

Chapter 8: Virtual-Memory Management



1

Chapter 8: Virtual-Memory Management

- Background
- Demand Paging
- Page Replacement
- Allocation of Frames

2

Objectives

- To describe the benefits of a virtual memory system
- To explain the concepts of demand paging, page-replacement algorithms, and allocation of page frames

3

Background

- Code needs to be in memory to execute, but entire program rarely used
 - Error code, unusual routines, large data structures
- Entire program code not needed at same time
- Consider ability to execute partially-loaded program
 - Program no longer constrained by limits of physical memory
 - Each program takes less memory while running -> more programs run at the same time
 - ▶ Increased CPU utilization and throughput with no increase in response time or turnaround time
 - Less I/O needed to load or swap programs into memory -> each user program runs faster

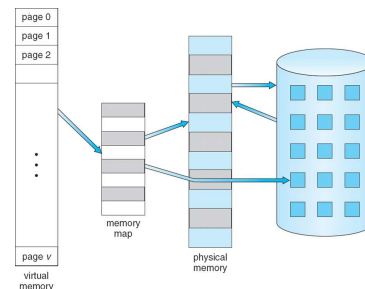
4

Background

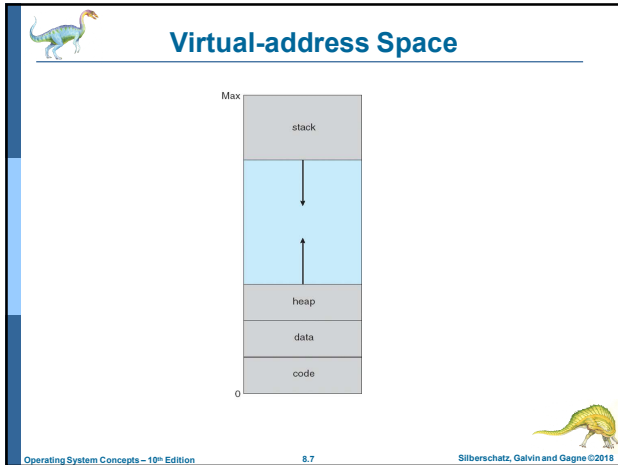
- **Virtual memory** – separation of user logical memory from physical memory.
 - Only part of the program needs to be in memory for execution
 - Logical address space can therefore be much larger than physical address space
 - Allows address spaces to be shared by several processes
 - Allows for more efficient process creation
 - More programs running concurrently
 - Less I/O needed to load or swap processes
- **Virtual address space** – logical view of how process is stored in memory
 - Usually start at address 0, contiguous addresses until end of space
 - Meanwhile, physical memory organized in page frames
 - MMU must map logical to physical
- Virtual memory can be implemented via:
 - Demand paging
 - Demand segmentation

5

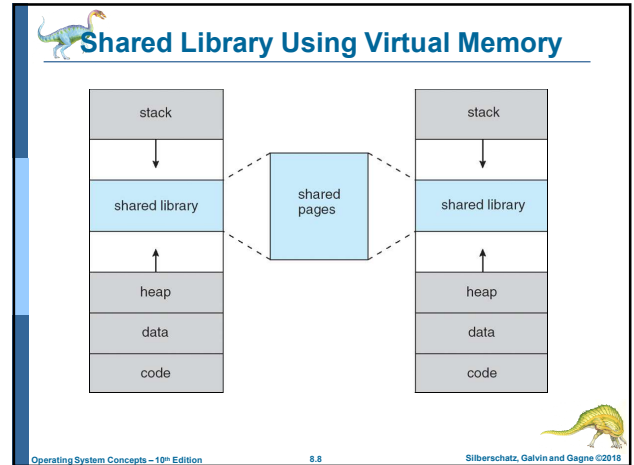
Virtual Memory That is Larger Than Physical Memory



6



7



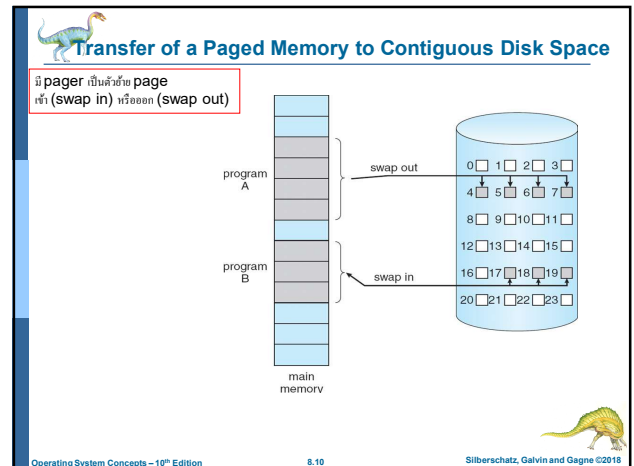
8

Demand Paging

(การจัดสรร **Paging** ตามความต้องการที่ร้องขอ)

