" MEMORY MANAGEMENT "

A VIDEOTAPE BY WENDELL HENRY

Extra Commentary by

Chris Moeller
HP Cupertino

HEWLETT hp PACKARD

MPE III
MEMORY RESOURCE
MANAGEMENT

MEMORY RESOURCE MANAGEMENT

OBJECTIVE

ATTEMPT TO MAXIMIZE THE PROBABILITY THAT A SEGMENT WILL BE IN REAL MEMORY WHEN IT IS NEEDED BY A PROCESS.

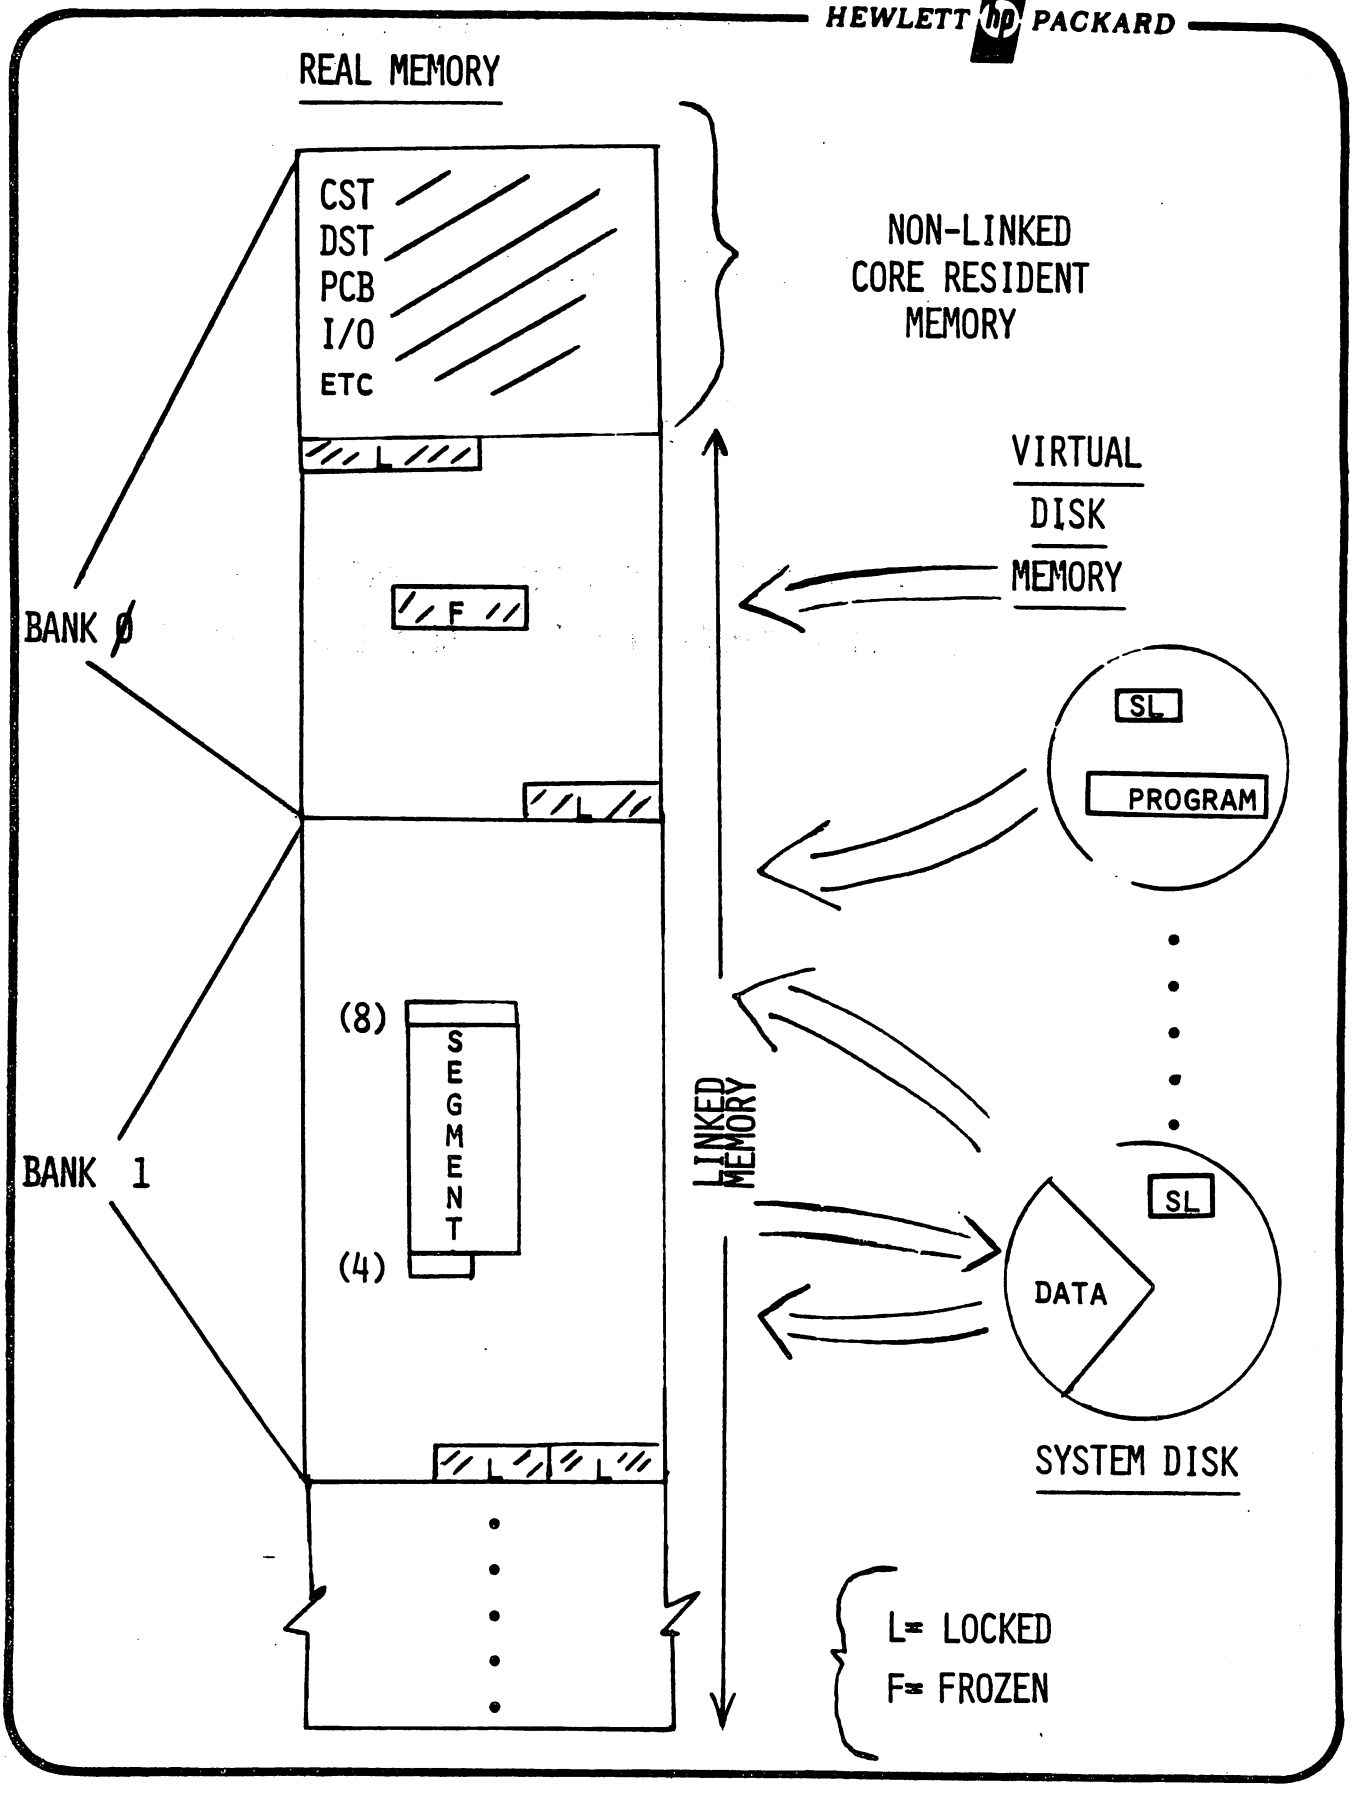
REAL MEMORY
CST
DST
PCB
I/O
ETC
NON-LINKED
CORE RESIDENT
MEMORY
VIRTUAL
DISK
MEMORY
F
BANK ø
SL
PROGRAM
(8)
SEGMENT
(4)
BANK 1
LINKED MEMORY
SL
DATA
SYSTEM DISK
L= LOCKED
F= FROZEN

