Hardware Memo 2

J. Holloway

PDP-1Ø Paging Device

2/20/70

1.0 PAGING

The MAC-1Ø Paging Device is a hardware device inserted between the PDP-1Ø processor and its memory buss. It maps address data from the processor (virtual address) into an address sent to the memory system (absolute address). Virtual and absolute memory space is divided into pages, each 1Ø24 words long. Virtual MA 18-25 are taken as the page number input to the paging box. The output consists of MA bits 17-25 mapped by the paging transformation, with MA 26-35 passed unmodified. Although the resulting address spans 512K of core, only up to 256K are addressable by one map.

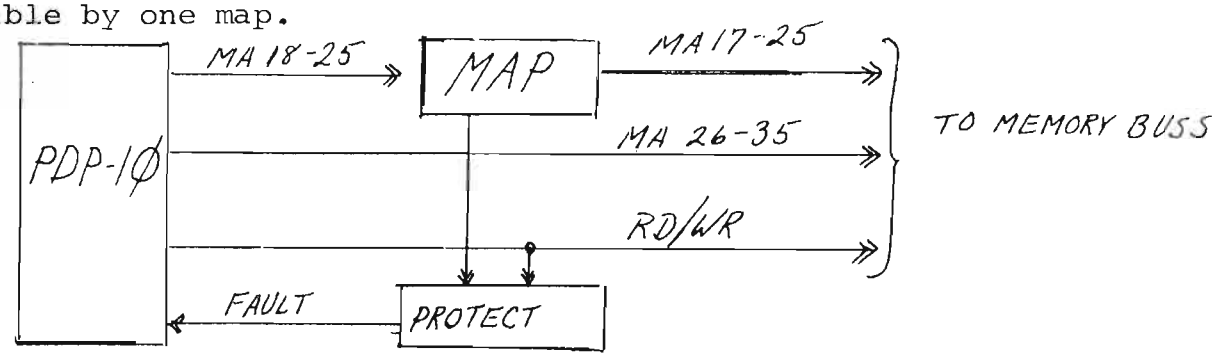

2.0 PAGE WORDS

The mapping information is stored in core memory in page tables. A page table consists of up to 128 18-bit entries called page words.

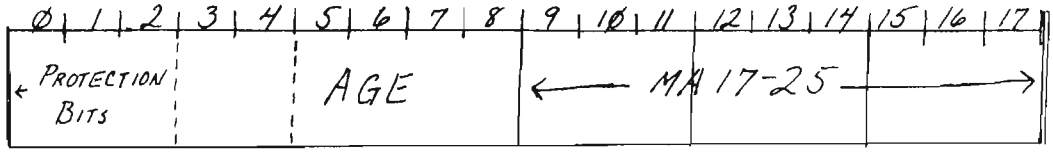

Ø-1 - Protection bits -

ØØ - Not Accessable

Ø1 - Read Only, executable by Pure code

1Ø - Read/Write/First

(Similar to Read Only, but not executable by Pure code)

2-4 - Unused

5-8 - Age - set from the Age register every time page word referenced.

9-17 - The absolute page number that this virtual page is mapped into.

2.1 DBR

The map contains three Descriptor Base Registers, which point to page tables. DBR1 maps the 128 pages in the lower half (Ø-377,777) of USER virtual memory. DBR2 maps the upper 128 pages 4ØØ,ØØØ-777,777) of USER memory. DBR3 maps the upper 128 pages of EXEC memory (4ØØ,ØØØ-777,777). EXEC addresses below 4ØØ,ØØØ are not mapped.

2.2 PAGE TABLES

The page table is addressed at its base by the corresponding DBR. Even page numbers are stored in the left halfword, in increasing addresses within the table. The maximum length of a page table is specified by a DBR length (DBRL) register associated with each DBR. The DBRL contains the number of full words in the page table (Ø-64).

2.3 ASSOCIATIVE REGISTERS

Each time an entry in a page table is referenced to determine the mapping and protection information, bits Ø, 1, 9-17 of the page

word are stored in an associative register (AR), keyed by the virtual page number plus a bit for USER versus EXEC pages. If another reference occurs to the same virtual page, the necessary information will be found in the AR, without reference to the page table in memory. The main group of associative registers. is called the A memory and consists of sixteen independent AR's. A four-bit counter (RBC) determines which AR is to be stored into when a new page table reference is made. The counter advances after each refill and acts like a ring buffer pointer.