- Bring 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 (หากต้องการใช้ page ให้วางตำแหน่งถึง page ที่ต้องการ)
 - invalid reference \Rightarrow abort (หากอ้างอิงถึง page ที่ไม่มีใน memory)
 - not-in-memory \Rightarrow bring to memory (หากอ้างอิงถึง page ที่ไม่มีใน memory ให้นำเข้ามาไว้ใน memory)
- Lazy swapper** – never swaps a page into memory unless page will be needed
 - Swapper that deals with pages is a **pager**

9



10

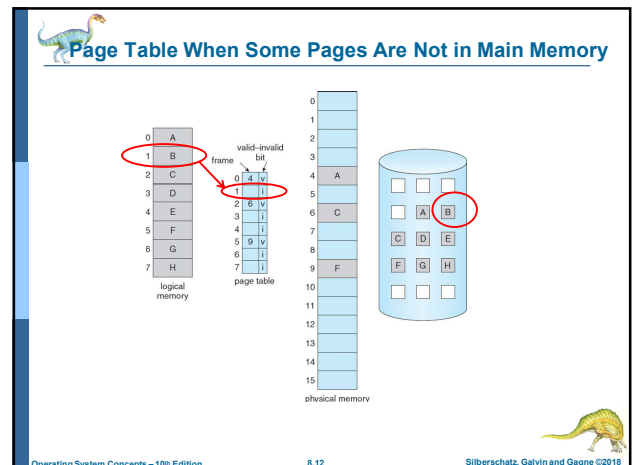
Valid-Invalid Bit

- With each page table entry a valid-invalid bit is associated ($v \Rightarrow$ in-memory, $i \Rightarrow$ not-in-memory)
- Initially valid-invalid bit is set to 1 on all entries
- Example of a page table snapshot:

Frame #	valid-invalid bit
0	v
1	v
2	v
3	v
4	v
5	v
6	v
7	i
...	...
...	i
...	i

During address translation, if valid-invalid bit in page table entry is $i \Rightarrow$ **page fault** (การอ้างอิงผิดหน้าหรือไม่พบหน้าที่ต้องการ)

11



12

Page Fault

เมื่อมีการอ้างอิงถึงหน้าหรือไม่พบหน้าที่ต้องการในหน่วยความจำหลัก (**page fault**) จะมีขั้นตอนดำเนินการดังนี้

- If there is a reference to a page, first reference to that page will trap to operating system:
 - page fault**
- 1. Operating system looks at another table to decide:
 - Invalid reference ⇒ abort
 - Just not in memory
- 2. Get empty frame
- 3. Swap page into frame
- 4. Reset tables
- 5. Set validation bit = **v**
- 6. Restart the instruction that caused the page fault

13

Page Fault (Cont.)

- Restart instruction (ทำต่อจากจุดที่ได้ห้ามมาแล้วล่าสุด หรือ ทำต่อจากจุดที่เกิด page fault ขึ้น)
 - block move
- auto increment/decrement location

14

Steps in Handling a Page Fault

The diagram illustrates the sequence of events during a page fault. It shows the interaction between the operating system, physical memory, and backing store. The steps are: 1. A reference to page M is made. 2. A trap occurs. 3. The page is found on the backing store. 4. The missing page is brought into a free frame in physical memory. 5. The page table is reset. 6. The instruction that caused the fault is restarted.

15

What happens if there is no free frame?

- **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

16

Page Replacement

- **Prevent over-allocation of memory** by modifying page-fault service routine to include page replacement
- Use **modify (dirty) bit** to reduce overhead of page transfers – only modified pages are written to disk
- Page replacement completes separation between logical memory and physical memory – large virtual memory can be provided on a smaller physical memory

17


Need For Page Replacement

The diagram shows two users' logical memory and page tables. User 1's logical memory has pages H, J, and M. User 2's logical memory has pages A, B, D, and E. Both users have a page M. The page tables show the mapping of logical pages to physical frames. The physical memory frames are numbered 0 to 7. The page M from user 1 is in frame 3, and the page M from user 2 is in frame 7. The diagram highlights the need for page replacement when both users have a page M in memory.

18

Basic Page Replacement

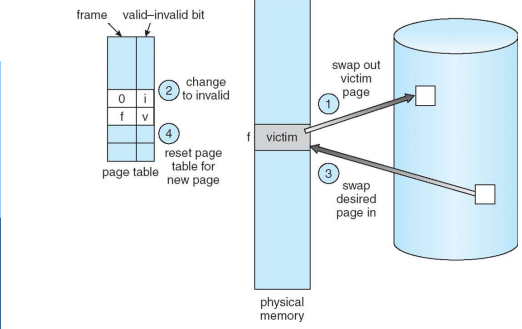
1. Find the location of the desired page on disk
2. Find a free frame:
 - If there is a free frame, use it
 - If there is no free frame, use a page replacement algorithm to select a **victim** frame
3. Bring the desired page into the (newly) free frame; update the page and frame tables
4. Restart the process



