

A Control Unit for a DEC PDP-8 Computer and a Burroughs Disk

AD VAN DE GOOR, STUDENT MEMBER, IEEE, AND C. GORDON BELL, MEMBER, IEEE

Abstract—The control unit maps a semicontinuous address space of the computer into a segmented BCD-addressed address space of the disk. An interesting buffering system provides for the maximum memory access time of the PDP-8 and high disk transfer rate. In principle, the buffering could be generalized to handle arbitrarily long waiting delays. The position of the disk may be read under program control. The execution of erroneous command sequences are not permitted (and hardware detected). Out-of-range disk addresses are detected.

Index Terms—Binary-to-decimal conversion, buffering, control unit, disk addressing, parity generation/checking.

A PROGRAM swapping and file storage device was needed for a PDP-8-based time-sharing system. Because of its high transfer rate, low access time, and low cost per bit, the Burroughs model 9370-2 disk was selected.

The control unit interfaces this disk with a DEC PDP-8 computer. Data transfers with the PDP-8 occur on a cycle stealing basis ($1.5 \mu\text{s}$ memory cycle/12-bit word).

The main characteristics of the control unit are the following.

- 1) The storage layout on the disk is transformed to look like a page-oriented structure (1 page = 128 12-bit words).
- 2) The semicontinuous binary address space of the computer is mapped into a segmented (BCD)¹ addressable address space of the disk.
- 3) A buffering system was incorporated because the maximum core-memory access delay time was greater than the time interval between words from the disk. This buffering also takes care of the conversion between the disk's 8-bit bytes and the computer's 12-bit words.
- 4) The current position of the disk can be read and transferred to the PDP-8 upon command.
- 5) Parity generation and checking is done.
- 6) Erroneous command sequences are prevented from being executed, and out of range address requests are detected.

Manuscript received June 9, 1969. This work was supported by the Advanced Research Projects Agency of the Office of the Secretary of Defense under Contract F 44620-67-C-0058 and is monitored by the Air Force Office of Scientific Research. This paper was presented at the 1969 IEEE Computer Group Conference, Minneapolis, Minn., June 17-19, 1969.

The authors are with Carnegie-Mellon University, Pittsburgh, Pa.
¹BCD is binary coded decimal.

DESIGN CONSTRAINTS

The design can be looked at as having to satisfy constraints.

The constraints can be divided into:

- a) Burroughs disk (fixed) constraints,
- b) PDP-8 (fixed) constraints,
- c) TSS/8 performance (variable) constraints,
- d) cost (variable) constraint.

Burroughs Model 9370-2 Disk Constraints

The disk unit consists of a single disk of which both sides (surfaces) are used. Each side has 100 data tracks together with 100 fixed READ/WRITE heads. Addressing a track has to be done with a BCD number (range 00-99). A track is divided into 100 data sectors addressable through a BCD number (range 00-99) and a maintenance sector which has the address BB (in hexadecimal).

A special track called the address track is provided. From this track sector numbers can be read (in BCD form) just prior to the occurrence of the corresponding sector.

A sector is divided into 100 8-bit bytes, a parity, and a space byte. The transfer rate is 3×10^5 bytes per second or 2.4×10^6 bits per second. The effective data rate is 2×10^5 words per second or 1 word per $5 \mu\text{s}$. The rotation time is 34.4 ms.

Some important control signals from the disk are:

- INDEX pulse, occurring once per revolution at the end of a track,
- SACF pulse, given just prior to reading a sector number when reading the address track, and
- FCLF pulse, given when a byte is available /requested depending upon READ/WRITE.

PDP-8 Constraints

The core memory of the PDP-8 has a cycle time of $1.5 \mu\text{s}$ and is logically divided into pages (1 page = 128 12-bit words). There are from one to eight core banks (called fields) of 4096 words.

The data break facility permits a word to be transmitted directly to memory. Two modes, 1-cycle and 3-cycle, control the word transfer and take one or three memory cycles using external or internal word count and address control registers, respectively. The maxi-

imum time delay between a break request and the start of the break cycle is 6 μ s.

TSS/8 Performance Constraints

Overhead due to disk activities should be minimal in order to allow simultaneous processing. This calls for a 1-cycle data break mode. Especially in a nonmulti-programmed system, the swap time of programs should be short (i.e., let the disk transfer at maximum transfer rate).

