


example: given LA which generates:

$$S = 2, d = 350$$

$$PA = B + d = 18,000 + 350 = 18,350$$

Segmentation with Paging


"Paging the Segments"



If Seg size = 64K.

→ Page size = 4K.


→ every Seg contains 16 page.



Chapter #9:

Virtual Memory Management


- How to run bigger programs on smaller memory.
- No need for the whole program to be loaded into memory, only part of it will do.
- We will discuss 'demand paging'



V/I (Valid/Invalid) Bits

1: in memory, 'Valid'

0: not in memory, 'Invalid'



- The LA is checked through the page table if it's a valid address, that's, if **v/i Bit = 1** (the page memory) that it continue execution as usual.
- What if the **v/i Bit = 0**, that's the page isn't in memory, then we say a page fault occurs.
- The OS looks for a free frame in memory, swap-in the required page from HD, it updates the page table & resumes execution.
- What if there's no free frame?
 - It selects a victim frame
 - maybe, it swap-out the page from memory to HD.
 - It swap in the required page from HD to memory
 - update the page table.
 - resumes execution.

Performance of demand pager:

Assume page fault rate (probability) = P

$$0 \leq P \leq 1$$

If $P=0$, All execution with no page fault. "perfect"

If $P=1$, Every execution of instruction, there's a page fault.

Assume memory access = M

$$EAT = (1-P) * M + P * \underbrace{\left[\text{Swap in Page} + \text{maybe Swap out page} + M \right]}_{\text{Page fault overhead}}$$

Example: Assume memory access (M) = 10 nsecs

Page fault rate (P)

Page transfer (swap) time = 10 msecs

40% of the time the page needs to be swap out. swap in

$$EAT = (1-P) * 10 + P * \left[10 * 1000 + 0.4 * 10 * 1000 + 10 \right]$$

$$EAT = 10 - 10P + 10,000P + 4,000P + 10P$$

$$= 10 + 14,000P \approx 14,000P$$

!! تحويل الوحدات الكبير لوحدات أصغر

mils, ~~no~~ mics $\rightarrow$ mils $\rightarrow$ mics

∴ Result:

Performance depends on P (page Fault Rate).

[-] Objective: Minimize the Page Fault Rate.

[!] Note: A new bit called "dirty Bit" is added to the page table to indicate if the page is modified.

1 = page is modified

0 = page isn't modified.

[-] Page Replacement Algorithms:

The OS must select the victim frame so that, to minimize the page fault rate.

1] FIFO

Replace the page which enters the memory First (oldest page in memory).

example: given the following page references.

1 2 3 4 1 2 5 1 2 3 4 5

1	1	1	4	4	4	5	✓	✓	5	5	✓
	2	2	2	1	1	1			3	3	
		3	3	3	2	2			2	4	

9 page faults

1 2 3 4 1 2 5 1 2 3 4 5

1	2	1	1	✓	✓	5	5	5	5	4	4
	2	2	2			2	1	1	1	1	5
		3	3			3	3	2	2	2	2
			4			4	4	4	3	3	3

10 page faults

2. Belady's Anomaly:

2. Optimal Replacement:

Replace the page which won't be used in the future for the longest period.

1 2 3 4 1 2 5 1 2 3 4 5

1	1	1	1	✓	✓	1	✓	✓	3	3	✓
	2	2	2			2			2	4	
		3	4			5			5	5	

7 Page faults

⚠ Major problem: How the OS knew in advance the next page in which execution occurs?!

* Optimal replacement is used as a benchmark

3 Least Recently Used [LRU]

Replace the page which have not been used in the past for the Longest period of time.

1 2 3 4 1 2 5 1 2 3 4 5

1	1	1	4	4	4	5	✓	✓	3	3	3
	2	2	2	1	1	1			1	4	4
		3	3	3	2	2			2	2	5

10 page faults "3 Frames"

1	1	1	1	✓	✓	1	✓	✓	1	1	5
	2	2	2			2			2	2	2
		3	3			5			5	4	4
			4			4			3	3	3

8 page faults "4 frames"


* Implementation:

(1) Add a Counter to the page table that contains the last time the page is referenced (used)

	time of reference
	10:3:20 → Integer
	10:3:11

page table

(2) Keep a queue of referenced pages



→ or we can use a stack, Every time we use a page, push it into stack.

↳ the stack top is the latest used page

LRU approximation:

Add one bit to the page table called reference bit

→ 1 ≡ The page is referenced (used),
Read or Written.

→ 0 ≡ Page isn't referenced.


	reference bit

Page table.

→ LRU approximation, Replace the page whose reference bit is 0.

Second chance:

- replace 1 0
- no change 0 1
- no change 0 1
- replace 1 0
- no change 0 1
- no change 0 1



Counter clock wise

Enhanced Second chance:

Use the reference bit & the dirty bit. (Reference bit, dirty bit)

- best**
- (0,0) = Page isn't referenced & not modified. (1)
 - (0,1) = Page isn't referenced but modified. (3)
 - (1,0) = Page is referenced but not modified. (2)
- worst**
- (1,1) = Page is referenced & modified. (4)

★ Counting Algorithms:

1. MFU (Most Frequently Used)

Replace the page which has been used for the maximum # of all times.

2. LFU (Least Frequently Used)

Replace the page which has been used for the minimum # of all times.