2.3.1 B AND C MEMORIES

In addition to the A memory, two more groups of sixteen AR's, called the B and C memories exist. These are keyed by a single five-bit key for each of the sixteen AR's. The key selects USER/EXEC pages, and associates a sixteen-page block of virtual memory. A contiguous 16K of virtual memory is mapped into sixteen separate pages.

The effect of this is invisible to the mapping process. The addition of the B and C memories merely diverts page refills for a specific area of virtual memory. They might be used to shade the A memory from the addresses generated by the instruction fetches or reference to a group of data tables. Since the A, B, and C memories provide 48K of mapping information contained in hardware associative registers, only programs with unusually

large working sets need incur any overhead referencing the page tables.

2.4 AGING

The paging box contains a four-bit program loaded register. Whenever a reference is made to a page word for refilling an AR, bits 5-8 of the page word in core are updated to the current age. This is a method of enabling the executive to differentiate between recently referenced pages versus those that have fallen out of the working set.

3.0 LOAD PAGE MEMORY

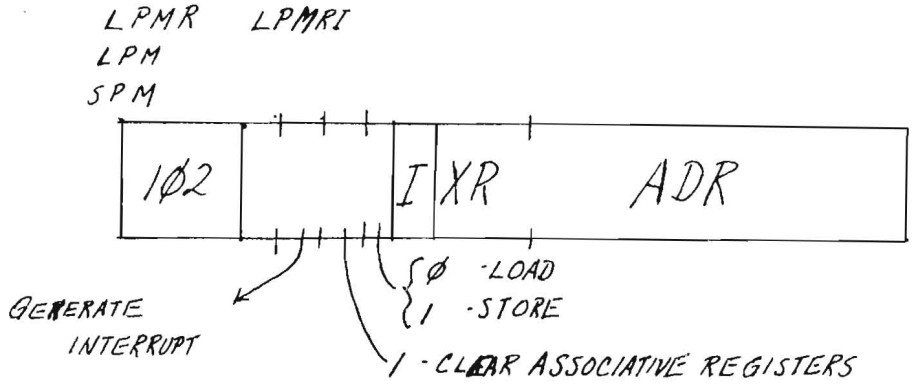

An instruction (OP code 12Ø,ØØØ,,) has been added to the PDP-1Ø. This instruction transfers eight words in a block pointed to by the effective address (absolute) to or from the paging device. Load Page Memory (LPM 1Ø2,ØØØ,,) transfers from memory to the paging device, and Store Page Memory (SPM 1Ø2,Ø4Ø,,) writes back in memory the same eight words of state information.

