

Ultra-Low-Cost Network for Personal Computers

Ken Clements and Dave Daugherty
Pacific Polytechnical Corp
POB 2780
Santa Cruz CA 95063

Ten years ago, computer "hackers" listened with glee to predictions that technological advances would soon allow them to buy their very own computers. Indeed, the seers predicted, the computers of the future would fit into a spare bedroom or basement and wouldn't even require air conditioning. The word went out: start saving \$100,000 to be ready when that great time came.

The time came with a vengeance. Today you can hardly take twenty paces around a technical organization, school, or office without bumping into or being addressed by yet another computer.

One of the sad outcomes of this exponential growth was creation of the computer junkie, the unfortunate soul who went out and bought each of the newest computers he or she

could afford. The junkie ended up with a basement full of equipment and a computer habit that could be satisfied only by more spending.

Just when the future was looking grim for these computer junkies, salvation took form and appeared on college campuses. Perhaps the best explanation came from a recruiter from the giant Xmumplex Corp, who took a young graduate aside and whispered, "I have just one word for you: *networks*."

The big-computer companies and an army of computer scientists apparently will be going network crazy for the next ten years. This development thrills the computer junkies because it provides more computer "stuff" to get excited about. And, the junkies calculate, if they could get their own personal networks going,

they might be able to string together all the "coldware" collecting dust in their basements.

What stops most people from going ahead with their own networks is complexity, both in terms of cost and technical considerations. A typical coaxial network "box" may be as difficult to build and interface as was the computer you wanted to network. This stumbling block is particularly large for the computer junkie who owns no two pieces of hardware that are the same. He must come up with a new interface for each one.

But almost all those pieces of hardware have at least one RS-232 serial port. RS-232 was designed to provide point-to-point communication, and it requires some central manager "box" to produce a network. But with as little as one diode per port, two resistors for the ends, and a -12V (V) source, you can turn RS-232 into ULCNET, the Ultra-Low-Cost Network.

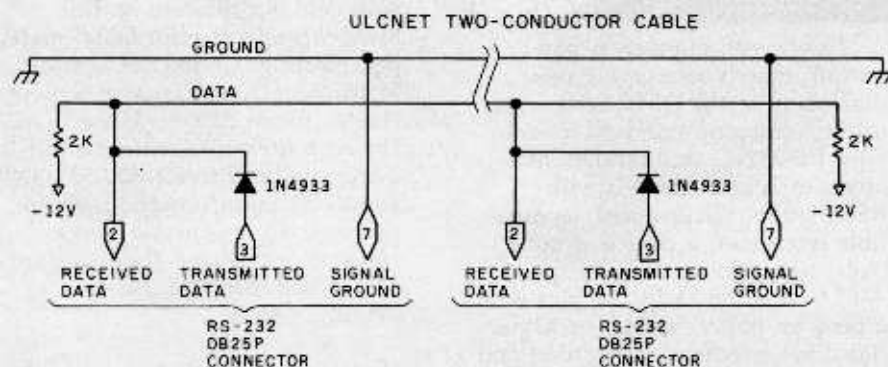


Figure 1: Simplest version of ULCNET. The addition of a diode, cable, and terminating resistors converts RS-232 ports into a basic network for personal computers.

Simple Technique

The primary technique for this transformation is shown in figure 1. It is amazingly simple: just connect a diode in series with the transmit line, then connect the receive line and the diode to your cable. At the ends of the cable you will need resistors to "pull down" the line to -12 V and to help soak up reflections. Serial communications via RS-232 are usually not too fast, so the type of cable and exact terminations are not critical.

