

## Memory Management Part II

- Virtual memory
- Paging and page replacement
- Modeling page replacement algorithms
- Design issues for paging systems
- Implementation issues
- Segmentation

Operating Systems

1

## Virtual Memory (VM)

- Basic idea
- Paging
- Demand Paging and its performance
- Page Replacement and Page-  
Replacement Algorithms
- Allocation of Frames
- Thrashing problem
- Segmentation: pure segmentation,  
Segmentation with paging

Operating Systems

2

## Basic Idea of Virtual Memory

- Basic idea behind the virtual memory is that the combined size of programs, data, and stacks may easily exceed the amount of physical memory available.
- This is more true in multiprogramming environment.
- The operating system uses physical memory together with secondary storage to solve this problem.
- **Virtual Memory** seems to be a state of art universal method. Processes are allocated physical memory at a point they need based on the availability.

Operating Systems

3

## Paging Concept in Virtual Memory

- Logical memory is divided into blocks of same size called **pages**.
- Physical memory is divided into fixed-sized blocks called **frames** (size is power of 2, between 512 bytes and 8192 bytes).
- Pages are needed to be mapped to frames during the execution.
- To run a program of size  $n$  pages, need to use  $n$  frames.
- OS need to manage the page-to-frame mapping dynamically.

Operating Systems

4

## Page Faulting

- During the execution, a reference to virtual address is first checked if it has already been mapped to a physical address.
- If not it is said to cause a “page fault”.

Operating Systems

5

## Page Fault Algorithm-1

- If there is a reference to a page (every address reference is a reference to a page),
  - It will cause a trap to OS, which will in turn cause “page fault” processing.
- Page fault processing:
  - Find an empty frame
  - Swap page into that frame
  - Reset page’s validation bit.
  - Give control back to the process causing the trap (restart instruction)

Operating Systems

6

## Page Fault-2

- If there is no free page, OS starts page replacement procedure:
  - Page replacement: find some page in the memory, the one which is not really in use, to replace.
  - Conduct the swap operations.
  - For performance reasons the algorithm used should result in minimum number of page faults.
  - It is possible that the same page is brought into memory several times, during execution of a program...

Operating Systems

7

## Address Translation-1

- The virtual address generated during the execution is composed of two parts:
  - page number ( $p$ )
    - used as an index into a *page table* which contains **base address** of the page in physical memory (**frame**).
  - Page offset ( $d$ )
    - combined with base address to define the physical memory address to be referenced.

Operating Systems

8


## Address Translation-2

- The base address is the address of the frame to which the page is mapped.
  - if the virtual page is already in the memory, the address mapping is straight forward and very efficient as it is done in firmware.
- If the page is not in the memory, first a **page fault** occurs, after which the address mapping is as before.

Operating Systems

9

## Virtual address and MMU-Memory Management Unit




## The position and function of the MMU

Operating Systems

10

## Page-frame mapping: Example

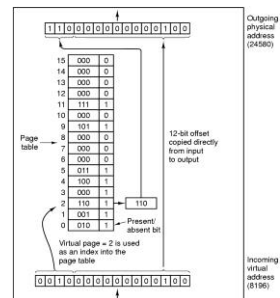
The relation between virtual addresses and physical memory address: page table and memory



Operating Systems

11

## MMU address mapping



## Internal operation of MMU with 16 4 KB pages

12

## Implementation of Page Table

- Page table is kept in main memory.
- *Page-table base register (PTBR)* points to the page table.
- *Page-table length register (PRLR)* indicates size of the page table.
- In this scheme every data/instruction access requires two memory accesses. One for the page table and one for the data/instruction.
- The two memory access problem can be solved by the use of a special fast-lookup hardware cache called *associative registers* or *translation look-aside buffers (TLBs)*

Operating Systems

13

## Associative Register

- Associative registers – parallel search

Page #	Frame #

- Address translation ( $A'$ ,  $A''$ )
  - If  $A'$  is in associative register, get frame #.
  - Otherwise get frame # from page table in the main memory

Operating Systems

14

## Associative Register

- *Valid-invalid* bit attached to each entry in the page table:
  - “valid” indicates that the associated page is in the process’ logical address space, and is thus a legal page.
  - “invalid” indicates that the page is not in the process’ logical address space.