CST/DST ENTRY

- CST/DST ENTRY DEFINES MEMORY STATE OF SEGMENT

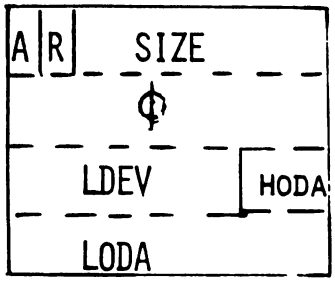

ABSENT SEGMENT

- A=1
- LOGICAL DEVICE # OF DISK
- HODA/LODA - DISK ADDRESS

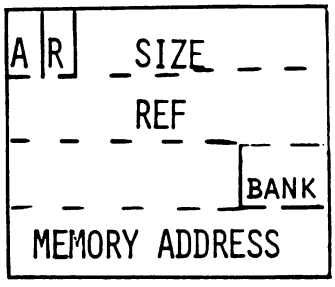

PRESENT SEGMENT

- A=0
- REF - REFERENCE COUNT
- BANK/ADDR - BANK AND MAIN MEMORY ADDRESS

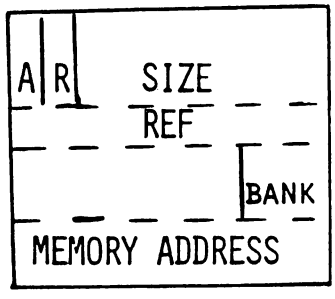

PSEUDO ABSENT SEGMENT

- A=1
- REF <> 0
- BANK/ADDR - BANK AND MAIN MEMORY ADDRESS

R - BIT DEFINES WHETHER SEGMENT HAS BEEN REFERENCED (PRESENT SEGMENT ONLY)

- SET BY MICROCODE DURING PCAL PROCESSING
- SET BY MIRCOCODE DURING MTDS/MFDS PROCESSING
- SET BY SOFTWARE DURING EXCHANGE DB

MEMORY STATES

- MEMORY IS MANAGED AS CODE AND DATA SEGMENTS
- AN ALLOCATED SEGMENT IS IN ONE OF THESE STATES

  — ABSENT — THE SEGMENT IS NOT IN MAIN MEMORY. IT IS IN VIRTUAL MEMORY (ON DISK)

  — PRESENT — THE SEGMENT IS IN MAIN MEMORY

    - AN UP-TO-DATE COPY OF A CODE SEGMENT ALSO EXISTS IN VIRTUAL MEMORY
    - AN OUT-OF-DATE COPY OF A DATA SEGMENT EXISTS IN VIRTUAL MEMORY

  — PSEUDO ABSENT — THE SEGMENT IS AVAILABLE FOR OVERLAY

MEMORY LINKS

- ALL REGIONS OF LINKED MAIN MEMORY ARE PRECEDED AND FOLLOWED BY CONTROL INFORMATION (MEMORY LINKS)
  - — MEMORY IS ALLOCATED AS MULTIPLES OF 4 WORDS
  - — THE LINK IN FRONT OF AN AREA IS 8 WORDS LONG
  - — THE LINE FOLLOWING AN AREA IS 4 WORDS LONG
- ALL FREE (UNALLOCATED) AREAS OF MAIN MEMORY ARE DOUBLY LINKED TOGETHER. THE MOST RECENTLY FREED REGION IS AT THE FRONT OF THE LIST
- ALL ALLOCATED AREAS OF MAIN MEMORY ARE POINTED TO BY THE CST/DST ENTRIES

PROCESS

- EXECUTABLE ENTITY
- CONTROL INFORMATION IN PCB
- PROPERTIES
  - — HAS WORKING SET
  - — HAS MEMEORY STATE
    - CORE RESIDENT
    - IN CORE
    - ABSENT
    - OUT OF MEMORY
  - — ON SCHEDULING QUEUE
    - READY LIST - NEEDS THE CPU
    - DISCARD LIST - WAIT FOR EVENT COMPLETION BEFORE NEEDING THE CPU
    - NO QUEUE

PROCESS

- BASIC EXECUTABLE ENTITY
- SEVERAL PROCESSES MAY BE RUNNING ON BEHALF OF ONE USER (JOB)

