

DEC 3 1975

Cray Research 819 Disk Controller

(expected into Spring 1976; may be changes then.) Not all true for interim controller

The 819 disk controller is a Cray Research product that is packaged in a small freon cooled cabinet. It must be located close to the CRAY-1 computer mainframe and operates synchronously with it. The controller may have from one to four Control Data 819 disk storage units attached to it. Any one of these disk storage units may be transferring data through the controller at any one time. The controller may be connected to a 16 bit mini computer station in addition to the CRAY-1 computer mainframe. If this additional connection is made the station and mainframe may share the controller operation on a function by function basis.

Each of the Control Data 819 disk storage units has two ports for controllers. A second independent data path may therefore exist to each disk storage unit through another Cray Research controller or through a Control Data controller.

The Cray Research 819 disk controller is designed to position and monitor the four disk storage units simultaneously. Separate counters exist in the controller for each of the four disk storage units. The controller may therefore monitor the angular position of each disk and read or write data on any unit without waiting for unit selection and resynchronizing of a common angular position counter. The controller contains two data buffers each of one disk sector capacity. These buffers are used in both reading and writing of data for the disk storage units. The buffers are used alternately as data streams to or from the disk so that data may be read or recorded continuously in every sector of a disk cylinder without losing a disk revolution time. An error correction code is recorded in each sector of disk data so that data recorded over a flaw in the disk surface may be recovered and a flaw mark recorded in a disk reservation table without loss of data at the time of discovery.

The CRAY-1 computer mainframe communicates with the disk controller over one input channel and one output channel. These channels are specially adapted for high speed data transfer in that the communication path is synchronous and the data is organized into units of 64 bit length. The channel data paths are 16 bit parallel paths and data flows in bursts of four 16 bit section . Output data from the CRAY-1 computer mainframe to the 819 disk controller may have one of two data formats.

1) Disk controller function message

This message consists of a 64 bit computer word in which the upper three 16 bit sections are ignored by the controller. The lowest order 16 bit section of the word is translated by the controller for function specification. The message is transmitted as a one word output block from the mainframe and is terminated by a disconnect signal to the controller.

2) Disk write data block

This output data format consists of 512 computer words of 64 bits each. The data is recorded in one sector of a disk storage unit along with control and verification data which is generated by the controller in the recording process. The 512 words of data in the disk write data block do not need to contain any control information for the disk storage unit. The 512 word output block from the mainframe is terminated by a disconnect signal to the controller.

Input data from the 819 disk controller to the CRAY-1 computer mainframe may have one of two functions.

1) Disk controller response message

This message consists of one to four 64 bit computer words containing status information for evaluation by the CRAY-1 monitor program. This message is in response to a disk controller function message from the CRAY-1 mainframe and completes the communication path for monitor control of the disk storage unit operation. This message is terminated by a disconnect signal from the 819 disk controller to the CRAY-1 mainframe.

2) Disk read data block

This input format consists of 512 computer words of 64 bits each. The data is read from one sector of a disk storage unit along with the control and verification data which was recorded with it. The control and verification data is not transmitted from the disk controller to the CRAY-1 mainframe. The 512 words of data are transmitted in a block and normally the block is terminated by a disconnect signal from the disk controller to the mainframe. If an abnormal condition is detected by the controller in the reading process an additional 16 bit word section is transmitted in place of the disconnect signal. This additional word section contains no data and is not intended to be read by the CRAY-1 computer program. It is intended to cause an error interrupt in the computer by overflowing a 512 word buffer boundary.

Disk Controller Function Messages

0000 xuu hhh hss sss Begin disk read
0001 xuu hhh hss sss Begin disk write
0010 xuu xxx xxx xxx Unit reservation
0011 xuu xxx xxx xxx Unit release
0100 xuu xxx xxx xxx Clear fault and return to zero
0101 xuu ccc ccc ccc Cylinder selection
0110 xuu xmm mmm mmm Margin selection
0111 xuu xxx fff fff Status readout

0000 xuu hhh hss sss Begin Disk Read

This function message causes the 819 disk controller to begin reading data from disk storage unit uu at the cylinder last selected for that unit. The first sector of data is read from head group hhhh and at sector ssss. The bit in the function code designated by x is not used.

This function message is transmitted from the requesting computer to the disk controller to begin the reading process. The proper cylinder in the disk storage unit must have been previously selected. The disk controller responds to this function message by transmitting a one word response message to the requesting computer. This response message contains 16 bits of data as follows.