ULCNET SHIELDED, TWISTED-PAIR CABLE

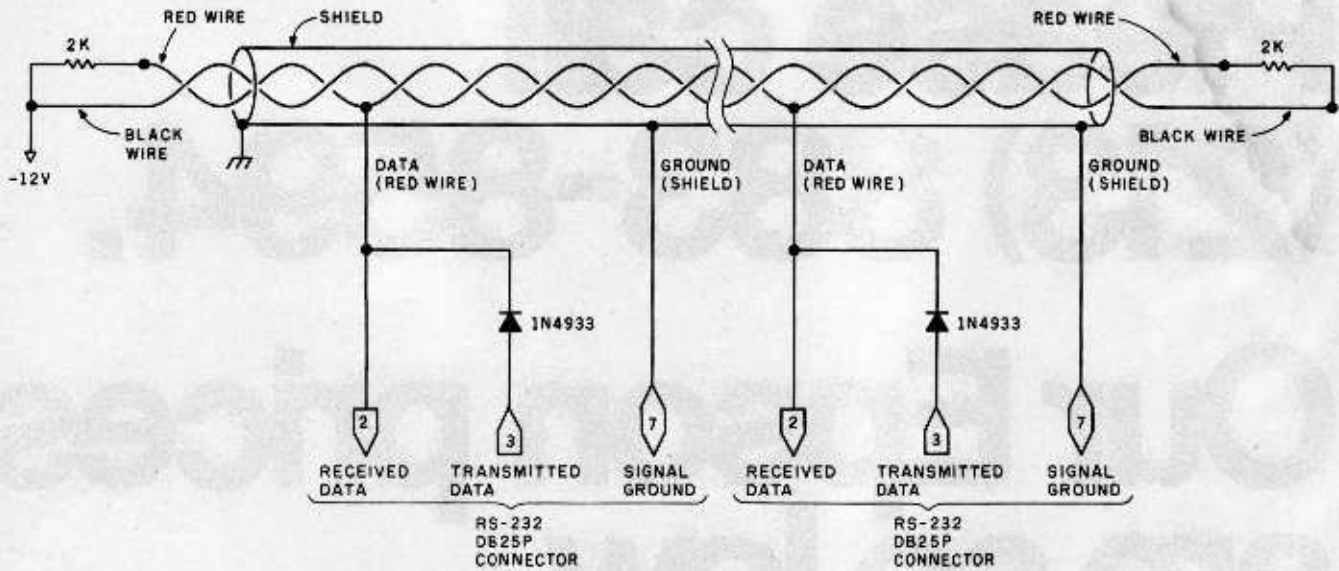


Figure 2: ULCNET for two-wire shielded cable. In this version, a single -12 V power supply is specified, and power is transmitted to the pull-up resistors via one conductor of the cable.

For most applications, it is easy to use shielded twisted-pair cable for the net. This allows one of the wires in the pair to carry the -12 V needed by the termination resistors at the end of

the cable. An example of wiring the termination is shown in figure 2. This technique assumes that somewhere along the line, the black wire in the pair is connected to -12 V and the


shield is grounded.

When characters are sent through an ULCNET port, they are received at all the ports on the net, including the port that did the sending. However, if two or more ports send different messages at the same time, the transmitting ports will each receive something other than what they sent: the logical OR of the two messages. This allows an extremely important property, namely *collision detection* (a property also used in Xerox's Ethernet).

The ULCNET uses the fact that an RS-232 port holds its transmit line at negative voltage when not transmitting, and then pulses the transmit line positive at the start of a character. The RS-232 standard defines a positive level as a transmitted 0 and a negative level as a binary 1. In other words, a character starts with a 0, followed by a byte of code transmitted low-order bit first. At least one binary 1 is inserted after each byte-long character, and it is called the stop bit.

The termination resistors on the ULCNET provide the negative level, and each port may "pull" the line to a positive level by the start pulse of a character. In terms of bits, the resistors supply the 1s, and the ports supply the 0s.

VT 100 Owners:
**THE KEY TO FAST
AND EASY EDITING**
(OF YOUR KED EDITOR)



INTRODUCING VT 100 KEYS FOR THE KED EDITOR... These quality 3-color keys will make editing faster and simpler. They're easy to install—no overlays—just pull out the old keys and push in a new set with the KED functions clearly stated. A set of 18 3-color keys — \$40.

Syntronics 4 Frost St./Cambridge, MA 02140
(617) 739-7710

The speed and distance limits of the ULCNET come from a combination of the drive-current limitations of an RS-232 port and the load each receiver puts on the net. The limits lead to a three-way trade-off of distance, speed, and number of receivers. For example, you might use the ULCNET at 19,200 bps (bits per second) for six devices separated by 20 feet, or you might connect three devices with two miles of wire and run at 300 bps.

Improvements

Some simple modifications can be made to expand the network capability. The first modification gets the number of receivers out of the trade-off equation. Figure 3 shows an alternate ULCNET connection in which an op amp (operational amplifier) is used to buffer the incoming signal. This reduces to almost nothing the load each node places on the network, thereby allowing as many connections as desired on the net.

Some RS-232 ports have +12 V and -12 V supplied on pins 9 and 10 of their DB25 connector (these can be used to power the op amp). Most,

however, do not, so the user will need to run a pair of wires to the power supply of the computer. If some other power source is used, the user must be sure its ground reference is the same as pin 7 of the RS-232 port.

Figure 3 also shows a circuit that drives the DTR (data terminal ready) input of the RS-232 port. This circuit is used to detect activity on the net, and it will assert (pull high) DTR if the net is busy. The circuit works by charging C1, a 0.1 μ F capacitor during the start bit of a character. The capacitor will then discharge through the 330-kilohm (k Ω) resistor R1 when characters are no longer being transmitted. The choice of values for these two components is set by the slowest data rate to be used on the net. The choice shown was picked for 1200 bps operation. If 2400 bps is desired as the lowest rate, then halve R1's value. The resistor can be scaled in this manner for the lowest transfer rate desired. Table 1 suggests resistor values for various data rates, but plan to experiment.

The purpose of the *busy flag* circuit shown in figure 3 is to relieve the software of checking the condition of the

Data Rate (bps)	Size of R, (k Ω)
1200	330
2400	160
4800	82
9600	39
19.2 k	22
38.4 k	10
76.8 k	5.1
153.6 k	2.2

Table 1: Suggested resistor values for running ULCNET at various speeds. Experimentation is suggested.

net, and to provide a signal that can be used with an interrupt-driven system. (These techniques are discussed later.)