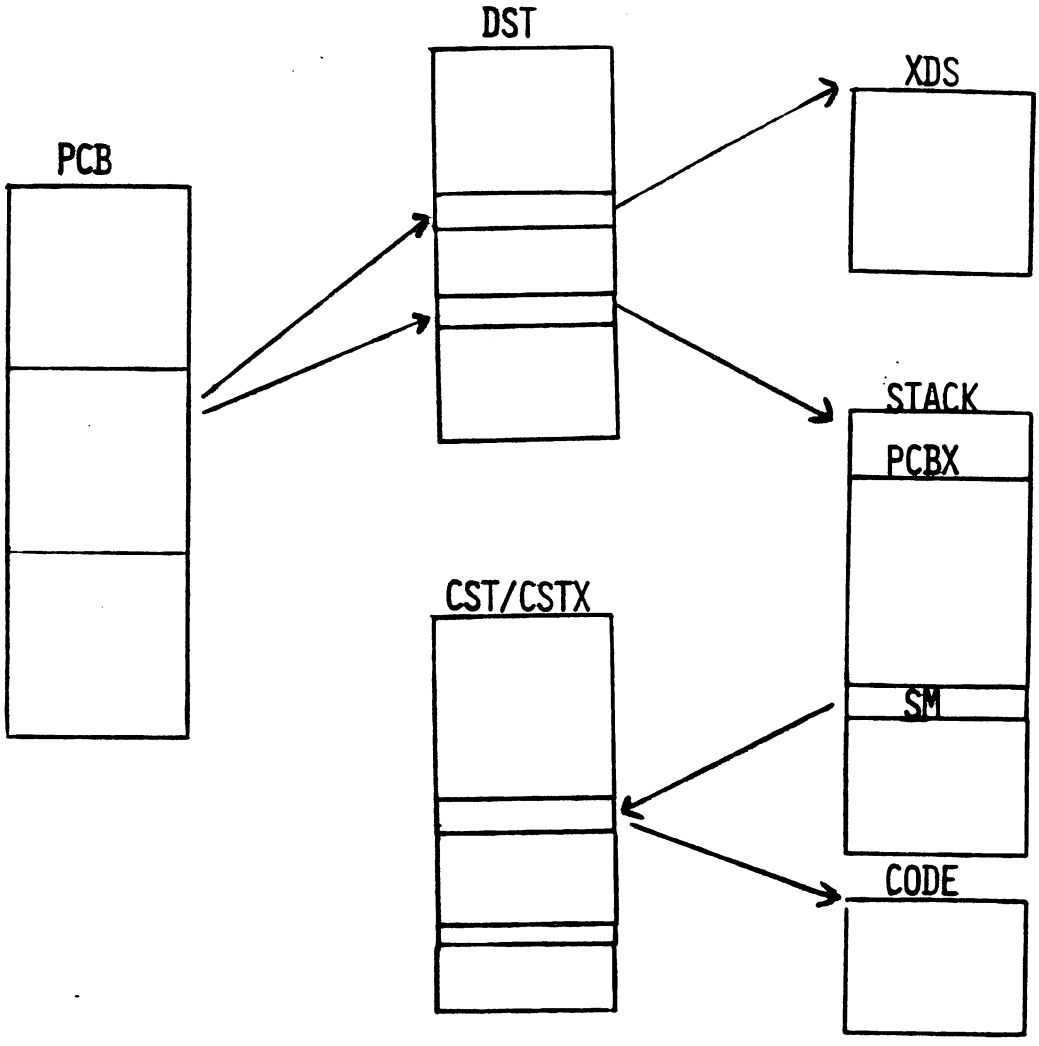

MEMORY ALLOCATION MANAGER (MAM)

- EXECUTES AS A PROCESS
- SCHEDULED BY DISPATCHER
  - — DISPATCHER HAS DETECTED THAT THE HIGHEST PRIORITY PROCESS IS ABSENT
  - — MAM WAS AWAKENED DIRECTLY BY ANOTHER PROCESS FOR HELP
- CONTROLS THE PLACEMENT/REPLACEMENT OF ALL SEGMENTS IN LINKED MEMORY

TABLES REQUIRED BY MAM

MTAB - CORE RESIDENT - CONTROL INFORMATION:

- FOR ALL CURRENT MAM I/O REQUESTS
- FOR ALL CURRENT ABSENCE REQUESTS

AVAIL'- CORE RESIDENT - THE CURRENT MEMORY AVAILABLE IN EACH BANK. AVAILABLE MEMORY IS THE SUM OF FREE AREAS AND AREAS ALLOCATED TO SEGMENTS WHICH ARE OVERLAY CANDIDATES