When swapping, complete fields are transferred. It is desirable to know the disk position so that transfers can be planned.

For simple file handling control in the software, it is desirable that the file space is binary and continuous. It is also important that the binary-decimal conversion not be done in software because of the frequency of use and the conversion time (see number convertor). Although the address structure of the disk could be changed, cost and delivery time would have increased and it would have become a nonstandard disk from a replacement (substitution) viewpoint.

CONTROL UNIT DESIGN

The control unit is shown in Fig. 1. It consists of the following main parts: data flow and parity generation/checking and control, core addressing and byte counting, disk addressing (number conversion), and a command decoder and general control taking care of the proper state transitions.

Layout of a Disk Track

For the storage of the information contained in a page (called segment), two sectors are required ($1/2$ segment = $1/2 \times 128 \times 12 = 768$ bits; 1 sector = $1 \times 100 \times 8 = 800$ bits). An important use of the disk is as a swapping device in which case 32 segments (a whole field) are transferred to and from the disk.

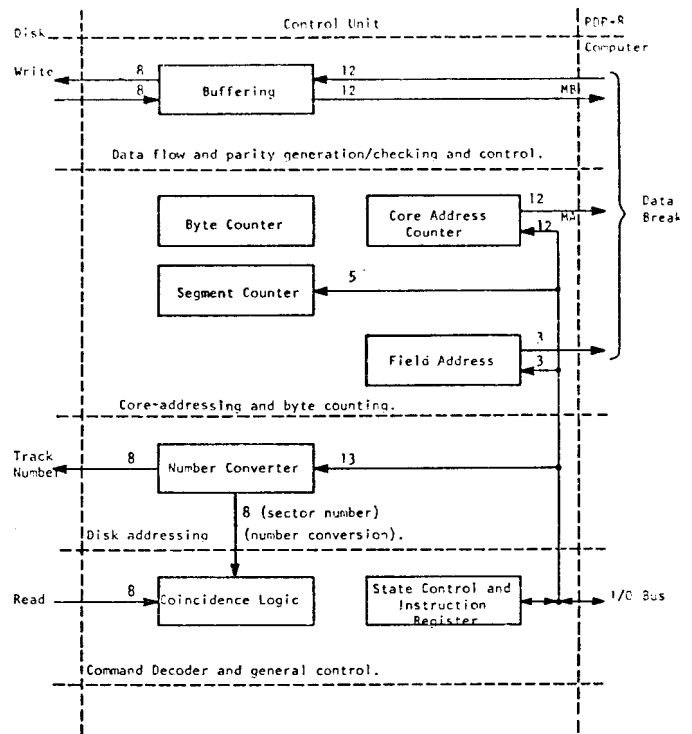
An empty or unused area (gap) at the end of a transfer is desirable in order to set up the conditions before the next transfer starts. This requirement suggested dividing a track into three areas of 16 segments according to Fig. 2.

Addressing Information on the Disk

Using the layout of Fig. 2 gives the disk a storage capacity of 48 segments per track, or $100 \times 40 = 4800$ segments per disk side.

Addressing a segment on the disk takes (except for the 1-bit disk side) a 13-bit binary number. A convention was made to determine the disk side bit and the thirteenth bit, such that only a 12-bit segment pointer had to be stored. The convention led to the disk address space layout of Fig. 3.

The disk, requiring a BCD track and sector number, necessitates process conversion hardware to convert the



Note: Only data-flow lines are shown. Numbers above lines show width of data-paths.

Fig. 1. Block diagram of control unit.





Fig. 2. Layout of a disk track.



Area	Range	Use
1	0 - 4095	User area for users with even user number
2	4096 - 4799	System area.
3	4800 - 8191	Nonexisting disk space.
4	8192 - 12287	User area for users with odd user number.
5	12288 - 12991	System area
6	12992 - 16383	Nonexisting disk space.

Fig. 3. Layout of disk-address space.

13-bit binary starting segment number. This process consists of a division by 48; the quotient is the track number in binary. The remainder of the division has to be multiplied by two and corrected for gaps in order to give the starting sector number in binary. After this the sector and track numbers have to be converted to BCD.