3.1 STATE BLOCK

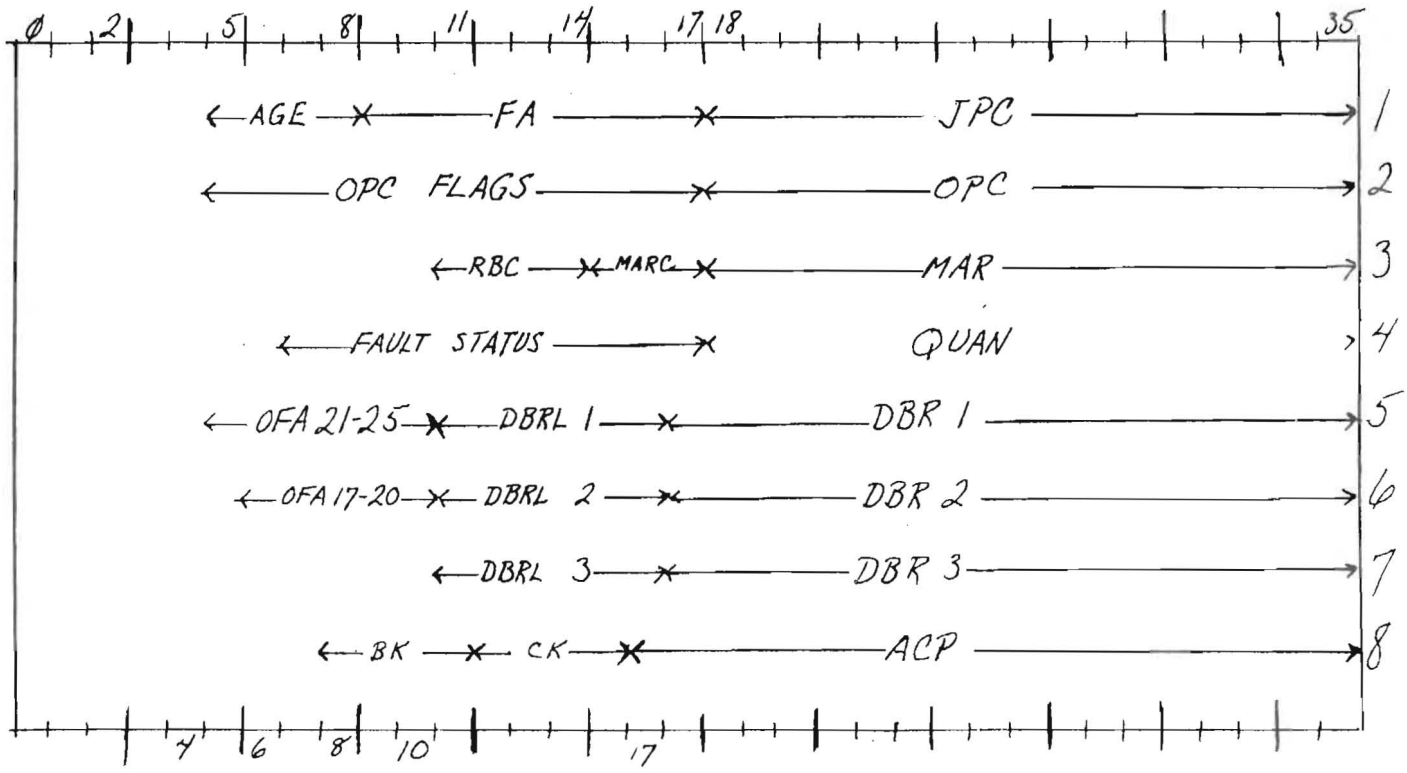

WD1

5-8 - Age - quantity stored in Age field (5-8, 23-26) of page word when referenced to refill AR.

9-17 - FAULT ADDRESS - virtual address referenced at time of page fault. Bit 9 is a "one" if reference was USER address, 10-17 correspond to MA 18-25.

WD2

5-17 - PC Flags - saved at the beginning of every instruction (except when Priority Interrupt is in progress).

5 - AR OVF (bit Ø of PC word)

6 - AR CRYØ (bit 1)

7 - AR CRY1 (2)

8 - AR FP OVF (3)

9 - BYTE interrupt (4)

10 - EX USER (5)

11 - EX IOT USER (6)

12 - PURE (7) Indicates instruction fetches must be made from Read Only pages.

13 - not used

14 - One Proceed (9) indicates that an interrupt will occur at the end of this instruction execution.

15 - not used.

16 - FXU (11)

17 - DCK (12)

18-35 - OLD PC - The contents of the Program Counter saved at the beginning of all instructions (except PI in progress). Used to restore state of machine after a page fault.

WD3

11-14 - Ring Buffer Couter - register that points to next AR in A memory to be refilled. Not reloaded by LPM, but reset to Ø.

15 - EXEC - if a "one" then the MAR address is a USER address

16-17 - MAR condition

ØØ - Never compare

Ø1 - Only the instruction fetches

1Ø - Write

11 - Any reference

18-35 - Memory Address Register (MAR) - a register that points to a location which will generate an interrupt when referenced (if the condition is correct). Does not function for fast memory references that occur through AC field or Index register field.

WD4

7-17 - Status bits

7 - (2ØØØ) One Proceed - set after instruction execution with One Proceed bit on in PC (bit 9). Does not set if instruction interrupted. Set after UUO even if transfers to EXEC mode.

8 - (1ØØØ) Fault - set if any page fault condition occurred. Not set if interrupt on fault inhibited (XCTR).

9 - (4ØØ) PG NXM - set if page table references or LPM, SPM reference caused non-ex-mem.

1Ø - (2ØØ) DBL - descriptor base length exceeded.

11 - (1ØØ) - Write attempted into Read Only page.

12 - (4Ø) - Write into Read/Write/First.

13 - (2Ø) - PURE code attempted instruction fetch from impure page.

14 - (1Ø) - No access

15 - (4) - EXR - enable mapping for EXEC addresses above 4ØØ,ØØØ.