WSET - CORE RESIDENT - CONTAINS THE WORKING SETS OF ALL EXECUTING PROCESSES

CHANGE'- CORE RESIDENT - USED TO DETERMINE WHETHER AVAILABLE SPACE IN A BANK HAS INCREASED

ORDER'- BANK NUMBERS SORTED IN ORDER OF MOST AVAILABLE SPACE

WORKING SET

- A WORKING SET PER EXECUTING PROGRAM
- ALL PROCESSES EXECUTING THE SAME PROGRAM SHARE THE SAME WORKING SET - POINTED TO BY PCB
- ALL "PRESENT" SEGMENTS MUST BE MEMBERS OF AT LEAST ONE WORKING SET
- A "SYSTEM" WORKING SET IS MAINTAINED FOR ALL MPE CODE AND DATA SEGMENTS SINCE THESE SEGMENTS ARE SHARED BY ALL EXECUTING PROGRAMS
- EACH WORKING SET IS DYNAMIC IN SIZE
- WHEN A SEGMENT IS REMOVED FROM A WORKING SET IT IS DEALLOCATED OR MADE AN OVERLAY CANDIDATE

WORKING SET

- INDICATES WHICH CODE AND DATA SEGMENTS ARE USED BY THE EXECUTING PROGRAM MOST FREQUENTLY
- IF AN EXECUTING PROCESS IS FAULTING FREQUENTLY THE WORKING SET IS ALLOWED TO GROW
- IF AN EXECUTING PROCESS IS FAULTING INFREQUENTLY THE WORKING SET IS REDUCED

  — IF THE LAST FAULT WAS < 100ms AGO THE NEW SEGMENT IS ADDED TO THE WORKING SET

  — IF THE LAST FAULT WAS > 100ms AGO ALL SEGMENTS WHICH WERE NOT REFERENCED DURING THE LAST 100ms ARE REMOVED FROM THE WORKING SET. THE NEW SEGMENT IS ADDED.

WORKING SET PARAMETERS

ALFA - TIME SINCE LAST REFERENCE COUNTER UPDATE. STORED IN EACH WORKING SET. EACH TIME A PROCESS EXECUTES ITS CPU USAGE WILL BE ADDED TO ALFA FOR HIS WORKING SET.

TAU - A VALUE OF 100 ms. PRESENTLY CONSIDERED TO BE AN OPTIMUM INTER-FAULT TIME

LAT - TIME SINCE LAST ABSENCE TRAP FOR A PROCESS. STORED IN EACH PROCESS' PCB.

FREE SPACE COMPACTION

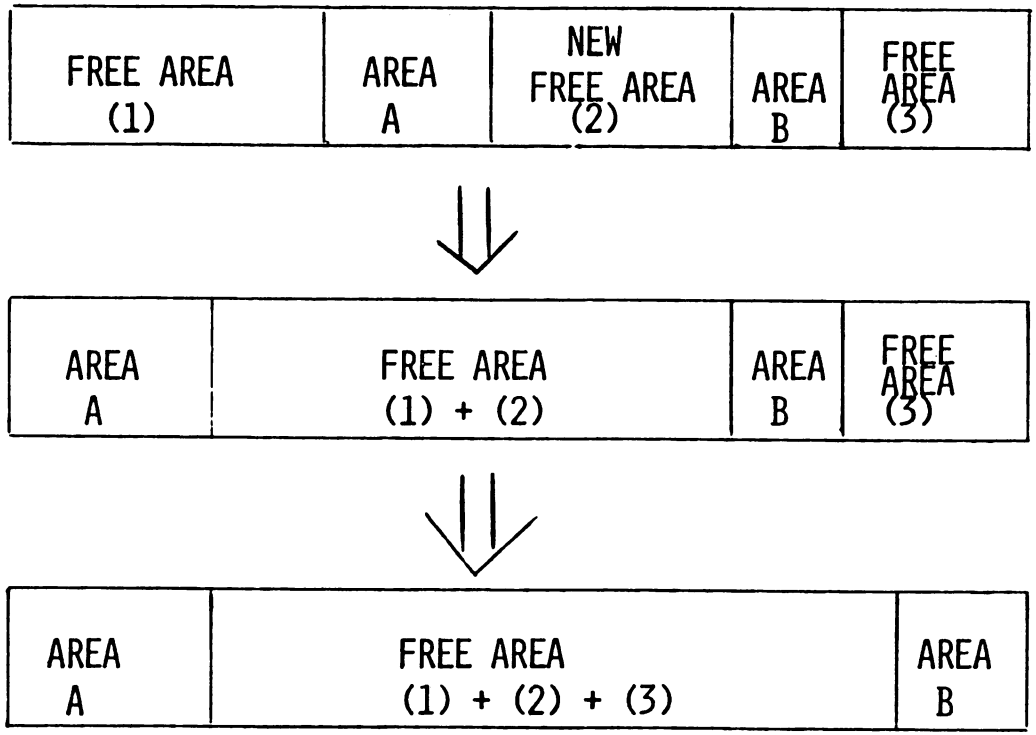

MEMORY MANAGEMENT ALGORITHM