This process could be done with a PDP-8 program. Such a program occupies 88 core locations and has an execution time ranging from 417 to 456 μ s. Considering its purpose, this program would have to be part of the resident portion of the monitor. Because of its size, execution time, and frequency of use, the hardware alternative was selected, especially because most of its registers were needed anyway. Fig. 4 shows the flow chart of the number convertor as implemented. Fig. 5 shows its organization at different process stages. The conversion process is done in 13 "clock" times under control of the counter C (see Fig. 4). The main states are:

- 1) C=0 no operation (idling)
- 2) C=1 to 7 serial division by 48
- 3) C=8 correction because of gaps
- 4) C=9 to 12 conversion binary to BCD
- 5) C=13 done, ready to determine coincidence of the disk position and the desired starting segment.

Because of the presence of gaps (during which no information should be transferred), the position of the disk has to be known while transferring.

At the end of a track when the requested number of segments has not yet been transferred, it is desirable to continue on the next higher track such that for the PDP-8 the disk looks like two (one for each side) contiguous strings of segments.

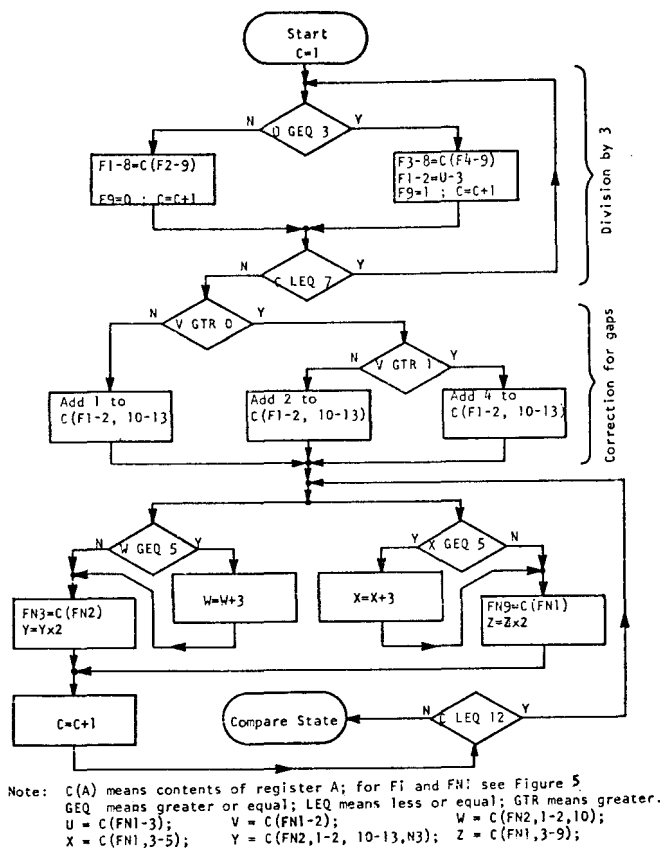
By making the registers containing the track and sector numbers to behave as BCD counters (after address coincidence has taken place), the two objectives mentioned above are realized.

Note the special way the binary to BCD conversion is accomplished. This method is based on the fact that a number $GEQ^2 5$ ($N=5+m$ where $m=0, \dots, 10$) after a right rotate (see Figs. 4 and 5) will be $GEQ 10$ (i.e., $2N=10+2m$), which would not be a BCD number. By adding +3 before rotating when $N \text{ } GEQ 5$, the following result is obtained: $2N=2 \times (5+m+3)=2m+16$. The 16 will not be visible in the BCD digit; it increments the next higher BCD digit position with 1.

Example:

1	2	4	8	16	32	64	128 ⁴			
27 ₁₀ ³	=	1	1	0	1	1	0	0	0	LSS ⁵ 5 rotate
		0	1	1	0	1	1	0	0	LSS 5 rotate
		0	0	1	1	0	1	1	0	GEQ 5
				1	1	0	0			add 3
				1	0	0	1			shift
		1	0	0	1	1	0	0		LSS 5 rotate
		0	1	0	0	1	1	1	0	result in BCD
		1	2	4	8	1	2	4	8 ⁴	
		2				7				

³ GEQ is greater or equal.
⁴ The tag 10 indicates the base of the number system.
⁵ Numbers indicate weight of the corresponding binary digits.
⁶ LSS is less.