Aiming for Speed

Figure 4 is included for those who crave speed. Here, the drive limitation is overcome by using a power FET (field-effect transistor) to drive coaxial cable. The cable can be either standard 50- Ω coax, or the 75- Ω coax commonly used in cable TV operations. Whichever you choose, you must use a matching resistor (50 Ω or 75 Ω) on each end of the cable.

In this form of the ULCNET, the logical 0 is represented by a +12 V level, and the logical 1 is at 0 V. The same busy-detect circuit is used, and all of the network techniques will remain the same. This version of ULCNET is included for those who have very fast controller devices on their ports and want to operate in the 50 kbps to 1 Mbps range.

To make this fast version work, it is important to have a very solid source of +12 V that can put out about one amp for a very short time. The fuse included in figure 4 is meant to shut down the connection if the computer turns on the power FET and leaves it on. If not corrected, this error condition would cause the entire net to halt.

One way to set up a network is shown in figure 5. This setup would allow all the computers to share the hard disk and the printer. The computer directly connected to the hard disk and printer would be partially

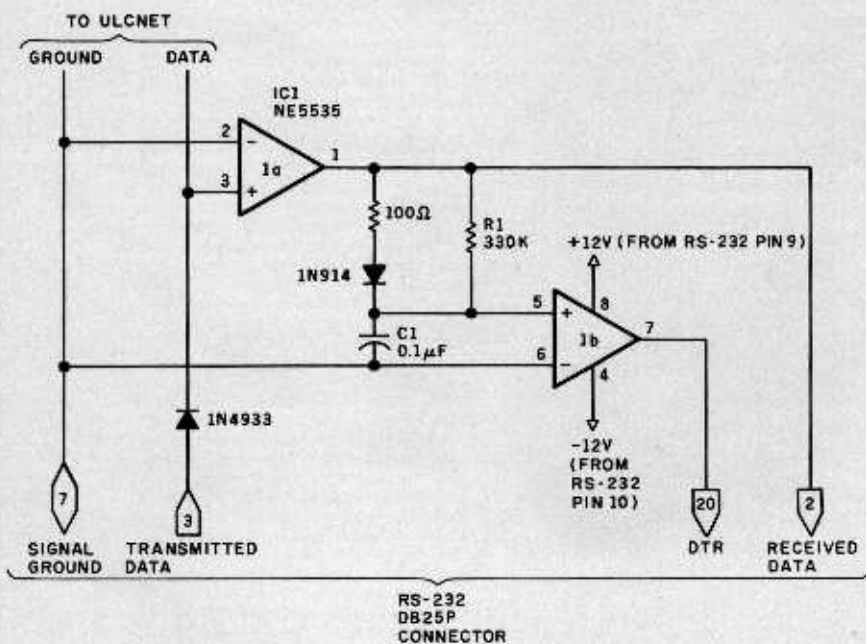


Figure 3: Simple modifications expand network capacity. An operational amplifier reduces the load placed on the net by each node, so that a virtually unlimited number of nodes can be used. Resistor R1 and capacitor C1 control the op-amp comparator to signal that the network is busy. The components shown can be used with speeds as low as 1200 bps; see table 1 for alternate selections.

