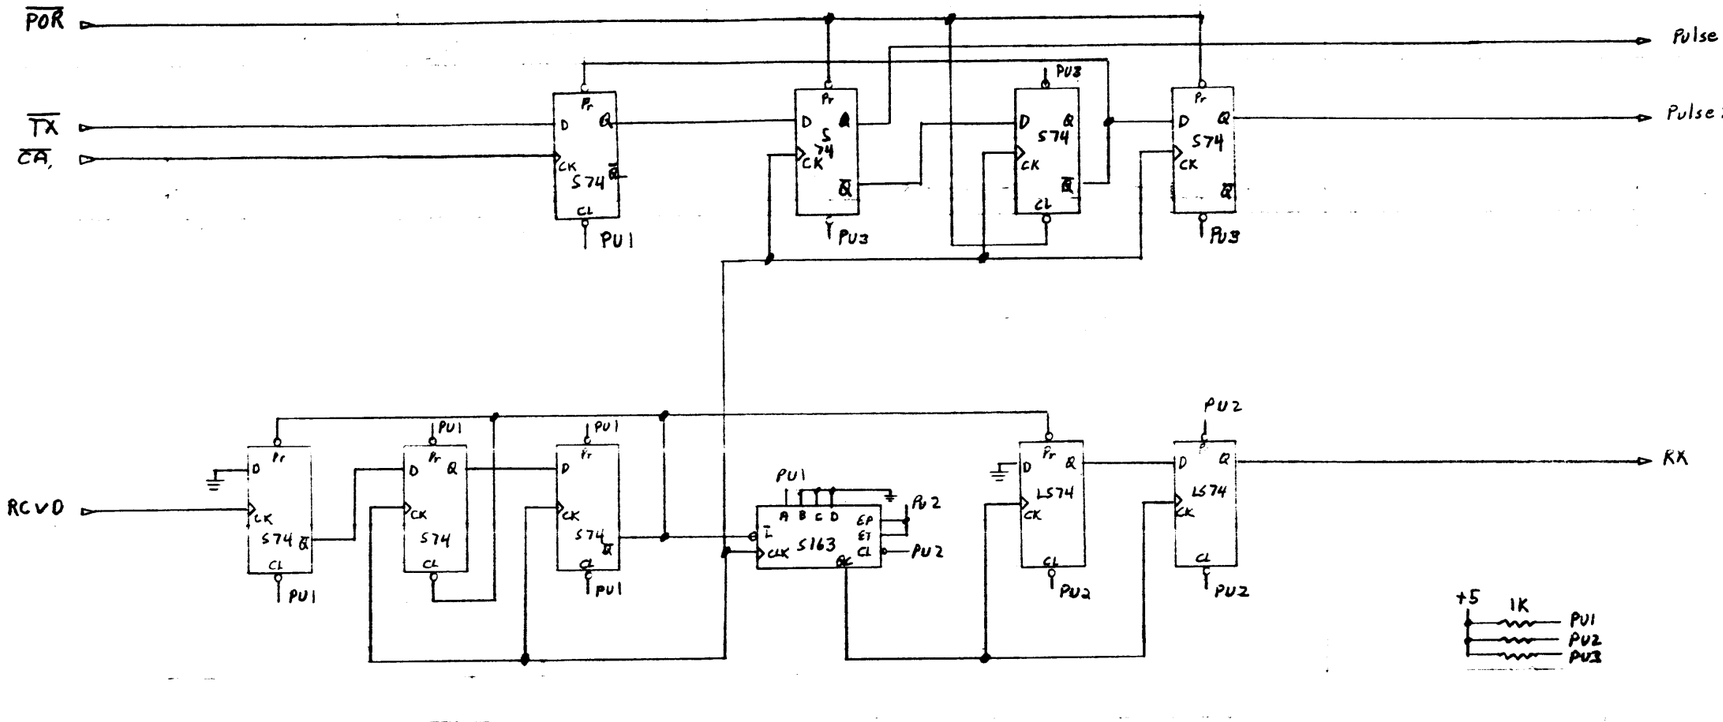

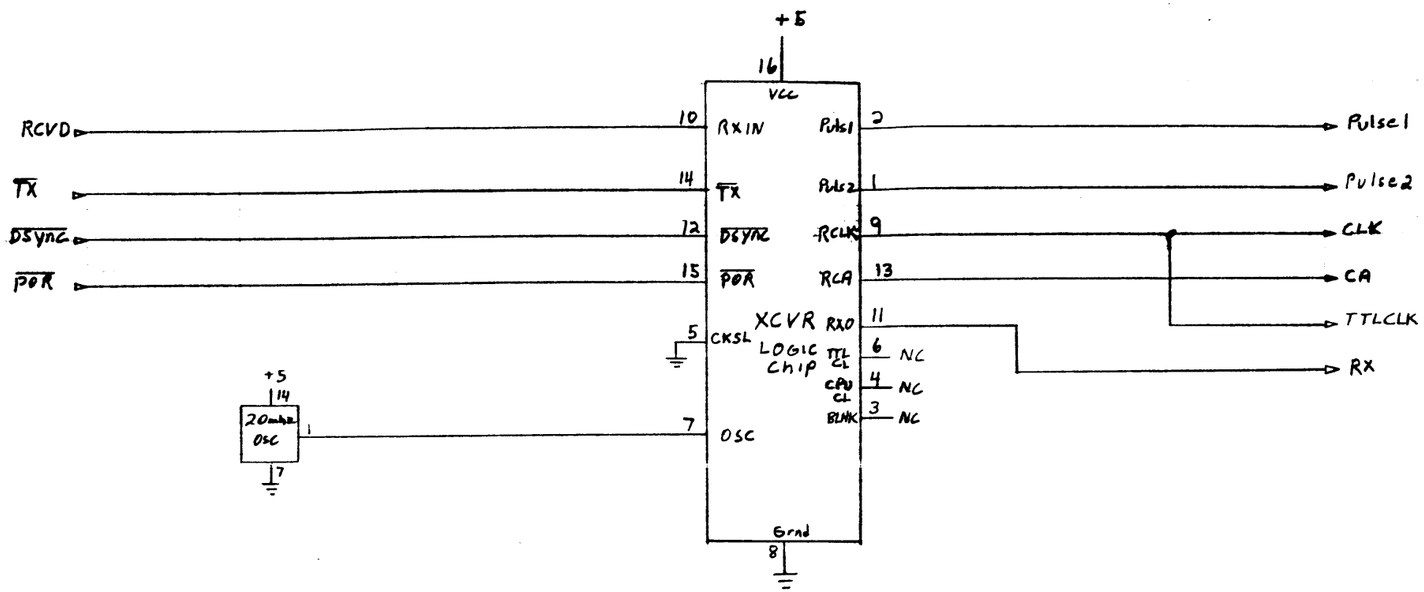

ARCNET LINK INTERFACE
XCVR LOGIC B

