NATIONAL RADIO ASTRONOMY OBSERVATORY

Charlottesville, Virginia

January 24, 1975

VLA COMPUTER MEMORANDUM #118

CPU LINK SOFTWARE PROGRAM SPECIFICATION

B. Clark

The software support described here for the Modcomp sixteen bit parallel link gives the capability of 1) using any peripheral device in the system as if it were connected to the CPU in which the program is executing, and 2) writing to or reading from the core of any CPU in the system as if it were a peripheral device.

USE OF REMOTE DEVICES

The use of the remote peripherals proceeds, from a user viewpoint, much as if they were connected to the CPU in which he resides. The User must include, at sysgen time, a DEVICE card for every such remote device to be used from the CPU in question. This DEVICE card must specify the link handler (defined by the appropriate CONTROLLER card) rather than the handler which handles the physical device. Appropriate option bits may be set in the device option word, but bit 15 ('disk') must be set to restrain the MAX I/O system from performing gratuitous and invidious tests on the link status word at I/O initiation time. Except for core-to-core transfers, record lengths must be limited to 256 bytes.

The controller card must have a controller name beginning with 'LK' in order to link with the link handler, which has entry names P$$$LK and D$$$LK. The other quantities to be specified are the device address and the DMP channel numbers. Some words of the PDT (Physical Device Table -- generated by the 'controller' card) are used in a non standard way. They are as follows:

Word 3 is a status word describing the operation in progress: 0--no operation 2--receive remote I/O request 3--originate I/O request 4--receive acknowledgement of previous request 5--send acknowledgement of previous request 6--receive core-to-core transmission request 7--send core-to-core request.

Word 8 is a word count--the first eight words sent are the REX, count, node address, and UFT parameters, this word is incremented once more for the transmission of the data block by DMP.

Word 9 is a pointer to the temporary storage area for the UFT parameters.

Word 11 points to the beginning of the temporary storage area. Bit 15 of word fifteen, if set, indicates master status in case of conflicting requests. Bit 14 indicates pass three. Bit 15 is set at sysgen time, bit 14 is set by the handler.

When an I/O request is originated to a remote peripheral, by means of a REX service, the MAX I/O system generates a service interupt into the link handler. The link handler then analyses the request, and stores the appropriate status word in word 3 of the PDT, as described above. It also places the same word on select lines 13, 14, 15 by means of a select command. It then sets pass 2, and exits.

The status word on the select lines causes a service interrupt in the remote CPU. The remote link handler is thus entered, and acknowledges the request by placing the same word, with bit 15 complemented, back on its select lines. It then sets pass 2 and exits.

The acknowledging status word causes a service interupt in the originating CPU. The originating handler then examines the select lines. If they do not contain an acknowledgement of his request, he must look at his priority status. If he is master, he simply waits until the acknowledgement arrives. If he is slave, he acknowledges the master's request and postpones any further action on his own request. This discipline of the parallel lines prevents both CPU's from trying to use the device at cross purposes. Once the request is properly acknowledged, the originating handler initiates transfer in the register I/O mode, and, with the same command, clears select lines 13, 14, 15 and sets select line 12.

The clearing of the select lines causes a service interupt in the receiving CPU. This CPU promptly sets pass 3, and initiates transfer itself, in receive mode. At this point the following information is transmitted in data receive mode. At this point the following information is transmitted in data interupt-register I/O mode:

The right hand byte of the REX command.

The count from the REX parameter area (valid if read or write)

The node address

The first five words of the UFT, with file name replaced by the logical device name with the initial character moved one forward in the alphabet.

When the receiving handler receives the last of these words, it examines the REX. If the request was a read or a write it obtains a 128 word

buffer from a buffer pool in the handler. If the request was a write, the receiving handler initiates a DMP transfer, input, to the buffer. It also clears select line 12. The action of the receiving handler in taking the last word generates a data interupt in the originating CPU. This CPU waits until the receiving handler has initiated transfer (if the REX was a write) and then itself initiates transfer to send the buffer.

When the DMP transmission is complete, a data interupt occurs in both machines. Both handlers terminate, causing a service interupt in both machines. At this time (pass 3) in the originating CPU, the node is removed from the queue, and operation of the link is cleared. In the receiving CPU, the presence of the received data is brought to the attention of the resident task 'LKT' by means of the common logical variable 'L$BUSY'. LKT then completes the UFT, constructs a calling sequence, and issues a REX for the requested I/O as if it were its own request. The REX is always issued in quick return mode, so that LKT is not held by slow I/O and thus may be used by several tasks in the other CPU.

