

PDP-1 COMPUTER
ELECTRICAL ENGINEERING DEPARTMENT
M.I.T.
CAMBRIDGE, MASSACHUSETTS 02139

PDP-40

SCHEDULING DATA
in the PDP-1-X EXECUTIVE

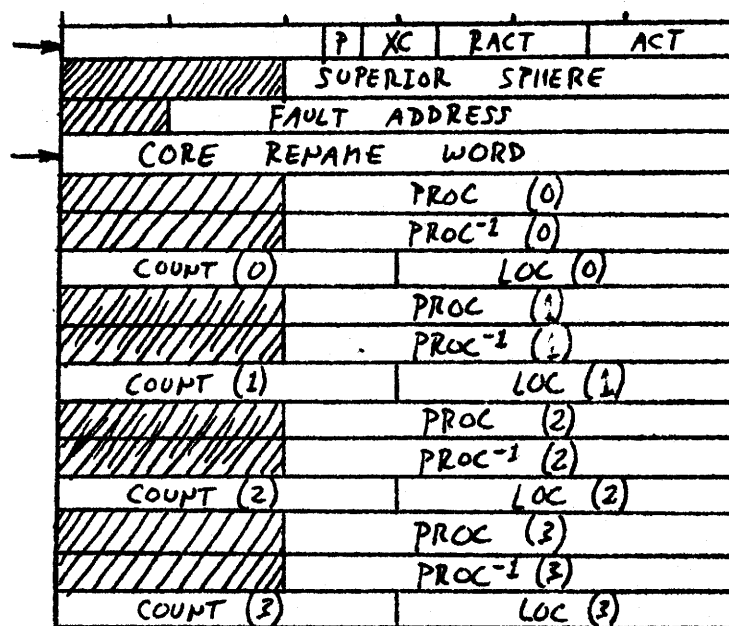
June 8, 1966

SCHEDULING DATA in the PDP-1-X EXECUTIVE

The scheduling data maintained by the PDP-1-X executive program consists mainly of process entities and computation entities. Other data objects used in scheduling are the i/o function started table, the queue head table, the queue population table, the target permit and target count tables, and the occupancy table; but these are covered adequately elsewhere (Project MAC Machine Structures Group Memo No. 19, the PDP Memo on Process Scheduling) and will only be mentioned in passing below.

Computations

A computation consists of a program image, a C-list (contained in the program image), and a set of processes. A computation being maintained by the PDP-1-X is represented by a computation entity in the executive. This is a 20₈-word block of storage with format as follows:



The first word in the computation entity is a status word. ACT is a 4-bit mask which tells which program image sections hold active processes. RACT is a similar mask which locates processes just reactivated after i/o. XC is the "expected cost" of restarting this computation, while P is a bit that tells when a computation has no C-list (program reference list). ACT, RACT, and XC are all manipulated by the executive, but P is determined when the computation is created, and is constant.

All computations (with maybe one exception) are created by some other computation, and this is the computation to which error conditions in this computation are referred. When a process in this computation runs into trouble, the process is suspended, a suspended process capability is created in the C-list of the computation at SUPERIOR SPHERE,

and a new process is started in this superior computation, at FAULT ADDRESS and with the index number of the suspended process capability as private data. SUPERIOR SPHERE and FAULT ADDRESS are determined when the computation is created, and are constant.

CORE RENAME WORD is the program image section reference translation map for all processes in this computation. (The word is divided into 6 fields of 3 bits each. A memory reference to core j is directed to the physical core named in field j (counting from the left) when core translation is active.) This word is part of the statewords of all processes in this computation. It is picked up on an unbreak, but is not deposited on a break since it cannot be changed by a user process.

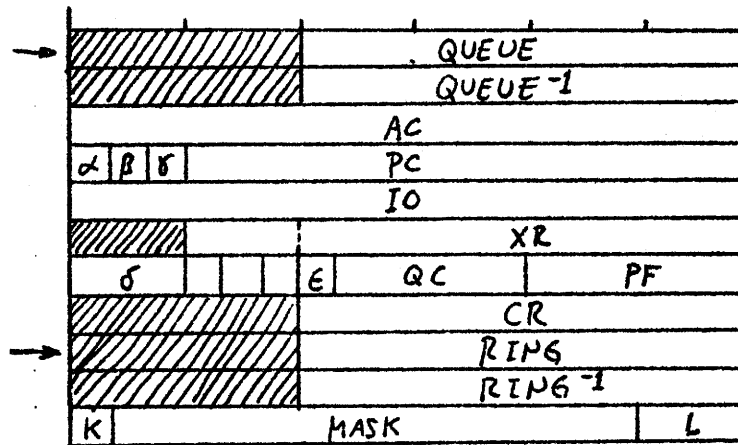
The following holds for $0 \leq j \leq 3$: PROC(j) and PROC⁻¹(j) are pointers to the head and tail, respectively, of the ring of processes running in program image section j. LOC(j) tells where program image section j is at any time: it may be a physical core name or a drum field name. COUNT(j) tells how many processes in program image section j are active.

When we start to run a new process, we remove that process from its PROC ring (and undex that ring's COUNT). When we get a PREEMPT or RND REN trap and have to run some other process, we return the running process to whichever PROC ring it belongs in at that time (and index that ring's COUNT). The PROC rings are maintained for the computation scheduler.

In the initial PDP-1-X executive, only the first six words of PROC and LOC information need be implemented.

Processes

A process consists of little more than a stateword, and is represented in the PDP-1-X executive by a process entity. This is a 13₈-word block of storage with format as follows:



QUEUE and QUEUE⁻¹ are pointers which tie processes to all sorts of things, most notably the run queue (embodied in the queue head table). RING and RING⁻¹ are pointers which tie the process to one of its computation's PROC rings.

The last word of the process entity is a status word. The bit K is on if the process is active. L is the level of the run queue which the process deserves to be on. (This is not maintained while the process is actually on the run queue.) The MASK will contain all sorts of interesting information when the process is inactive or suspended.

The stateword contains the following information:

- AC - accumulator
- PC - program counter
- IO - input/output register
- XR - index register
- QC - quantum counter
- PF - program flags
- CR - address of the core rename word
- α - overflow bit
- β - extend mode indicator
- γ - privileged mode indicator
- δ - addressing mode
- ϵ - program reference list indicator

Partitions of the Process Space

The following is a classification of the states a process might be in. This is intended to provide a useful partition of the set of all processes; the data in the PDP-1-X executive will reflect this partition.

A process can be:

I. Active

A. Runnable - on the run queue

B. Not runnable

1. because its program image section is not in core
2. because it does not belong to the computation whose quantum is currently in progress
3. (both of the above)