Note: C(A) means contents of register A; for F1 and FN1 see Figure 5. GEQ means greater or equal; LEQ means less or equal; GTR means greater. U = C(FN1-3); V = C(FN1-2); W = C(FN2,1-2,10); X = C(FN1,3-5); Y = C(FN2,1-2, 10-13,N3); Z = C(FN1,3-9);

Fig. 4. Flow chart of number converter.

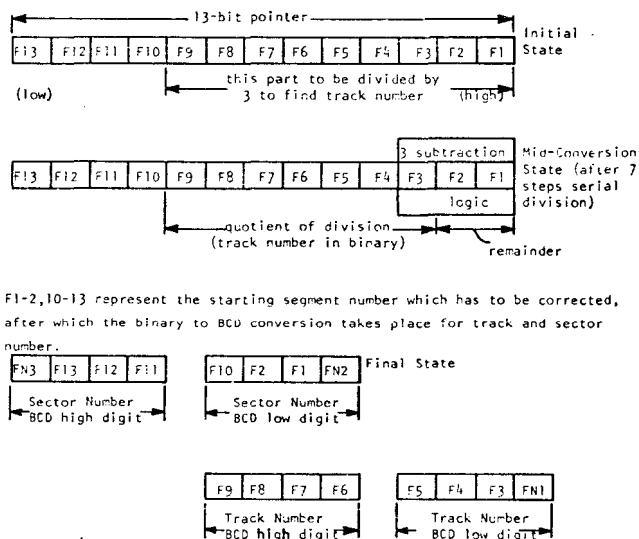


Fig. 5. Flip-flop configuration of the number converter at various stages (states) in the conversion process.

Buffering

Because of the disks' high data rate (1 word per 5 μ s), the delay in granting a data-break request (max. 6 μ s + 1.5 μ s), and the delay in the control unit, more than two words of buffering are needed. The buffering is also used for the conversion between words and bytes. Because of the difference in internal buffering in the disk hardware, an extra stage of buffering is required when writing on the disk (see Fig. 6). The numbers in the registers indicate their size, and the numbers on the

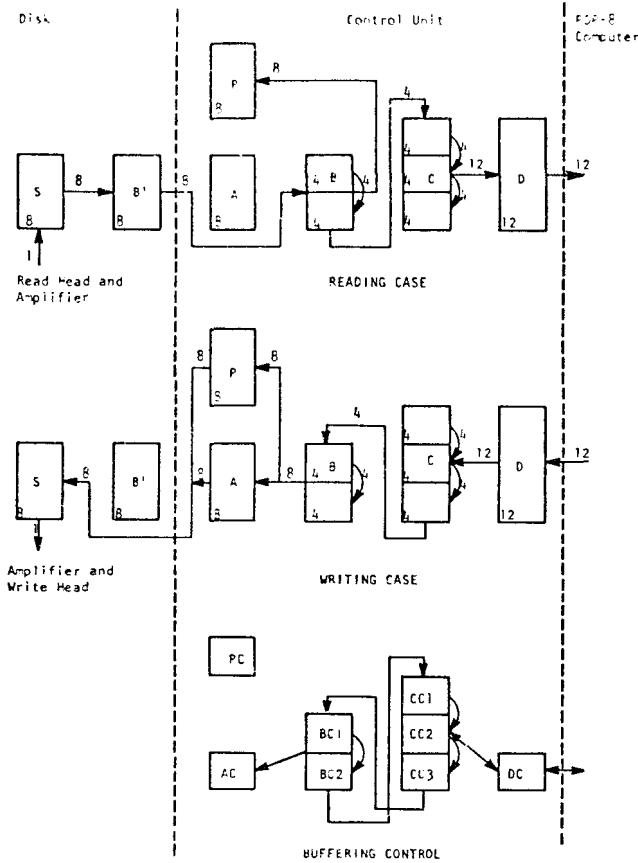


Fig. 6. Buffering, parity generation, and control.

lines indicate the width of the transfer paths. Registers B and C are divided into subregisters in order to allow for word-length conversion.