16 - (2) MAR - MAR interrupt occurred.

17 - (1) QUAN - quantum overflow. This bit inhibits counting of the quantum timer. Carried into by QT18.

18-35 - Quantum Counter that counts except when machine is PI in progress. 1 MHZ.

WD5

5-9 - Bits 21-25 of Output Fault Address (OFA). The mapped address of the reference that faulted. Corresponds to FA 9-17.

1Ø-16 - DBRL1 - Length of first DBR in words. A zero indicates no entries.

17-35 - DBR1 - pointer to USER Ø-4ØØ,ØØØ page table. Absolute pointer - full 19-bit address.

WD6

6-9 - Bits 17-2Ø of OFA.

1Ø-16 - DBRL2

17-35 - DBR2 - USER 4ØØ,ØØØ-777,777

WD7

1Ø-16 - DBRL3

17-35 - DBR3 - EXEC 4ØØ,ØØØ-777,777.

WD8

7-11 - B key selects 16K block associated with B memory.

7 - "one" - USER address

8-11 - corresponds to MA 18-21

12-16 - C key

12 - "one" - USER addresses

13-16 - MA 18-21

17-35 - AC Pointer - contains address (absolute 19 bit) of block of 16 words that are to be referenced in place of fast memory when mapping addresses less than 2Ø with an XCTR instruction.

3.2 LPMR, LPMI

Bit 11 of the AC field of an LPM instruction causes all associative registers to be cleared (LPMR Load Page Memory and Reset). Bit 1Ø sets the Memory Protect Flag and causes and interrupt (LPMI Load Page Memory and Interrupt).

4.1 FAULT HANDLING

Page reference faults (such as NO ACCESS) abort the execution of the current instruction. In all cases the instruction is terminated before the contents of core or fast memory are affected. By restoring the program counter to point at the fault instruction, the state of the process before the fault is preserved. All faults (plus MAR, QUAN, One Proceed, and PAGE NXM) set the Memory Protect

Flag (bit 22 of CONI APR_1) and request an interrupt on the processor channel.

4.1.1 OLD PROGRAM COUNTER

The contents of the Program Counter (plus MISC Flags) are preserved in the OPC register at the beginning of every instruction not PI in progress. The flags are also loaded at IT1 to reflect the change in the Byte Increment Suppression Flag during ILDB, and IDPB. Thus, if the instruction that generates the fault transfers or skips before the fault is detected, the original PC is made available.

4.1.2 FAULT PAGE NUMBER

The eight-bit page number plus a bit that is a "one" for USER references, "zero" for EXEC references is saved when a fault occurs.

4.1.3 OUTPUT FAULT ADDRESS

MA bits 17-25 that were result of mapping an address which faulted (due to protection bits). Useful for testing A,B, and C memories.

4.1.4 PI TRAP MODE

The PDP-1Ø normally checks for interrupt request after the fetch for the instruction (or indirect word). After a fault it is

necessary to avoid further instruction fetches which may cause more faults or change the OPC. A special mode called PI TRAP mode is entered after the fault occurs. Interrupt requests are detected before the beginning of the instruction. This mode must stay in effect through interrupt on channels higher than the processor. When the interrupt routine recognizes the request and clears the Memory Protect Flag, PI TRAP mode is turned off.