Operating Systems

15

## Effective Access Time

- Associative Lookup =  $\epsilon$  time unit
- Assume memory cycle time is 1 microsecond
- Hit ratio – percentage of times that a page number is found in the associative memory =  $\alpha$
- Effective Access Time (EAT)
 
$$EAT = f(\text{PageTableAccessTime}, \text{MemoryAccessTime})$$


$$EAT = (1 + \epsilon)\alpha + (2 + \epsilon)(1 - \alpha)$$

$$= 2 + \epsilon - \alpha$$

Operating Systems

16

## Two-level page tables



- 32 bit address with 2 page table fields

Operating Systems

17

## Multilevel Paging and Performance

- Since each level is stored as a separate table in memory, mapping a logical address to a physical one may take four memory accesses.
- Even though time needed for one memory access is theoretically four times as much, caching permits performance to remain reasonable.

Operating Systems

18

### Multilevel Paging and Performance: Example


- Cache hit rate of 98 percent yields:  
 $\text{mem.access} = 100 \text{ nsec}$ ,  $\text{cache.access} = 20 \text{ nsec}$ ,  
 $\text{effective access time} = 0.98 \times (100 + 20) + 0.02 \times (400 + 20)$   
 $= 126 \text{ nanoseconds}$ .

which is only a 26 percent slowdown in memory access time.

Operating Systems

19

### Page Tables Entry Format



Typical page table entry

Operating Systems

20

### TLBs – Translation Lookaside Buffers

Valid	Virtual page	Modified	Protection	Page frame
1	140	1	RW	31
1	20	0	R X	38
1	130	1	RW	29
1	129	1	RW	62
1	19	0	R X	50
1	21	0	R X	45
1	860	1	RW	14
1	861	1	RW	75

A TLB to speed up paging

Operating Systems

21


### Inverted Page Table

- One entry for each real page of memory.
- Entry consists of the virtual address of the page stored in that real memory location, with information about the process that owns that page.
- Decreases memory needed to store each page table, but increases time needed to search the table when a page reference occurs.
- Use hash table to limit the search to one — or at most a few — page-table entries.

Operating Systems

22


### Inverted Page Table Architecture



Operating Systems

23

### Inverted Page Tables



Comparison of a traditional page table with an inverted page table

Operating Systems

24


## Shared Pages

- **Shared code**
  - One copy of read-only (reentrant) code shared among processes (i.e., text editors, compilers, window systems).
  - Shared code must appear in same location in the logical address space of all processes.
- **Private code and data**
  - Each process keeps a separate copy of the code and data.
  - The pages for the private code and data can appear anywhere in the logical address space.

Operating Systems

25

## Shared Pages Example: Editor is shared



Operating Systems

26

## DEMAND PAGING

Operating Systems

27

## Demand Paging

- It means bringing a page into memory only when it is needed.
  - Less I/O needed
  - Less memory needed
  - Faster response
  - More users
- Page is needed  $\Rightarrow$  reference to it
  - invalid reference  $\Rightarrow$  abort
  - not-in-memory  $\Rightarrow$  bring to memory
- Page replacement – find some page in memory, but not really in use, swap it out:
  - algorithm
  - performance – want an algorithm which will result in minimum number of page faults.
- Same page may be brought into memory several times.

Operating Systems

28

## Performance of Demand Paging

- **Page Fault Rate  $0 \leq p \leq 1.0$** 
  - if  $p = 0$  no page faults
  - if  $p = 1$ , every reference is a fault
- **Effective Access Time (EAT):**  

$$EAT = (1 - p) * \text{memory access time}$$

$$+ p * (\text{page fault overhead} + \text{swap page out}$$

$$+ \text{swap page in} + \text{restart overhead time})$$

Operating Systems

29

## Demand Paging: Example

- **Assumption: No associative memory**
  - Memory access time = 1 microsecond
  - Page fault rate  $p=50\%$
  - $q (=50\%)$  of the time the page that is being replaced has been modified and therefore needs to be swapped out.
  - Swap Page Time = 10 msec = 10,000 microsec
- **Computation of Effective Access Time:**  

$$EAT = (1 - p) * 2 + p (2 + q * 10000 + (1 - q) * 20000)$$

$$= (1 - 0.5) * 2 + 0.5 (2 + 0.5 * (10000) + (1 - 0.5) * 20000)$$

$$= 1 + 1 + 2500 + 5000$$

$$= 75002 \text{ micro sec}$$

Operating Systems

30

## Page Replacement Algorithms

- Page fault process:
  - Find the page to be removed or
  - make room for incoming page
- Modified page must first be saved before being overwritten
  - If unmodified can be overwritten
- Better not to choose an often or recently used page.
  - It will probably need to be brought back in soon!