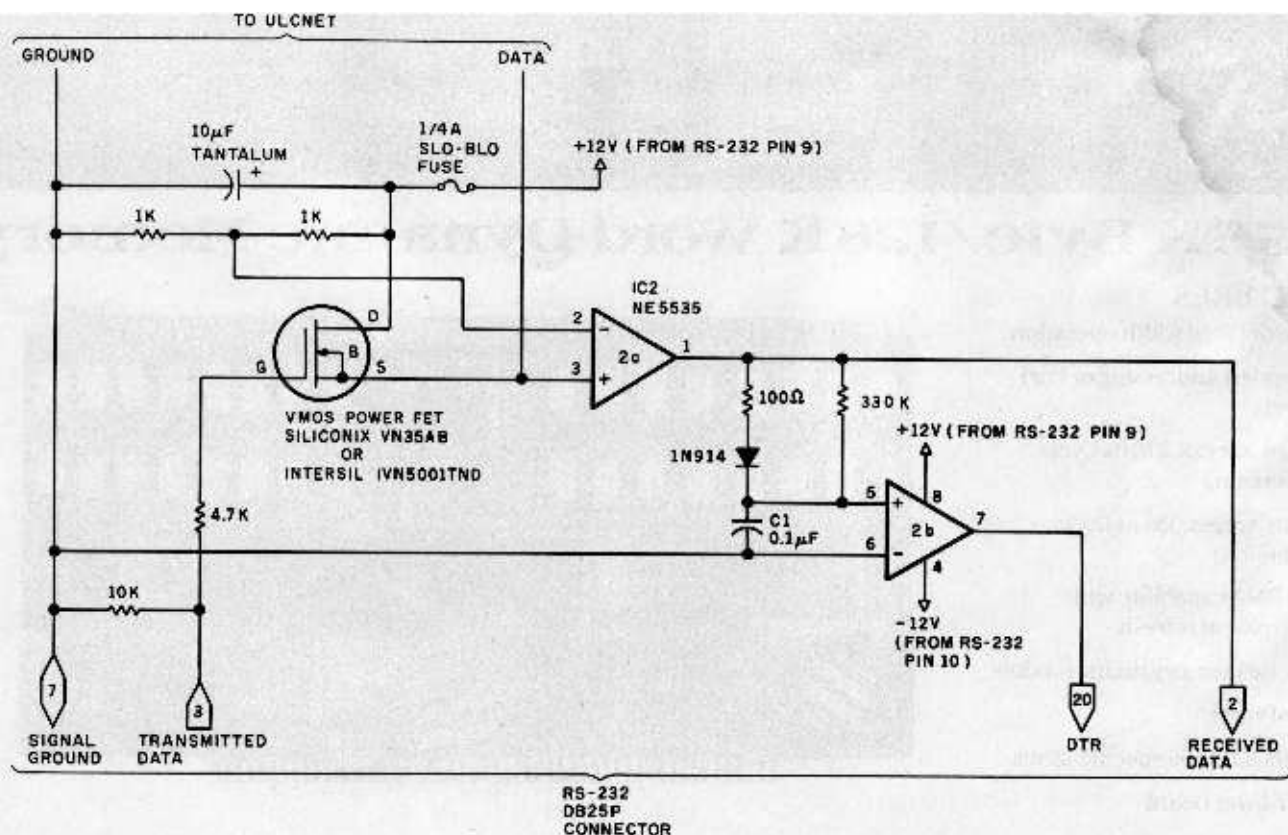


Figure 4: Fast version of ULCNET. The primary limitation of driving power is overcome by installing an output transistor at each port. The transmitter shown may draw as much as 1 A from the -12 V supply, for short periods.

dedicated to servicing the requests for these resources.

Design Issues

Now that we've discussed the hardware for the ULCNET, let's look at some of the issues involved in designing software for the network. These issues are: node-addressing concepts, message formats, task layering, low-level transmission and reception, communication protocols and error recovery, dialogue pipes, special types of networking communica-

tions, and networking under multi-tasking operating systems.

First let's define a node as any device connected to the ULCNET that has the ability to transmit information, receive information, or both.

If there are more than two nodes on a net, some mechanism is needed to uniquely specify the destination of transmitted information. This need is fulfilled by assigning to each node a unique numeric address. A single digit may be sufficient to specify the node for which a message is intended.

Many mechanisms can be used to inform the node's software of its particular address. The possibilities include establishing a switch setting on an input port, including the information in the software for each node (but each node would then need a unique version of the network software), or having the software query the user for an address during initialization.

An address does not necessarily have to be a number, as long as it can be uniquely recognized. It could be a character string such as EVA or SHIRLEY, but you must be willing to pay the cost of pattern matching in order to adopt this scheme.

A *nameserver* mechanism allows the nicety of character strings for addresses without sacrificing the advantage of number matching for decoding addresses. The nameserver consists of a file and a program on a node with mass storage that associates an ASCII (American Standard Code for Information Interchange) string with an address number. The nameserver

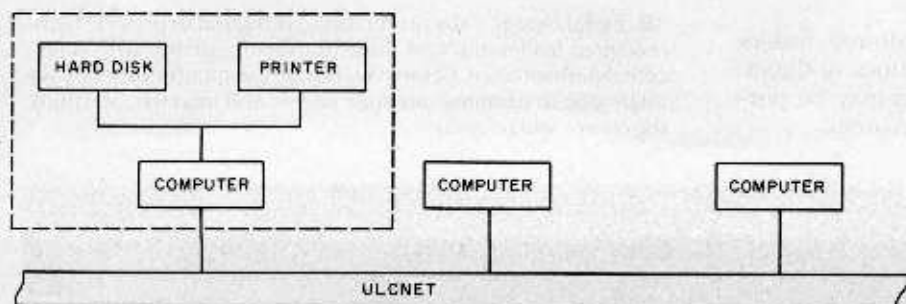


Figure 5: One possible ULCNET configuration. In this example, the mass storage and printer on one computer system can be shared by several other systems.

SOT	TO ADDRESS	FROM ADDRESS	MESSAGE NUMBER	BYTE COUNT	MESSAGE I D	DATA	CHECKSUM
-----	------------	--------------	----------------	------------	-------------	------	----------

Figure 6: Proposed message format. Various fields are included in each message to help the network software ensure reliability of the system.

accepts requests for registration, deregistration, and name queries.

Special generic addresses also can be set aside for special purposes. For instance, the nameserver could be assigned a generic address to be used by all nameserver-related messages, making it unnecessary to know which node the nameserver is actually on.

Another generic address could be set aside to represent a broadcast message—one that all nodes on the network would want to receive. A typical use of a broadcast message is sending a company-wide memo to all employees on the network. The generic address eliminates the need to address the same memo to each person on the net.

Special types of nodes such as mass-storage nodes or printers can have their own addresses. For exam-

ple, the address M might be reserved for the printer node. If there is only one printer on your net, M would mean that printer. If there is more than one printer on the net, an additional field called the logical printer number could be used to specify the printer for which the message is destined.

Message Formats

A message is a predetermined sequence of fields by which two nodes communicate. A message normally consists of several parts: the header, the body, and some kind of error-checking mechanism, such as checksum, at the end.