5.1 EXECUTE RELOCATED

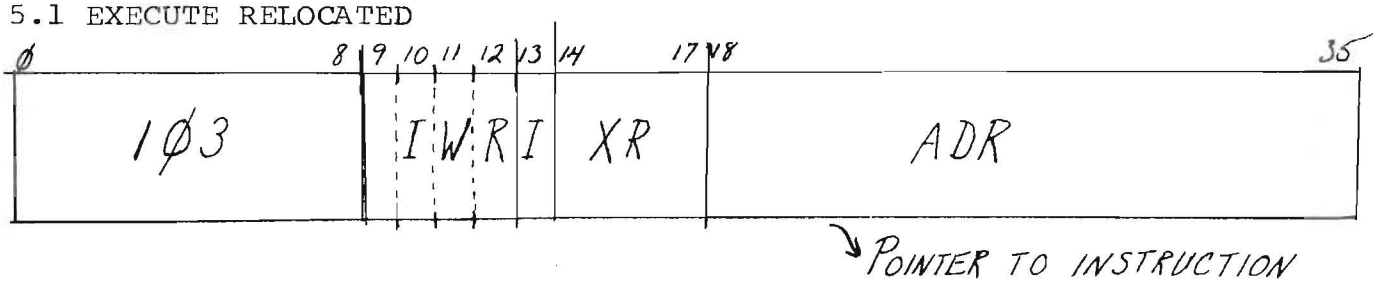

XCTR (OP code 103) is identical to XCT, except AC11, AC12 control the enabling of mapping on Write and Read references of the instruction executed. If AC12 is a "one" then all Read references after the effective address computation are mapped as if in USER mode (DBR1 and DBR2). If AC11 is a "one" then the Write references are mapped. References to the AC's by AC Field and Index Field are not mapped. Bit 10 of the AC Filed of the XCTR instruction inhibits the generation of a MEM PROT interrupt. Instead the next instruction is skipped. If traps are inhibited by this bit, the Fault bit in the Status word (bit 5) is not set, but bits 10-14 are set according to the type of fault.

5.1.2 AC POINTER

Locations below 2Ø which are relocated during an instruction being XCTR'd are not mapped to Fast memory or USER pages. Instead a 19-bit pointer in the State Word (ACP) is used to specify the location of a 2Ø word block that contains addresses Ø-17. (AC memory or absolute location Ø-17 depending on the Fast memory enable switch on the machine.).

6.1 MEMORY ADDRESS REGISTER (MAR)

All memory addresses fed to the paging device are compared against an eighteen bit register, the MAR. A three bit condition (bits 15-17 of MAR word) enable addresses equal to the MAR to set the MEM PROT Flag. Bit 15 is a "one" for USER mode addresses, a "zero" for EXEC addresses. Bits 16-17 are decoded: ØØ - never interrupt, Ø1 - only if address is referenced to fetch an instruction (not executed by SCT however), 1Ø - if location is being written into, 11 - under any condition. The MAR interrupt does not abort the instruction and doesn't set the FAULT bit (8) of the Status word. The PI TRAP mode is asserted, however, and the processor will interrupt before the next instruction and the OPC register will contain a pointer to the instruction that generated the interrupt. USER addresses referenced by means of the XCTR instruction will be compared against the MAR.

6.2 JUMP PROGRAM COUNTER (JPC)

In USER mode, any instruction that causes a transfer will save the contents of the Program Counter at the beginning of that instruction in the JPC register. This register is not updated when in EXEC mode.

6.3 QUANTUM COUNTER (QT)

An eighteen bit counter (WD4 bits 17-35) increments at 1MHZ rate except when the processor is PI in progress (or the PI system is turned off). When the counter sets bit 17 it stops and sets the MEM PROT Flag (PI TRAP mode is not turned on).

6.4 PURE PROCEDURE

Bit 7 of the PC word is the PURE bit. It is saved and restored along with the other bits in the left half of the PC. Trapping to EXEC mode (Interrupt or UUO) will clear this bit after it is stored. If this bit is on, then instruction fetches which address pages other than READ ONLY will fault. The PURE CODE bit (13) of the Status word is set. This bit can be turned on or off in USER mode by JRST 2.

6.5 ONE PROCEED

Bit 9 of the PC word is the One Proceed bit. When the current instruction increments the PC or transfers, then the One Proceed bit is cleared, and if it was set an interrupt is generated. MEM

PROT is set and Bit 7 of the PAGING STATUS word is set. PI TRAP mode is entered and an interrupt should occur before the next instruction. Instructions that are interrupted out of before advancing the PC will not clear this flag. A UUP that transfers to EXEC mode will interrupt after the JSR to the system UUO handlers. The system should simulate the interrupt after the execution of the call.

6.6 PAGE NON-EX-MEM

NXMs detected during the fetch of a page word will set the Page NXM bit (9) and the Processor NXM and MEM PROT flags. The instruction will continue with a zero fetched if a Read was requested. NXM's detected during an LPM or SPM instruction will set the Page NXM bit and the processor MEM PROT Flag.

6.7 PARITY ERRORS

Page table references and words referenced by LPM are checked for parity error by the processor. The error will set the processor Parity Error Flag. Currently, however, the parity stop switch will halt the machine only on the reference that caused the page table reference. LPM and SPM parity errors will not halt the processor.