Operating Systems

31

## Page Replacement Algorithms

- The real algorithms in use are generally combination of several approaches. The discussion of few of the ones listed below is more an academic one:
  - Optimal
  - Not recently used
  - FIFO
  - Second Chance
  - The Clock
  - Least Recently Used
  - Simulating LRU

Operating Systems

32

## Optimal Page Replacement Algorithms

- Replace the page needed at the farthest point in the future
- Normally, given a program, this is theoretically impossible to know. However it can be predicted by some means.
  - For example, logging page use on the previous runs of a process can be replayed for the future executions of the same program with same data...

Operating Systems

33

## Not Recently Used Page Replacement Algorithm: Implementation

- Each page has Reference bit, Modified bit
  - R and M bits are **set** when page is referenced and or modified
- They are classified according to the combination.
  1. not referenced, not modified
  2. not referenced, modified
  3. referenced, not modified
  4. referenced, modified
- NRU removes page at random from lowest numbered non empty class

Operating Systems

34


## FIFO Page Replacement Algorithm: Implementation

- Maintain a linked list of all pages
  - in order of their coming into the memory
- When required, the page at beginning of list is replaced.
- Disadvantage
  - page in the memory the longest time, may be the most often used one! Yet it will be replaced, causing extensive page fault...

Operating Systems

35

## Second Chance Page Replacement Algorithm: number indicate the last reference times

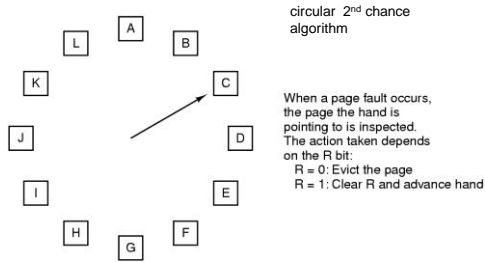


- Operation of a second chance
  - The candidate page to be replaced is the one at the tail of the list. The replacement depends on the reference bit.
  - The page referenced, has its R bit set to 1. If a tail page has its R bit set, it is reset and moved to the head.
  - If the tail page has its R bit 0, it is replaced, R bit is set to 1 and moved to the head.
  - At page fault at time 20, A has R bit is reset: it is treated as if it has just been brought in (numbers above pages are reference times), with R bit cleared.

Operating Systems

36

## The Clock Page Replacement Algorithm



Operating Systems

37

## Least Recently Used (LRU)

- It is based on the assumption that the pages used recently will be used again soon
- Write out the page that has been unused for the longest time
- Must keep a linked list of pages
  - most recently used at front, least at rear
  - update this list at every memory reference !!
- Alternatively keep counter in each page table entry
  - choose page with lowest value counter
  - periodically zero the counter

Operating Systems

38

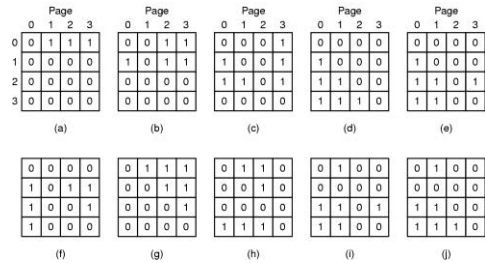
## Least Recently Used (LRU):bit matrix

- LRU maintains an nxn bit matrix in hardware
- On page reference, the corresponding row is set to 1, column to 0
- At any instant the lowest value row is the LRU

Operating Systems

39

## n x n LRU in hardware



LRU using a matrix – pages referenced in order 0,1,2,3,2,1,0,3,2,3

Operating Systems

40


## Simulating LRU in Software (1)

- Hardware may not exist for the architecture
- So a software solution is more practical
- Implementation of a counter that includes the effect of the aging
- Shift the counters right before the R bits are added on a reference
- This is repeated at every clock tick: each time reference bits are cleared, reference counters are shifted...