The structure allows for much variation. The basic component for constructing a message usually is a byte. A field is defined as one or more

bytes that designate a particular section of a message. Typical fields in a message are shown in figure 6 and explained below.

• **SOT:** start of transmission. This byte is useful for informing all receivers that the beginning of a message is now on the net and that the next byte will be the address byte. Obviously, the byte must not be confused with bytes in the middle of a message.

• **To Address:** the address of the intended receiver.

• **From Address:** the address of the node that transmitted the message. As will be shown later, this field is important for sending acknowledgments back to the transmitter.

• **Message Number:** a unique number that distinguishes one message from the next. The usefulness of this field will be illustrated in the sections of this article dealing with duplicate messages.

• **Bytecount:** tells a receiver how many bytes to expect in the message body. It can be used as a receive loop counter, to be decremented each time a byte is received. When the counter equals zero, the user knows the checksum byte will follow immediately.

• **Message ID:** distinguishes three types of messages within a network system. The *data message* contains the essential information to be transmitted from one node to another. The *message acknowledgment* acknowledges a data message, and the third type of message, *ACKACK*, acknowledges a message acknowledgment.

• **Data:** zero or more bytes of information that follow the Message ID.

• **Checksum:** the error-checking byte, computed as the n -bit sum of all the bytes in the message (except the SOT byte and the checksum itself). The transmitter sums up all the bytes in its transmitted message and "ships out" the lower n bits of that sum as the last byte of the message. Meanwhile, the receiver does the analogous operation on the message it receives. If all the characters were received correctly, the receiver's lower n -bit sum should match the transmitter's checksum.

NEW! Speech synthesis using the Votrax SC-01 with the **ads SYNTHETALKER!**

- 64 PHONEMES AND 4 INFLECTIONS PROVIDE AN UNLIMITED VOCABULARY THAT IS USER PROGRAMMABLE
- PROGRAMMABLE OSCILLATOR FOR PITCH CONTROL GIVES YOU MORE NATURAL SPEECH OR MULTIPLE VOICES
- EXTENDED DEVICE OPTION
- UP TO TWO WAIT STATES
- EASILY PROGRAMMED IN BASIC!

Write or call for complete details


Ackerman Digital Systems, Inc.
 110 No. York Rd., Suite 208
 Elmhurst, IL 60126
 (312) 530-8992

*Votrax is a trademark of Federal Screw Works

68000 μ P on the S-100 Bus?

YES, AVAILABLE NOW FROM DUAL SYSTEMS!

- 8 MHz 68000 microprocessor.
- 16-megabyte direct addressing.
- 32-bit internal arithmetic.
- Minicomputer type instructions including MULTIPLY.
- FULL IEEE-696 S-100 compliance. Runs with all 4 MHz S-100 boards and automatically runs faster when accessing Dual Systems memory boards, for FULL SPEED OPERATION OF THE 68000.
- Powerful vectored interrupts. 7 Vectored interrupts, including NMI, as well as alternate mode having up to 256 interrupts.
- On board monitor ROM for immediate use.
- Connector for future addition of memory management unit for multi-user operating systems.
- Built to the highest industrial standards with 200 hour burn-in.



CPU/68000 CPU board. . . . \$1195
 32K-byte 8/16-bit NONVOLATILE RAM board, for secure storage of programs you are developing. Allows FULL SPEED CPU operation. CMEM-32K, per 32K-bytes . . \$895
 32K-byte 8/16-bit EPROM board, EPROM-32K \$395
 Serial I/O board, SIO-2 \$285
 All of the above with cabinet, power supply and backplane. \$3685

OEM and Dealer pricing is available.

Sales representatives in most metropolitan areas.



system reliability/system integrity

DUAL SYSTEMS CONTROL CORPORATION
 1825 Eastshore Hwy, Berkeley 94710
 (415) 549-3854 • TWX 910 366-2035

Layering the Tasks

The network software can be broken up into three separate layers for implementation (see figure 7). These layers are the basic transmitter and receiver subroutines, the protocol layer, and the application program. Breaking up the network software in this manner is useful because it allows the implementer to concentrate on a subset of network functions without having to give much consideration to the rest of the functions. As an added benefit, the layered structure limits the software modifications needed in order to bring up networking capability for particular network tasks and particular machines.

As an example, let's say network software is to be brought up on two of the same type of microcomputers, each having a different serial interface. Subroutines in the transmitter/receiver layer that specifically deal with the serial interface are the only parts of the network software that need changing. On the other end, a printer-application program and a disk-write program should be able to use the same protocol layer and transmitter/receiver layer.

The Transmitter

A buffer and a byte count are the necessary parameters this routine needs from the protocol layer. The transmitter should neither know nor care what type of message is in the buffer. First, the transmitter will need to know if anyone else is currently using the network. In an interrupt environment, this can be determined by a flag set when a character is received and reset when a carrier-detect interrupt occurs. If the flag is reset, therefore, it shows that the network is not in use.

If the transmitter is to be implemented without the aid of interrupts, it will be necessary to wait the length of time needed to receive one character (based on the data-transfer rate). If no characters are received in this time, it is assumed no one is in the middle of transmission.