- 1) Error flag
- 2) Unit number
- 9) Cylinder number
- 4) Head group number

In the case of a CRAY-1 computer request the 16 bits of response data are right justified in the 64 bit computer word. If the requested disk storage unit is not on cylinder at the time of the function request the response message is delayed until the on cylinder condition is met and verification data can be read from the disk recording. The error flag bit in the response message is set if there is a fault flag set in the

disk controller for the requested disk storage unit. The fault flags are those defined for the status readout function. The unit number, cylinder number, and head group number in the response message are the values read from the disk recording for the cylinder currently under the read/record heads for the requested disk storage unit. The computer monitor program can verify the values in the response message against the values requested.

After transmitting the response message to the requesting computer the disk controller begins monitoring the sector number for the requested disk storage unit until the requested sector is positioned for reading. That sector data is then read from the proper head group into disk controller buffer A. Transmission of the data in buffer A to the requesting computer begins as soon as the buffer has been filled and the computer has requested input data. The rate of this data transfer is controlled by the requesting computer. Data is sent by the disk controller as appropriate for each input data request. A disconnect is normally sent by the disk controller when 512 words of 64 bit length have been transmitted and another input data request is received. This terminates the block of data and empties buffer A in the disk controller. An abnormal termination of the data block occurs if a check sum error is detected in the data in buffer A or if a fault condition occurred during the reading process. In this case an additional 16 bit word section is transmitted instead of the disconnect signal.

The disk controller advances the requested sector number by one count when buffer A has been loaded with the first sector of data. If the sector number advance overflows the sector modulus (18) the sector number is returned to zero and the head group designator is advanced by one count. The disk controller begins reading data from this newly designated sector into buffer B as soon as the sector is in position and the requesting computer has requested the first word of data from buffer A. If the computer does not request the first word of data from

buffer A in time a disk revolution is lost and the controller waits for the sector to come into position on the next disk revolution.

The disk controller continues to read data from consecutive sectors as long as the computer requests input data and no error occurs in the disk controller. Data alternates between buffers A and B. If an error occurs and an abnormal transmission termination results the disk controller stops reading data and holds the error data for interrogation by the requesting computer. If the sector/head group designators are advanced beyond the head group modulus (10) the disk controller continues by reading the first sector in the disk cylinder. The disk controller remains in this reading mode until a new function message is sent to the controller by a requesting computer.

0001 xuu hhh hss sss Begin Disk Write

This function message causes the 819 disk controller to begin writing data onto disk storage unit uu at the cylinder last selected for that unit. The first sector of data is written on head group hhhh and at sector ssss. The bit in the function code designated by x is not used.

This function message is transmitted from the requesting computer to the disk controller to begin the writing process. The proper cylinder in the disk storage unit must have been previously selected. The disk controller responds by transmitting a one word response message to the requesting computer. This response message contains 16 bits of data as follows.

- 1) Error flag
- 2) Unit number
- 9) Cylinder number
- 4) Head group number

In the case of a CRAY-1 computer request the 16 bits of response data are right justified in the 64 bit computer word. If the requested disk storage unit is not on cylinder at the time of the function request the response message is delayed until the on cylinder condition is met and verification data can be read from the disk recording. The error flag bit in the response message is set if there is a fault flag set in the disk controller for the requested disk storage unit. The fault flags are those defined for the status readout function. The unit number, cylinder number, and head group number in the response message are the values read from the disk recording for the cylinder currently under the read/record heads for the requested disk storage unit. The computer monitor program can verify the values in the response message against the values requested.

After transmitting the response message to the requesting computer the disk controller waits for the first block of data. The controller remains in a write mode until data is received or until another function message is received. The requesting computer transmits a disk write data block of 512 computer words of 64 bits each. This data is loaded into the disk controller buffer A and the transmission rate is controlled by the computer. Error correction codes are generated in the disk controller as the data is loaded in buffer A and these codes are entered in the buffer along with the data at the end of the data block. The write data block is terminated by a disconnect signal from the transmitting computer.

When buffer A in the disk controller has been loaded with the data and the generated error correction words the disk controller begins monitoring the sector number for the requested disk storage unit for the proper angular position. When the requested sector is positioned properly for writing the disk controller streams the data from buffer A to the disk storage unit and records the 512 words of data and the error correction words. When this recording is complete the disk controller advances the requested sector number by one count. If the sector number advance overflows the sector modulus (18) the sector number is returned to zero

and the head group designator is advanced by one count.