Operating Systems

41

## Simulating LRU in Software(2)



- The aging algorithm simulates LRU in software
- Note 6 pages for 5 clock ticks, (a) – (e)


Operating Systems

42

### The Working Set Page Replacement Algorithm (1)

- WS is the set of the pages that are currently in use
- WS page replacing algorithms make use of the locality of the reference
- If the entire WS of a process is in the memory, no page fault occurs until the process moves to a new locality.
- The pages will remain in a WS as long as their reference  $R$  bit is 1.
- Otherwise, the difference between timestamp and the current time is used to decide which one to push out of WS.

### The Working Set Page Replacement Algorithm (2)

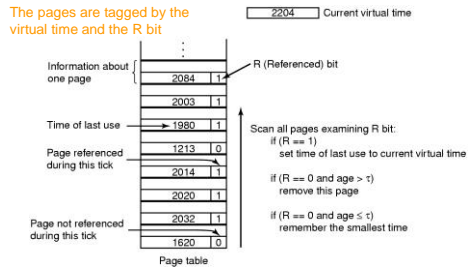


- The working set is the set of pages used by the  $k$  most recent memory references
- $w(k,t)$  is the size of the working set at time,  $t$

### WS based PRA

- Which page to be excluded from a WS, in case of a page fault:
- Remember that R and M bits are automatically set
- R bits are cleared periodically
- Each entry in the page table has a time stamp field and an R bit together with other info about the page...
- On page fault, if the R bit is set, the clock time is written to that entry, otherwise it is candidate for eviction, based on the difference between the time stamp and the current time...

### The Working Set Page Replacement Algorithm (3)



### The working set algorithm

### The WSClock Page Replacement Algorithm(1)

- Similar to clock replacement algorithm, the pages are arranged in a circular data structure with a pointer.
- The test is conducted from the pointer onward, which may remove the page from the WS or keep it as is, with R bit is set to 0.
- If it is referenced its time is updated.
- If R is 0 and it is not dirty and it is not within the time  $\tau$  evict it.  $\tau$  is the age of the page, relative to the current time (=currentTime-timestamp)...
- If it is dirty and outside the window it is put on hold for eviction...

### Review of Page Replacement Algorithms

Algorithm	Comment
Optimal	Not implementable, but useful as a benchmark
NRU (Not Recently Used)	Very crude
FIFO (First-In, First-Out)	Might throw out important pages
Second chance	Big improvement over FIFO
Clock	Realistic
LRU (Least Recently Used)	Excellent, but difficult to implement exactly
NFU (Not Frequently Used)	Fairly crude approximation to LRU
Aging	Efficient algorithm that approximates LRU well
Working set	Somewhat expensive to implement
WSClock	Good efficient algorithm





## Design Issues for Paging Systems

### Local versus Global Allocation Policies

- Local page replacement: the replacement is local to the process, i.e., one of the pages of the process causing page fault is replaced.
- Global page replacement: the oldest page in the system is replaced

Operating Systems

55

## Design Issues for Paging Systems

- Local vs Global allocation policies
- Thrashing is eliminated using load control, externally
- Page size can be optimized
- Separate instruction and data spaces
- Cleaning policy
- Sharing memory pages

Operating Systems

56

### Local versus Global Allocation Policies (1)

Page	Age
A0	10
A1	7
A2	5
A3	4
A4	6
A5	3
B0	9
B1	4
B2	6
B3	2
B4	5
B5	6
B6	12
C1	3
C2	5
C3	6

(a)

A0	
A1	
A2	
A3	
A4	
A5	CAB
B0	
B1	
B2	
B3	
B4	
B5	
B6	
C1	
C2	
C3	

(b)

A0	
A1	
A2	
A3	
A4	
A5	
B0	
B1	
B2	
B3	CAB
B4	
B5	
B6	
C1	
C2	
C3	


(c)

- a) Original configuration
- b) Local page replacement
- c) Global page replacement

Operating Systems

57

### Local versus Global Allocation Policies (2)



- Page fault rate as a function of the number of page frames assigned
- A: high page fault rate, B: Low page fault rate

Operating Systems

58

## Load Control

- Despite good designs, system may still thrash
- Solution :
  - Reduce number of processes competing for memory
  - swap one or more to disk, divide up pages they held
  - reconsider degree of multiprogramming

Operating Systems

59

## Advantages of small page size vs disadvantages

- Advantages
  - less internal fragmentation
  - better fit for various data structures, code sections
  - less unused program in memory
- Disadvantages
  - programs need many pages, larger page tables