1 2 3 1 2 1 4 3 4 5

1	1	4	✓	✓	✓	4	MFU
	2	2				2	
		3				3	
						1	
						2	LFU
						4	

Global vs. Local Replacement

- (1) Global replacement: $(0, 0)$
- (2) Local replacement: $(1, 0)$
- (3) Global replacement: $(0, 1)$
- (4) Local replacement: $(1, 1)$

Global replacement

(0, 0) (1, 0) (0, 1) (1, 1)

Global replacement: $(0, 0)$

Local replacement: $(1, 0)$

Global replacement: $(0, 1)$

Local replacement: $(1, 1)$

Global replacement: $(0, 0)$

May 19, 2018

Allocation of Frames for processes :-

(1) Equal Allocation: Each process gets the same number of pages (frames).

example: memory 100 frames, 5 processes

— Every process gets $\frac{100}{5} = 20$ frames.

— unfair $\rightarrow$ not good performance.

(2) Proportional Allocation according to size:

example: Assume the size of process $P_i = S_i$

Assume we have m frames of memory

Process P_i gets $\frac{S_i}{\sum S_i} * m$

— Assume memory 100 frames, 3 processes with sizes 100, 400, 700 KB.

— $P_1 = \frac{100}{1200} * 100 \approx 8$ frames

— $P_2 = \frac{400}{1200} * 100 \approx 34$ frames

— $P_3 = \frac{700}{1200} * 100 \approx 58$ frames

(3) Proportionally according to priority:

example:

process	priority
---------	----------

P₁

2

P₂

3

P₃

7

→ highest priority

— If we have 100 frames of memory


$$P_1 = \frac{2}{12} * 100 = \frac{1}{6} * 100$$

$$P_2 = \frac{3}{12} * 100 = \frac{1}{4} * 100 = 25 \text{ Frames}$$

$$P_3 = \frac{7}{12} * 100$$

Thashing: The OS is busy swapping page in & out.

page faults



Allocated frames
for process.



Good performance.



Poor performance.



— Sometimes the number of allocated frames for the process is low, which leads to poor performance (low system utilization).

The OS thinks that the degree of multiprogramming is low,

→ OS increases the degree of multi programming


→ The system gets worse.

→ Most of the time the system OS is doing swapping of page.




27.7 (A)

The time to move from track 18 → 225 is called "seek time" mechanical move.



HD access time = Seek time + Latency time




→ we could improve the seek time by selecting a good disk scheduling algorithm. (minimize seek time)

example: Assume HD has 200 tracks (0 - 199)


given the queue of HD requests as follows:

98, 183, 37, 122, 14, 124, 65, 67.

Assume the R/W head is currently serving a job at track 53 & just finished serving a track job at track 40.




(1) FCFS



→ Average Head Movement =

$$\frac{(183 - 53) + (183 - 37) + (122 - 37) + (122 - 14) + (124 - 65) + (67 - 65)}{8}$$

(2) Shortest Seek Time First (SSTF)




AHM $\rightarrow$ 29 tracks / Job.

SSTF gives the optimum solution
minimum


⚠ Problem: Starvation

(3) SCAN (Elevator)




AHM $\rightarrow$ 41 tracks / Job.

(4) C-SCAN (Circular SCAN)






AHM $\rightarrow$ 47 tracks / Job.

(5) Look:



(7) C-Look:





Hard Disk

File Access:

(1) Sequential Access.

(2) Direct Access.

Location: Address of First block in the file (Constant, can't be changed).

File Pointer: The current location where the file is located.

> Sequential Access:

The ability to read the next block in the file.


example:

IF FP at block (7), then sequential access is the ability to read/write block (8).

> Direct Access:

The ability to read/write block # in the file.

(n is relative Address 'Block #')



Allocation Methods:

How the disk blocks are selected and allocated for the file.

Note: Every storage device, has 'device directory', which contains the information about the files stored on that device. Generally, it's a 'hash table'.

Device directory

Name	Location	FP	Size
Sam.txt	11	11	6
All files on the device			


(1) Contiguous Allocation:

The OS selects contiguous blocks of the file.

Device Directory-

Name	Location	Size	FP
test.dat	13	7	13

test.dat



Hard Disk

Advantages:

It supports both sequential & direct access easily.

→ Sequential Access:

If FP at address X ,

∴ The next block address is $X+1$.

→ Direct Access:

Assume we want to read block #5 (5 relative) ^{in the file}

The physical address of block #5 in the file =

$$\text{Location} + 5 - 1 = 17$$

∴ To access Block $\neq n$,

$$PA = \text{Location} + n - 1$$

!! Disadvantages:

(1) Fragmentation (holes).

→ Solution: Compaction (defrag).

(2) "Major Disadvantage"

— what if the file grows up.


— what if we delete a block.

— what if we insert a block.


(2) Linked Allocation:

— The File Blocks are selected as a linked list of blocks.

Name	Location	Size
test.dat	13	7



location



Advantages:

No Fragmentation

- File can grow, insert, delete as you like

Disadvantage: "Major"

- It only supports Sequential Access easily.


- Direct Access isn't supported & needs other tools such as an "index file".

13
4
14
15
6
1
7

test_index.dat

(3) Indexed Allocation:

- Each File has at least one index block which contains the addresses of the File data blocks.



Advantages:

It supports both Sequential & Direct Access.

→ Sequential Access:

IF FP at $IND[n]$;

next block address equals $IND[n+1]$;

→ Direct Access:

To access block # n in the file, its $PA = IND[n]$.

Disadvantage:

The wasted index block(s). each file needs at least one index Block.