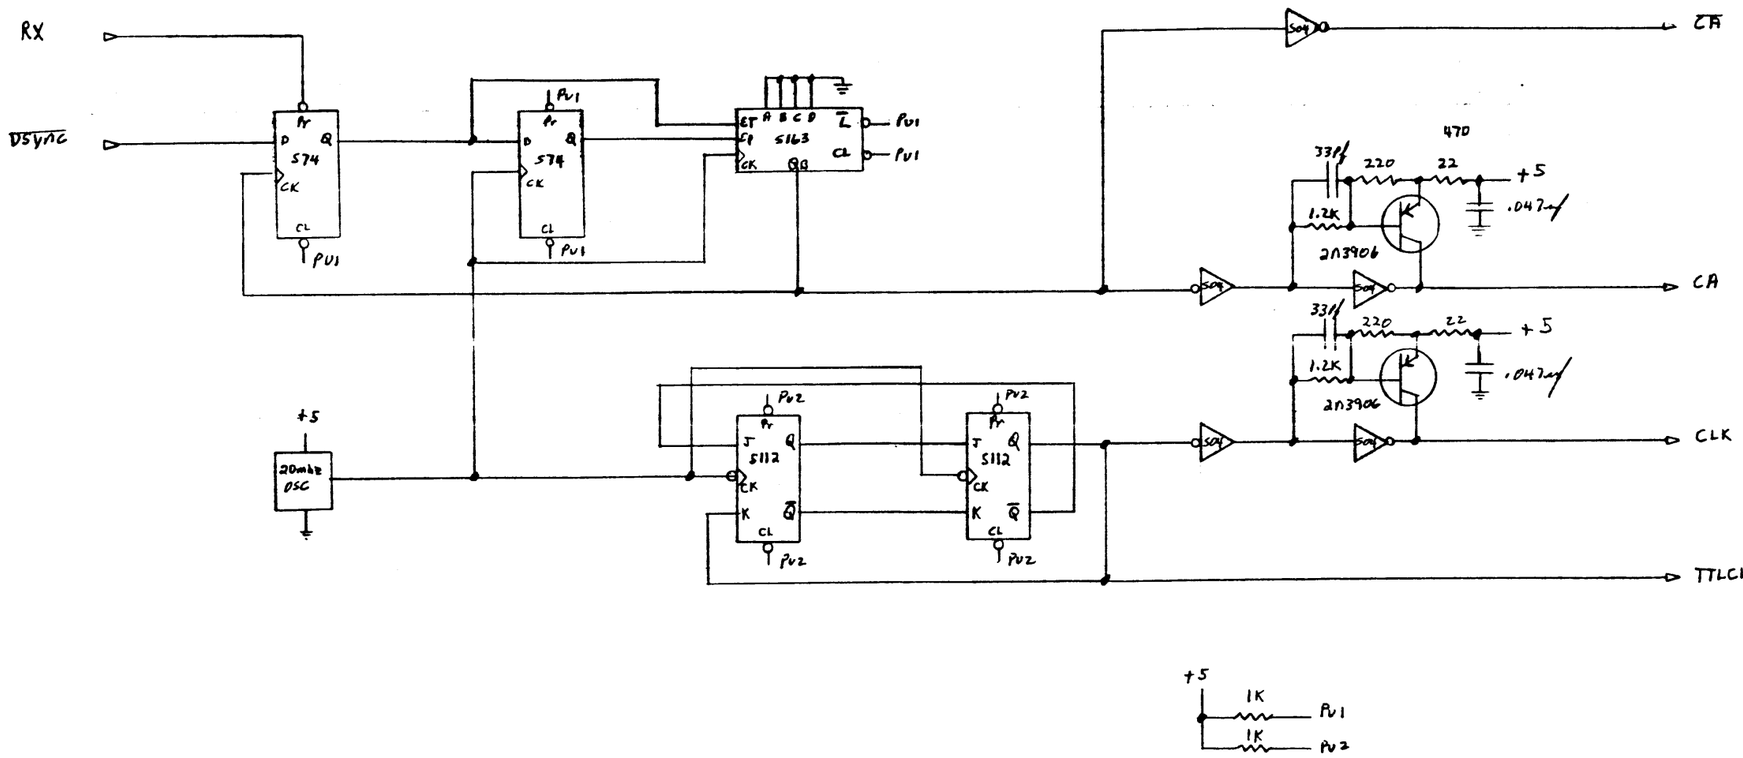

ARCNET
CLOCK LOGIC

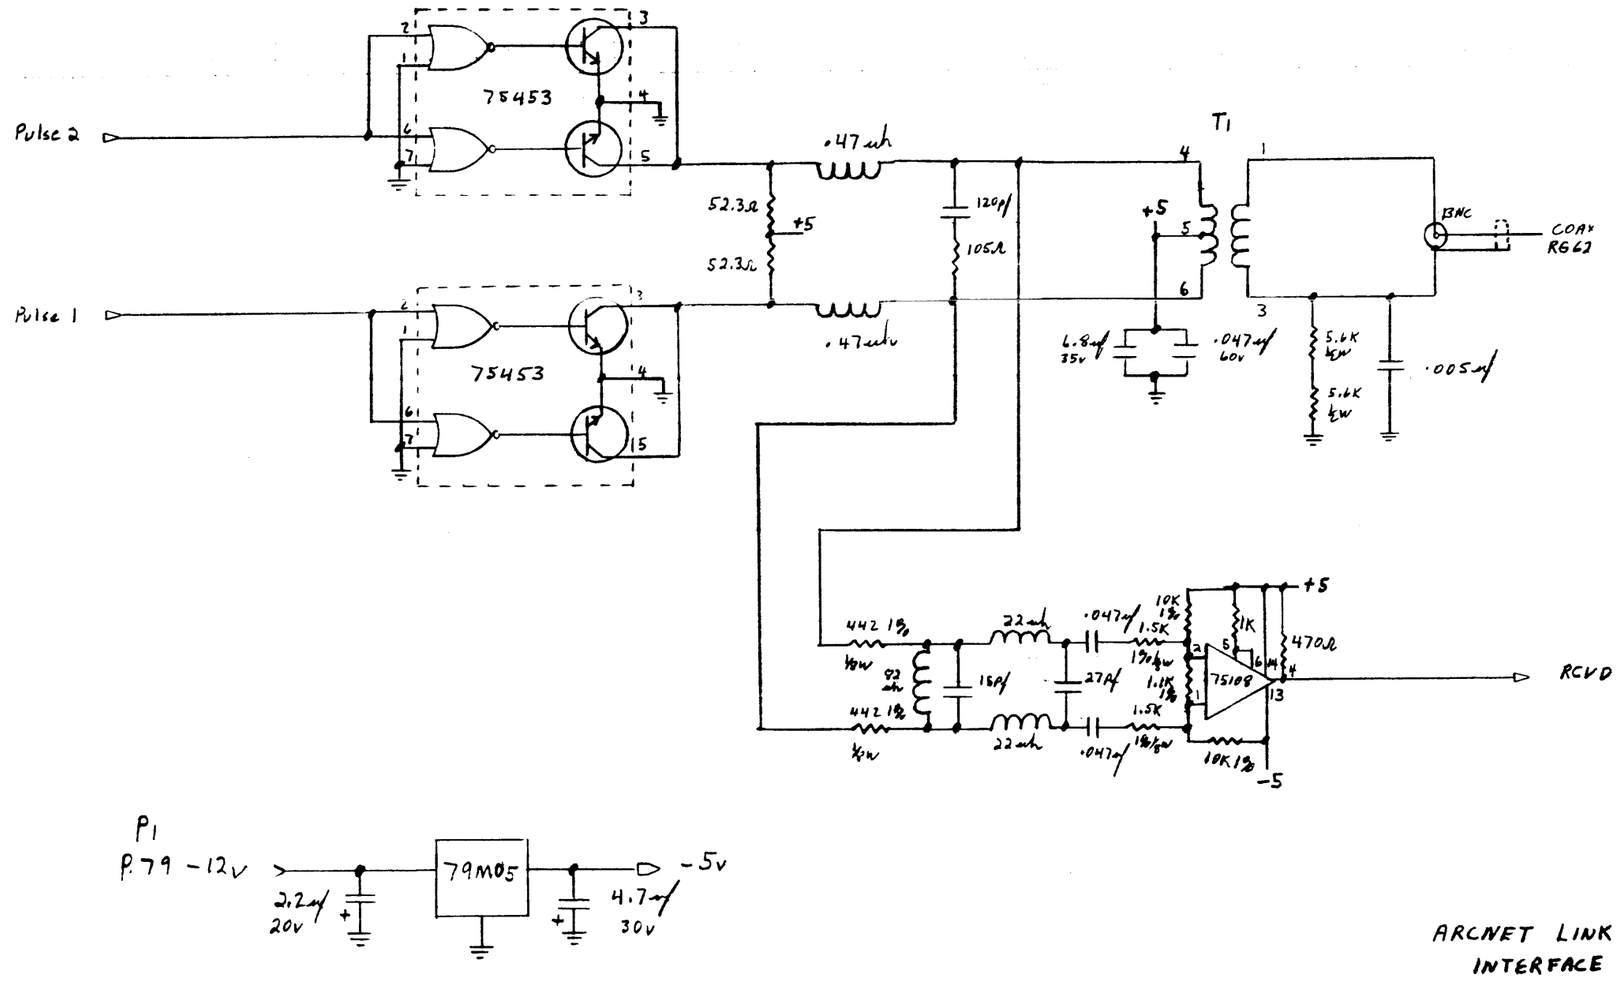
Pulse 2
Pulse 1
75453
75453
.47 uh
.47 uh
52.3Ω
52.3Ω
+5
120pf
105Ω
T1
+5
6.8uf 35v
.047uf 60v
5.6K ½w
5.6K ½w
.005uf
BNC
COAX RG62
442 1%
½w
22 uh
.047uf
10K 1%
1.5K 1%
1K
+5
470Ω
16Af
27Af
75108
RCVD
1.5K 1%
10K 1%
-5
P1
P.79 -12v
2.2uf 20v
79M05
-5v
4.7uf 30v
ARCNET LINK INTERFACE