Once it has been determined the network is not busy, the transmitter must send out the SOT field. A




Figure 7: Network protocol is based on the layer concept. Applications programs deal at a high level by letting the underlying layers do the "dirty work."

potential "race" problem resulting in a collision could occur at this point, since two transmitters could conceivably start this transmission simultaneously.

Because the network is set up so that transmitters receive what they transmit, the received character should always be compared to the character that was just transmitted. If the two characters do not match, a collision has occurred. Later, we will decide how to recover from such a collision.

Assuming the transmitter received what it transmitted, it continues to send out bytes until all, including the checksum, have been sent. If the transmitter is interrupt-driven, it may want to set a flag to inform the protocol layer that transmission was successful. For a transmitter running without interrupts, this information could be returned as a parameter to the routine that called the transmitter.

The Receiver

A receiver activated by interrupts will be able to synchronize with the beginning of a message by the carrier-detect interrupt that occurs after the end of any message. Receivers without interrupts or latched carrier-detect pulses must repeatedly wait until a whole character time has gone by without receiving anything. The next field to be received should be the SOT field. If it is not, it will be necessary to go back to the previous step until an SOT is detected.

Once the SOT is detected, the next field should be the Destination Address. When this field is received, it should be compared with the receiver's own address to determine whether the message is intended for this receiver. If your network supports broadcast messages, all

The Ultimate IEEE S-100 Memory Would...

- BE NONVOLATILE** holding data for up to eight years with the power off.
- RUN AT 6 MHZ** without wait states.
- RUN IN 8 OR 16-BIT** systems with 8 or 16-bit wide data paths.
- HAVE EXTENDED 24-BIT ADDRESSING** and bank select.
- HAVE DYNAMICALLY MOVABLE WRITE PROTECT AREAS** to prevent accidental erasure of programs and critical data.
- GENERATE POWER-FAIL** interrupts for orderly system shutdown & power failure recovery.



...Available Now from Dual Systems

The Dual Systems CMEM memory boards combine high-speed CMOS memories with new 5-10 year lithium batteries to give you the nonvolatility of an EPROM board while retaining the instant writability of a high-speed read/write RAM. These industrial grade boards are ruggedly built and are burned-in for 200 hours.

- CMEM-32K, 32K-bytes ... \$895
- CMEM-16K, 16K-bytes.... \$795
- CMEM-8K, 8K-bytes..... \$695

OEM and Dealer pricing is available.

Sales representatives in most metropolitan areas.



system reliability/system integrity

**DUAL SYSTEMS
CONTROL CORPORATION**

1825 Eastshore Hwy., Berkeley 94710
(415) 549-3854 • TWX 910 366-2035

receivers must check to see if the message is a broadcast message. Additionally, printer and disk storage nodes must also check to see if the destination address is their generic address. If no address match exists, the receiver should go back to hunting for an SOT field (unless this receiver is a gossip monger).

If the message is addressed to a particular receiver, the address and all subsequent bytes should be received and summed together for comparison with the checksum byte at the end of the message. If your particular network uses parity, the message should also be checked for each character received. The receiver should not care what type of message was received; it should simply inform the protocol layer of receipt. With an interrupt-driven receiver, a flag can be set at completion to inform the protocol layer. Additional information, such as whether any errors occurred during the message, could also be communicated to the protocol layer via common memory. If the receiver is not interrupt-driven, this information can be passed back as parameters to the protocol layer.

The Protocol Layer

For the following discussion, *source* will be defined as the node that transmitted the original message, and *destination* as the node to which the message was addressed.

When computer A sends a message to computer B, there is no guarantee that computer B will receive it. Many things could go wrong. There might be a loose connection somewhere. Computer B might not be running, or it might not be listening to the net. Computer C could start transmitting at the same time as computer A.

Protocol schemes detect and correct such situations. Protocol is basically a conversation between a source and a destination, trying to ensure that what the source transmitted was actually received by the destination.

The simplest protocol is one in which the source sends a message to a specific destination and assumes the message arrived. If your network is in good working order and you know

that a particular destination is running properly, this protocol will be sufficient most of the time. You probably would want to use this protocol, for example, when you are sending messages to your friend Carol, who is using computer B. If she is there, she will probably send a message back, thereby acknowledging that she received your message. You'd also use this protocol for broadcast messages, to prevent the net from getting jammed by everyone trying to send acknowledgments at the same time.

When you are doing things on your net, such as writing a file to a disk, assuming the file got there is not enough. You need some real acknowledgment that the file got to the disk. If no acknowledgment comes back from the destination, or if the destination returns to the source an acknowledgment stating that the disk is full, the source will have to take some error-recovery measures. These are discussed later.

What happens if the destination receives a correct message and sends back an acknowledgment that is not received by the source? In this case, the source thinks its original message did not get through, but it actually did. To avoid this situation, an acknowledgment of an acknowledgment received (ACKACK) can be added to the protocol. If after sending an acknowledgment, the destination does not receive the ACKACK, it will have to take some kind of error-recovery action.