Operating Systems

60

## Minimization of overhead: Optimum Page Size

- Overhead due to page table and internal fragmentation

$$overhead = \frac{s \cdot e}{p} + \frac{p}{2}$$

page table space
internal fragmentation

- Where

- s = average process size in bytes
- p = page size in bytes
- e = size of each entry in the PT in bytes


Optimized when

$$p = \sqrt{2se}$$

Operating Systems

61

## Separate Instruction and Data Spaces may be preferred




- Program and data address spaces together: one address space
- Program address space
- Data address spaces

Operating Systems

62

## Shared Pages



- Two processes may share the same program related page table with separate data and related page tables

Operating Systems

63

## Cleaning Policy: maintenance of an acceptable level of free frames

- Need for a background process, paging daemon
  - periodically inspects state of memory to maintain an acceptable level of available pages
- When too few frames are free
  - selects pages to evict using a replacement algorithm
- Implementation example:
  - use same clock type page handling: two handed clock, front hand does cleaning, back hand does replacement..

Operating Systems

64

## Further Reading

## Segmentation

Operating Systems

65

Operating Systems

66


## Segmentation vs. Linear memory

- So far virtual memory was linear, but frames were in ad-hoc positions
- For many problems two or more memory address spaces are more convenient.
- Compilers are good examples for this with many unrelated components such as text, symbol table, parse tree, etc.

Operating Systems

67

## Segmentation (1)




- With one-dimensional address space with growing tables, one table may bump into another

Operating Systems

68

## Segmentation (2)



Allows each table to grow or shrink, independently

Operating Systems

69

## Segmentation (3)


Consideration	Paging	Segmentation
Need the programmer be aware that this technique is being used?	No	Yes
How many linear address spaces are there?	1	Many
Can the total address space exceed the size of physical memory?	Yes	Yes
Can procedures and data be distinguished and separately protected?	No	Yes
Can tables whose size fluctuates be accommodated easily?	No	Yes
Is sharing of procedures between users facilitated?	No	Yes
Why was this technique invented?	To get a large linear address space without having to buy more physical memory	To allow programs and data to be broken up into logically independent address spaces and to aid sharing and protection

Comparison of paging and segmentation

Operating Systems

70

## Implementation of Pure Segmentation




- (a)-(d) Development of checkerboarding
- (e) Removal of the checkerboarding by compaction

Operating Systems

71

## Segmentation with Paging: MULTICS (1)




- Descriptor segment points to page tables
- Segment descriptor – numbers are field lengths

Operating Systems

72

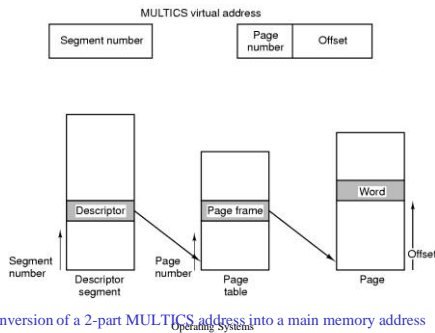
### Segmentation with Paging: MULTICS (2)

- Segment page table may not be in the memory
- Page referenced may not be in the memory
- TLB is used to keep the address of the most recently used pages



A 34-bit MULTICS virtual address

### Segmentation with Paging: MULTICS (3)




Conversion of a 2-part MULTICS address into a main memory address

### Segmentation with Paging: MULTICS (4)

Comparison field					Is this entry used?
Segment number	Virtual page	Page frame	Protection	Age	
4	1	7	Read/write	13	1
6	0	2	Read only	10	1
12	3	1	Read/write	2	1
2	1	0	Execute only	7	1
2	2	12	Execute only	9	1


- Simplified version of the MULTICS TLB
- Existence of 2 page sizes makes actual TLB more complicated

### Segmentation with Paging: Pentium (1)




A Pentium selector

### Segmentation with Paging: Pentium (2)



- Pentium code segment descriptor
- Data segments differ slightly

### Segmentation with Paging: Pentium (3)




Conversion of a (selector, offset) pair to a linear address

Operating Systems

79

### Segmentation with Paging: Pentium (4)




Mapping of a linear address onto a physical address

Operating Systems

80

### Segmentation with Paging: Pentium (5)



Protection on the Pentium

Operating Systems

81