The requesting computer may transmit the next block of data as soon as disk controller begins streaming the data from buffer A to the disk storage unit. The new block of data is entered in buffer B. The rate of the transmission is controlled by the computer after the disk controller accepts the first word. Error correction codes are generated in the disk controller and added at the end of the data block in the same manner that was described for buffer A.

When buffer B has been loaded with the data and the generated error correction words the disk controller monitors the sector number for the disk storage unit until the angular position is proper for recording the data in the newly designated sector. The data then streams from buffer B to the disk storage unit and is recorded in the sector following the buffer A data. If the computer is late in filling buffer B with data the disk controller will miss the proper time slot for consecutive sector recording and a disk revolution will be lost. The buffer B data will then be recorded on the next disk revolution.

The disk controller continues to read data from the computer and advance the sector designator for each data block as long as the computer sends data. The disk write data blocks alternate between buffer A and buffer B. If the sector/head group designators are advanced beyond the head group modulus (10) the disk controller continues recording data beginning at the first sector in the disk cylinder. The disk controller remains in a write mode after the last block of data is recorded until a new function message is sent by a requesting computer.

0010 xuu xxx xxx xxx Unit Reservation

This function message causes the 819 disk controller to request the reservation of disk storage unit uu. The bits designated by x in the function message are not used. If the requested disk storage unit is not busy with the other controller port the reservation is made in the unit. The unit then remains reserved until it is released by this controller. If the unit is busy with the other controller port at the time of this function message the reservation cannot be made and an error flag is set in the disk controller for this disk storage unit. This unit reservation error flag remains set until a subsequent 0010 function message when the test is made again.

0011 xuu xxx xxx xxx Unit Release

This function message causes the 819 disk controller to release the disk storage unit uu reservation. The bits designated by x in the function message are not used. The unit reservation error flag for this disk storage unit is set in the disk controller indicating that the unit is not available to this controller.

0100 xuu xxx xxx xxx Clear Fault and Return to Zero

This function message causes the 819 disk controller to clear the fault indicators in disk storage unit uu and position the read/record heads to cylinder zero. The bits designated by x in the function message are not used. If the unit reservation error flag is set in the disk controller for this unit then this function message will have no effect.

0101 xuu ccc ccc ccc Cylinder Selection

This function message causes the 819 disk controller to position the read/record heads to cylinder ccc ccc ccc in disk storage unit uu. The bit designated by x in the function message is not used. If the unit reservation error flag is set in the disk controller for this unit then this function message will have no effect.

0110 xuu xmm mmm mmm Margin Selection

This function message causes the 819 disk controller to select a margin condition in disk storage unit uu. The bits designated by x in the function message are not used. If the unit reservation error flag is set in the disk controller for this unit then this function message will have no effect. The significance of the bits designated by m in the function message are as follows.

- bit 7) Set read late clock
- bit 6) Set read early clock
- bit 5) Set servo offset direction
- bit 4) Servo offset value 2^4
- bit 3) Servo offset value 2^3
- bit 2) Servo offset value 2^2
- bit 1) Servo offset value 2^1
- bit 0) Servo offset value 2^0

Bits 6 or 7 (but not both) in the function message may be used to recover marginal data by shifting the read strobe time early or late in the disk storage unit. This selection is effective for the next read operation only. If bits 6 and 7 are both set an error flag will be set in the disk storage unit.

Bits 0 thru 5 are used to offset the read/record heads from the normal cylinder position. If bit 5 is set the arm moves in the forward direction. If bit 5 is clear the arm moves in the backward direction. The lower order bits in the function message select one of 32 offset positions in each direction. This selection is effective for the next read operation only.

0111 xuu xxx fff fff Status Readout

This function message causes the 819 disk controller to respond with a status message to the requesting computer. The response may be one or four computer words of 64 bit length. The response is essentially immediate and gives the current status of the requested conditions. The response indicates the status of disk storage unit uu and the bits designated by f in the function message select the status type. The bits designated by x in the function message are not used. If none of the bits in the f field in the function message are set the disk controller responds with the status of the error flags in the controller itself. The bits in the f field select various status readouts from the disk storage units as follows.