The control is made to simulate the behavior of information flow in the registers. It has one flip-flop for each register or subregister, and operates completely asynchronous from both the disk and the PDP-8 under control of its own 4-MHz clock. Control signals from the PDP-8 and disk are not sampled but used to trigger flip-flops in the control unit, thus eliminating such requirements as specific sampling rates and minimum and maximum durations of these control signals.

The design of the buffering is such that, in principle, arbitrary long waiting times can be handled by adding more registers.

Read Current Disk Position

A small change in the buffering structure of Fig. 6 allowed for the transfer of the position of the disk, read from the address track, to the PDP-8 core memory.

Parity Generation/Checking

The guaranteed error rates of the disk are: for recoverable errors better than 1 in 10^{10} ; for nonrecoverable errors better than 1 in 10^{12} .

The maximum number of bits transferrable in one hour is $60 \times 60 \times 2.4 \times 10^6 = 0.85 \times 10^{10}$, which means that the number of recoverable errors is better than one per hour and nonrecoverable errors better than one per 100 hours. For these reasons, longitudinal parity

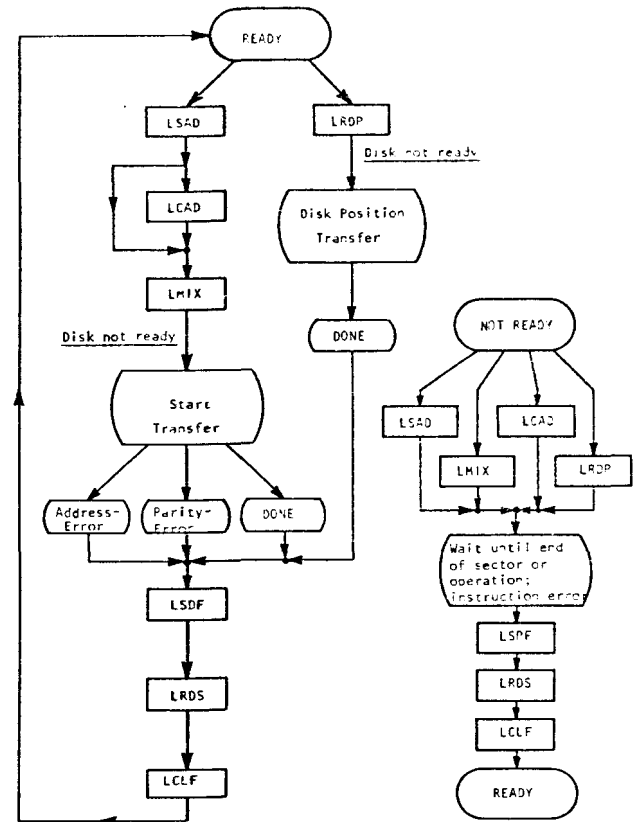


Fig. 7. Flow chart of IOT-command sequences.

(written as the 97th byte of every sector) was considered to be sufficient.

The P register of Fig. 6 is part of the parity generation/checking logic.

Instructions to the Disk

PDP-8 disk communications are done through input-output (IOT) commands⁶ and interrupts.

The IOT commands implemented are the following.

- 1) LSAD clear interface registers and transfer 12 bits of the starting segment address.
- 2) LRDP clear interface registers and transfer core-memory address (of field) after which the disk position is transferred to the specified location in core.
- 3) LSPF skip on disk flags (done + error flags).
- 4) LRDS read disk status. Done + error flags are read into the accumulator of the PDP-8.
- 5) LCLF clear done and error flags.
- 6) LCAD transfer starting core address.
- 7) LMIX transfer several bits (requested number of segments to be transferred, field number, READ/WRITE, disk side, thirteenth bit of starting segment address). After this command the actual transfer starts.

Interrupts are given when the requested task is done or when any of the three error conditions occurs. (These are: parity error, address error, and instruction error).

Fig. 7 shows the possible sequences of commands.

⁶ See PDP-8 Users Handbook.