What happens if the source receives the acknowledgment and sends the ACKACK, but the destination does not receive the ACKACK? Somebody has got to have the last word, and there can be no guarantee that a message and all its associated protocol are transmitted and received successfully. Especially on a low-speed network, the criterion for deciding how much protocol to use is "as little as possible for a particular application." An intelligent system might provide all three types of protocol (ie: message, message-ACK, and message-ACK-ACKACK) and allow the application program to decide which one to use.

Toughest Boards in Town... IEEE 696/S-100

NON-STOP CLOCK

Keeps time with power off. Our industrial clock utilizes a new lithium battery for 3-9 years use. Easiest clock to program you'll ever see. Runs in all S-100 systems.

- Year, date, hrs, mins, secs, msecs.
- Uses new LSI CMOS chip.
- Vectored interrupts.
- CLK-24 \$250

A/D CONVERTER

IEEE696/S-100 AIM-12 industrial standard module designed for industrial analog-to-digital use.

- Runs in all S-100 systems.
- 32-channel, 16-differential 12-bit resolution/accuracy. 25-microsecond conversions.
- Instrumentation amplifier.
- BASIC program provided. AIM-12, \$695 or \$785 w/1-1000 gain transducer amplifier.

D/A CONVERTER

AOM-12 IEEE696/S-100 industrial level digital-to-analog (D/A) converter.

- 12-bit $\pm 1/2$ L.S.B. accuracy over full 0-70°C temperature range.
- Outputs 0-10, ± 5 , or ± 10 volts.
- Short circuit protection, all outputs.
- Switch-programmable for multiple boards.
- AOM-12, \$575.

VIC 4-20

Standard output for industrial control 4-20 mA D/A converter. Used in conjunction with the D/A board.

VIC4-20, \$445.

DUAL 77 Data Acquisition and Control System-

Built to industrial standards; designed for severe environments. BASIC language makes programming easy. Access to hundreds of sensors. Expandability to meet your increased needs. Nonvolatile memory. Power interruption recovery with automatic restart. DUAL 77 is economical; \$5985 & up.



system reliability/system integrity

DUAL SYSTEMS CONTROL CORPORATION

1825 Eastshore Hwy., Berkeley 94710
(415) 549-3854 · TWX 910 366-2035

Error Recovery

What should be done when a message was sent and no acknowledgment came back? Or when an acknowledgment was sent but no ACKACK came back? Both these cases call for a timing mechanism. A source that transmitted something and is expecting a reply from the destination must wait a certain amount of time for that reply to come back. If the reply does not come back within that time, it will be assumed an error condition exists.

How long should this time be? There is no way to guarantee that a destination really did receive the message and will transmit an acknowledgment within the time the source has set. The waiting time, then, should be more than long enough to cover any reasonable situation.

Once the source has waited a set amount of time without receiving a reply, a reasonable action would be to retransmit the original message at least once more, and again wait the specified amount of time for a reply. The same strategy could be used by the destination when it sends acknowledgments and waits for an ACKACK. If you are doing your network without the aid of a hardware timer, you will need a time-counting subroutine that continually checks to see if a reply was received, and decrements the counter. If the counter reaches 0 before a reply is received, then a *timeout error* exists. If your software has access to a hardware timer, you can use it to set an interrupt.

If no reply is received after repeated attempts to transmit a message, there is nothing to do but give up and report the problem to the program that initiated the network call.

This retransmission scheme introduces another problem. Suppose the source sends a message that is received by the destination, but the destination sends back an acknowledgment that is never received by the source. After timing out, therefore, the source retransmits the original message, and the destina-

tion receives it a second time. The Message Number field, along with the From Address field, can be used to correct such situations.

All receivers should keep a list of the last *n* messages received. The list need contain only the message number and the From Address. When a new message is received, the list should be examined for a match. If a duplicate is detected, the message should be "dumped," but the appropriate response should be sent back to the transmitter of the duplicate message. If the duplicate was an original message, an acknowledgment should be sent back, or if the duplicate message was an acknowledgment, an ACKACK should be sent back.

Collisions are another issue. Assuming that all transmitters check the state of the network before starting transmission, collisions can happen only when two or more transmitters start their transmissions within one character time of each other. When collisions happen, all transmitters involved should immediately stop transmitting and allow the network to return to the "not busy" condition.

Now some kind of mechanism is needed to tell colliding transmitters when they can start transmitting again. If they all wait an equal amount of time, they will collide again. Therefore, they must all wait different lengths of time.

One way to ensure this setup is to establish a priority order based on node address. If a node with the address of 1 collides with a node with the address of 3, then node 1 will wait one unit of time before attempting retransmission, while node 3 will wait three units of time. One problem with this scheme is that under heavy load conditions where collisions are more frequent, nodes with high address numbers may never be able to get a message through because they must wait so long after each collision.

A fairer scheme would be one in which each node has a random-number generator guaranteed to create a unique sequence of random numbers. All nodes would then have

equal priority in retransmissions after collisions.

A Typical Application Program

As an example of a typical application program, let's consider a request to a filing system on a hard-disk node.