The link associated task relinquishes control, after initiating I/O, returning to execution at every system event, and checks whether one of its I/O operations has completed. If so, it initiates a REX write to the handler to put it back into operation. The fact that this is an acknowledgement of a remote I/O request rather than an origination is signaled by putting, in the REX call sequence, the address of the UFT in the UFT work area less the start of the work area, instead of the buffer address. The handler recognizes any buffer address less than 100 as being an acknowledgement rather than a new request.

The acknowledgement is transmitted to the originating CPU using the same proceedures and disciplines as the original request. The only exceptions are that on acknowledgement, data is transmitted (by DMP) if the original request was a read (the original request transmits data to the receiving CPU's buffer area if it was a write), and the UFT and other parameters returned are those previously received rather than those used to signal the handler. On the final service interput, the receiving CPU frees the UFT work area and, if one was used, the 128 word buffer. The handler in the originating CPU takes the node address given it when the UFT work area was returned, and places it back on the task's unused node queue. Note that certain embarassing events (for instance aborting the link associated task, LKT) cause a node to be trapped in the receiving CPU, and there is no way to get it back except by master clearing both computers.

CORE-TO-CORE TRANSMISSION

Core-to-core transmission occurs with much the same rules that peripheral device I/O invokes. The core of the remote computer is treated as an I/O device. Core-to-core transmission is flagged by

the logical device name that the request is sent through. Logical device names of COR, COS, COT, and COU signify core-to-core transmission. The connection of a logical device to the core of a particular computer is done by the sysgen 'DEVICE' card, pointing to the proper controller.

Core-to-core transmission differs from I/O transmission in the fact that the DMP data transfers are set up directly between the two core areas involved, and thus the end of DMP ends the entire transaction, with no need for an acknowledgement.

The I/O request for core-to-core transmission modifies the usual meanings of the UFT. The meaning of the UFT words are as follows:

Word 1 is the usual flag word.

Word 2 is the usual file name (this file will be assigned to 'COR' or 'COS', etc.

Word 3 is a three character CAN coded array name, identifying the area in the receiving CPU to be read or written.

Word 4 is a word offset from the array named in word 3, pointing to where the transmission is to begin.

Word 5 is the usual data length transferred (this will always be set to the count parameter of the last REX).

Word 6 is the usual pointer to the file assign list. The user must set words 2, 3, and 4; the system sets the others.

Arrays in the receiving computer are identified to the handler by means of a special REX service. This special REX has been put into the system by a 'ADREX #44,CCT' card in the sysgen. The user loads register 2 with the three character CAN code array name and register 3 with the address of this array, and executes a REX,#44. This array is now identified to the link handler. If the option bit is set in the REX call, this array is instead removed from the list of recognized arrays. If a blank array name is given (register 2 = 0), and the option bit set, all arrays within the task body area of the calling task are removed from the list.

MODIFIED LINK HANDLER FOR USE WITH A NON-MAX OPERATING SYSTEM

A slightly modified version of the handler is provided to reside in the computers without a MAX system, which interfaces with the MAX handler in boss. The modifications are as follows: 1) No provision is made to operate a peripheral in the non-MAX computer - the handler is only for I/O requests originating there and for core-to-core operations. 2) The MAX I/O node queue handling and task master operations are eliminated. 3) The handler is entered by a subroutine call rather than a REX service. 4) The I/O request queuing is strictly first in-first out.

In this version, the user must supply the node as well as the usual UFT, and the convention is established that the node must immediately follow the UFT. Device independent I/O is not provided, so the file assign list and logical device table are superfluous, and have been eliminated; the user provides the logical device name in his UFT instead of the file name.

The (reentrant) routine which queues an I/O operation on the handler is named 'LINK' and has the following calling sequence:

```
        LDI,2     UFTE
        BLM,8     LINK
        DFC       OPERATION CODE (0-8, the same as the I/O REX
                      services)
        DFC       BUFFER
        DFC       COUNT (bytes, as usual)
         :
         :

UFTE    DFC       0,0(logical device name)
        DFC       ((options) or (array name))
        DFC       ((file position) or (array offset))
        DFC       0,0
        RES       5,0       FOR NODE
```

LINK uses registers two through eight.

The array identification for the non-MAX system is provided by reference to an external symbol 'L$TAB' which is a data area of the form

```
          INT     L$TAB,L$END
L$TAB     EQU     $
          DFC     @ (array name 1), ARRAY1
          DFC     @ (array name 2), ARRAY2
           .
           .
           .
          DFC     @ (array name N), ARRAYN
L$END     DFC     0
          END
```