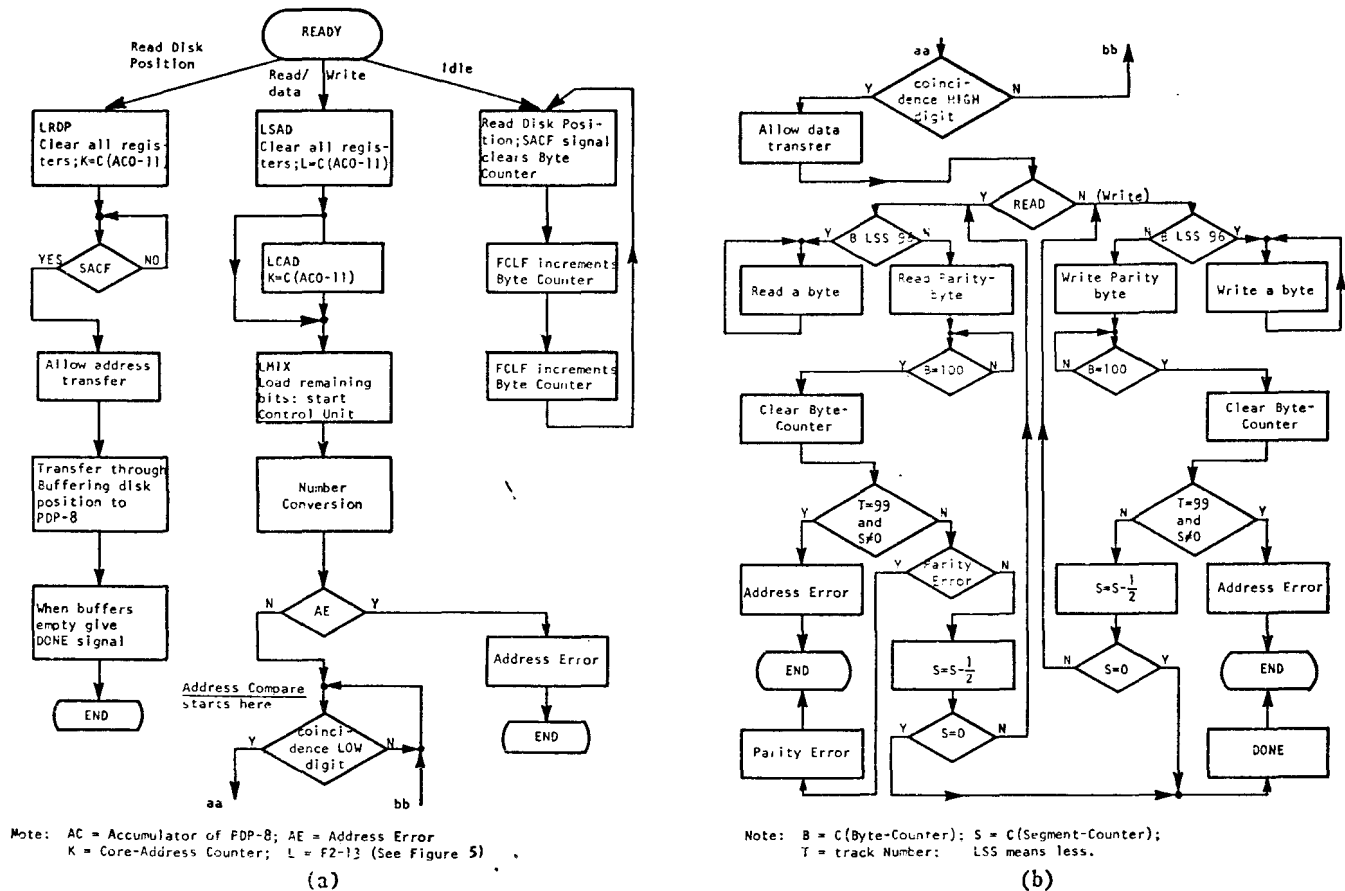


Fig. 8. (a) Flow chart of complete control unit. (b) Flow chart of complete control unit.

Instruction and Address Error Detection

Through the addition of some extra logic, the execution of commands disturbing the current execution when the disk is active was prevented and resulted in an instruction error (see Fig. 7). Addresses out of range result in an address error.

MODES OF OPERATION

A flow chart of the complete interface is given in Fig. 8. The basic three modes of operation are very apparent:

- 1) idling mode,
- 2) READ disk position mode,
- 3) READ/WRITE data mode.

CONCLUSIONS

The state of the art of control unit design, though, is well developed in practice. (Little has ever been directly reported.) This article is partially written to show the design protocol. The most interesting part is the buffering which, in principle, could have been generalized to handle arbitrary long service delays.