The "save" request would first want to send to the filing system a message containing the file name and the number of sectors to be saved. The request probably would ask the protocol layer to expect an acknowledgment and allow the protocol layer to take care of retransmissions if necessary. Along with the acknowledgment would come information from the filing system indicating whether or not the request can be accommodated. If it cannot be accommodated, the request program must report the failure to its caller.

If the request can be accommodated, the save request program must break up the file to be saved into convenient blocks (probably a disk sector). When errors occur during transmission, it is more economical to retransmit small blocks than large ones. In either case, the save request should send an ACKACK to the filing system to say it agrees to what the filing system considers the state of the request.

Once the file has been partitioned into blocks, the save request should hand them in sequence to its protocol layer for transmission to the filing system. The request should ask its protocol layer to expect an acknowledgment for each block transmitted. Each block should have

a unique number that can be checked by the filing system against block numbers already received. In this manner, duplicate blocks can be dumped.

By the value of the last block number, both parties know when the file transfer is completed. If implementation is done in a straightforward manner, the last block number should equal the corresponding field in the original request message.

The save request should ask the protocol layer to send an ACKACK to the filing system when it submits the last block for transfer. Upon receipt of this ACKACK, the filing system can be sure it will not be getting a retransmission of the last block, and it can close the file and forget about the request.

When extended conversations are taking place between two nodes on the net (as in the previous file transfer examples), the network can be made to appear constantly busy by never allowing more than a character time to elapse between messages. In this way, no other user on the network can interfere with the conversation.

If the data rate is controlled by software on the two conversing nodes, you might consider increasing the rate after the initial conversational link has been established. The rate could be increased beyond what's normally acceptable to every node on the network, but it must be changed back after the conversation is completed. While the process is going on, every other node on the network should recognize it as a network-error condition. Because the nodes have

not seen a transition from a busy net to a nonbusy net, they will not be looking for an SOT field anyway. This scheme can get a little tricky when attempting to end a conversation, especially if the last acknowledgment or ACKACK did not get through but the data rate on one node has already been reduced to its former value.

Multitasking Environments

Networking in multitasking environments raises many issues that cannot be considered here, but a few obvious ones should be pointed out.

The protocol layer probably should be set up as a process by introducing another parameter to indicate whether the application program will "go to sleep" waiting for a reply or acknowledgment. The protocol layer would then have to give the application program a "wake up" by indicating whether the message got through to the receiving process.

Since messages could in this way be addressed to one of several processes on a node, the address fields for To and From addresses would need to be extended to include a Process ID number.

The software design presented in this article reflects only one of many possibilities. For more information, or for software if you don't want to write your own, contact Cheshire Software, POB 2780, Santa Cruz CA 95063.

Now that you have a taste of what networking is all about, you can experiment and enjoy implementing your own ULCNET. ■



ADVANCED MICRO SYSTEMS • 26188 ADAMOR ROAD • CALABASAS, CA 91302

MORROW DESIGNS PRODUCTS AT LOW LOW PRICES!

Tel.: (213) 880-4670

OFFER
EXPIRES
Oct. 31, 1981

*HARD DISK SUBSYSTEMS

M-26 DISCUS M26	\$3397.00
M-20 DISCUS M20	\$3653.00
M-10 DISCUS M10	\$2780.00

ADDITIONAL HARD DISK

A-26 DISCUS M26	\$3037.00
A-20 DISCUS M20	\$3193.00
A-10 DISCUS M10	\$2400.00

FLOPPY DISK SUBSYSTEMS

F-1118 DISCUS 1	\$ 816.00
F-1218 DISCUS 2D	\$ 888.00
F-2218 DISCUS 2 - 2	\$1106.00

TWO DRIVE SUBSYSTEMS

F-1128 DISCUS 1	\$1389.00
F-1228 DISCUS 2D	\$1450.00
F-2228 DISCUS 2 - 2	\$1896.00

DISK CONTROLLERS (A & T)

DJ-M26 M26 HARD DISK	\$720.00
DJ-M20 M20 HARD DISK	\$720.00
DJ-M10 M10 HARD DISK	\$720.00
DJ-220B DOUBLE DENSITY FLOPPY	\$349.00
DJ-110B SINGLE DENSITY FLOPPY	\$219.00

STATIC MEMORY (A & T)

16K SUPERAM	\$248.00
16K MEMORY MASTER	\$284.00
24K MEMORY MASTER	\$378.00
32K SUPERAM	\$450.00
65K STATIC RAM	\$870.00

*DECISION I BASIC UNIT \$1342.00

NOTE:

- The disk subsystems include Digital Research CP/M 2.2 and Microsoft Basic V5.2
- If you own a MITS controller and disc drive, we have the interface to run them with our new subsystems, and a reformatter for your Perdec drives and software to translate your Mits files to your new operating system.

SB-2411 SWITCHBOARD I/O	\$219.00
SB-2411-4K 4K RAM OPTION	\$ 80.00
MB-3200 MULT I/O BOARD	\$309.00

WB-0800 8 SLOT MOTHER BRD.	\$109.00
WB-1200 12 SLOT MOTHER BRD.	\$129.00
WB-2000 20 SLOT MOTHER BRD.	\$173.00

All shipments FOB Calabasas — 90 days warranty on all products

Prices, Terms, Specifications subject to change without notice.