- bit 5) Read error correction data
- bit 4) Read disk unit interlock register data
- bit 3) Read disk unit offset register data
- bit 2) Read disk unit head register data
- bit 1) Read disk unit cylinder register data
- bit 0) Read disk unit fault register data

The readout data to the requesting computer for a function message with a blank f field is as follows.

- bit 4) Read data checkword error
- bit 3) Parity error in computer data
- bit 2) Cell counter error at index mark
- bit 1) Sector verification error
- bit 0) Unit reservation error flag

The readout data to the requesting computer for a function message with bit 5 set in the f field consists of four computer words of 64 bits each. These words contain the error correction data which may be used by the computer to modify the data in the read data block which was in error.

Disk Recording Information

The disk storage units each hold 2,424,176,640 bits of useful computer data when operated by a Cray Research disk controller. The data is located on cylinders, head groups, and sectors with parameters as follows.

Number of cylinders	411
Number of head groups	10
Number of sectors	18
Words per sector	512
Bits per word	64

The disk surfaces rotate at 3600 revolutions per minute, or a period of approximately 16.6 milliseconds. Maximum cylinder selection time is 80 milliseconds. The data is recorded on the disk surface using four read/record heads operating in parallel. Average data transfer rate when streaming data continuously over a given disk cylinder is 35.4 million bits per second. The total amount of data which can be streamed in a unit in this manner is a disk cylinder capacity of 5.9 million bits. The disk cylinders each contain 10 head groups and a given read/record head scans 18 sectors in a disk revolution. The disk surface under one read/record head is divided into 18 uniform sectors of 8960 cells each. Each sector contains a verification portion, followed by an intra-sector gap and then a data portion. The verification portion of each sector is recorded in a formatting operation when the disk storage unit is placed in operation in the system. This data is not altered during normal disk use. The data portion of the sector, which follows the intra-sector gap contains the data from the computer plus the error correction code generated by the disk controller. This portion of the sector is recorded, read, and rerecorded in normal operation. The cell assignment to the various portions of the sector under each read/record head are as follows.

Pre-recorded portion of sector.

300) preamble

2) sync byte

6) data

6) postamble

58) intra sector gap

Alterable portion of sector.

300) preamble

2) sync byte

8224) data

6) postamble

58) inter sector gap

20.26

The verification data pre-recorded in the first portion of each sector contains 24 bits per sector as follows.

2) Unit number

9) Cylinder number

4) Head group number

5) Sector number

4) Parity check

There is a manual switch in the disk controller for formatting the disk storage units. This switch changes the begin write position for each sector to allow the recording of the verification data and disables the error detection mechanisms.

Error Correction Code

not the same as CDC's - should be better. is s. to have a flag.

The error correction code is a 32 bit code generated for each recording head and associated with 8192 bits of computer data. This code is generated by a Fire code generator at the input to the A and B buffers in the disk controller. As data is received from the computer and entered in the buffer the Fire code generator accumulates a 32 bit quantity which is then added to the data in the buffer. As data is received from the disk storage unit and entered in the buffer the Fire code generator accumulates a new 32 bit quantity which results from summing the data as was done during the writing sequence plus the summing of the error correction words that were recorded with the data. The result should be a zero quantity for each of the four channels of recorded data. If any of these four quantities is not zero an error flag is set in the disk controller for that read operation. This causes the streaming of data to the computer to stop after the erroneous data block and allows the computer to read the four 32 bit quantities for error correction of the data.

This particular error correction code allows the correction of a burst of 11 errors or the detection of any two errors of 11 or less. The code is generated by a 32 bit shifting register which shifts one bit position for each new bit of data processed. The register is cleared at the beginning of the operation. Each bit processed causes the register to shift right circularly by one bit position. If the incoming bit is different from checkword bit 2^{31} then bits 2^0 , 2^2 , 2^{11} , 2^{21} , and 2^{23} are complemented before the shift.

CRAY-1 Channel Parity

The 16 bit output channel from the CRAY-1 computer contains four parity bits associated with the 16 data bits. These parity bits are checked by the disk controller as each word arrives from the computer. If a parity error occurs the parity error flag is set in the disk controller and this information will appear as an error flag in the next response to a read or write function.

The 16 bit input channel to the CRAY-1 computer from the disk controller contains four parity bits associated with the 16 data bits. If a parity error occurs during a CRAY-1 input buffer operation the channel control module in the CRAY-1 computer interrupts the input buffer and the monitor program then responds to a short block of input data.