Operating System Concepts – 10th Edition 8.19 Silberschatz, Galvin and Gagne ©2018

19

Page Replacement




Operating System Concepts – 10th Edition 8.20 Silberschatz, Galvin and Gagne ©2018

20

Page Replacement Algorithms

- Want lowest page-fault rate
- Evaluate algorithm by running it on a particular string of memory references (reference string) and computing the number of page faults on that string
- In all our examples, the reference string is

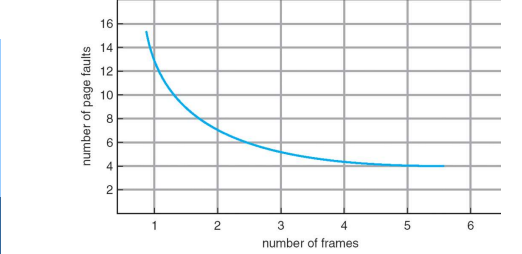
1, 2, 3, 4, 1, 2, 5, 1, 2, 3, 4, 5



Operating System Concepts – 10th Edition 8.21 Silberschatz, Galvin and Gagne ©2018

21

Graph of Page Faults Versus The Number of Frames



Operating System Concepts – 10th Edition 8.22 Silberschatz, Galvin and Gagne ©2018

22

First-In-First-Out (FIFO) Algorithm

- Reference string: 1, 2, 3, 4, 1, 2, 5, 1, 2, 3, 4, 5
- 3 frames (3 pages can be in memory at a time per process)

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


เกิด 9 page faults

□ 4 frames

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

เกิด 10 page faults

- Belady's Anomaly: more frames ⇒ more page faults



Operating System Concepts – 10th Edition 8.23 Silberschatz, Galvin and Gagne ©2018

23

FIFO Page Replacement

มี 0 อยู่แล้วใน memory จึงไม่ต้องไปดึงข้อมูลมาใหม่อีก จึงไม่เกิด page fault สำหรับ page 0


reference string
7 0 1 2 0 3 0 4 2 3 0 3 2 1 2 0 1 7 0 1

7	7	7	2	2	2	4	4	4	0	0	0	0	0	0	0	7	7	7
0	0	0	0	3	3	3	2	2	2	3	3	2	1	1	1	1	0	0
			1	1	1	0	0	0	3	3	2	2	2	1	1	2	2	1

page frames

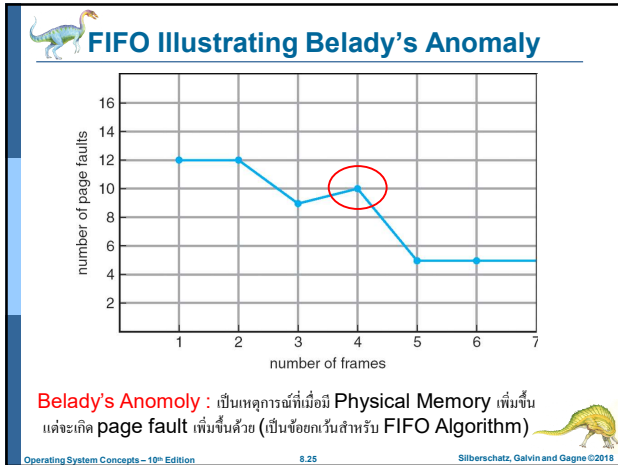
เมื่อ physical memory 3 frames และใช้ reference string ตามที่กำหนดให้ จะเกิด page fault ทั้งหมด 15 page faults

- ครั้งแรกไม่มี page อยู่ใน Physical Memory จะเกิด page fault



Operating System Concepts – 10th Edition 8.24 Silberschatz, Galvin and Gagne ©2018

24



25

Optimal Algorithm

- Replace page that will not be used for longest period of time (มองใน list ของ reference string ถ้าไปในอนาคตว่า page ไตอีกนานที่จะถูกใช้งาน จะโดน replace)
- 4 frames example
1, 2, 3, 4, 1, 2, 5, 1, 2, 3, 4, 5

1	4
2	
3	
4	5

เกิด 6 page faults

ไม่ make sense เราไม่สามารถรู้อนาคตได้

- How do you know this?
- Used for measuring how well your algorithm performs

Operating System Concepts – 10th Edition 8.26 Silberschatz, Galvin and Gagne ©2018

26

Practice: Optimal Page Replacement

reference string
7 0 1 2 0 3 0 4 2 3 0 3 2 1 2 0 1 7 0 1

7	7	7	2	2	2	2	2	2	7
0	0	0	0	0	4	0	0	0	0
1	1	1	3	3	3	3	1	1	1

page frames

เกิดกี่ page faults ???

Operating System Concepts – 10th Edition 8.27 Silberschatz, Galvin and Gagne ©2018

27

Practice: Optimal Page Replacement

reference string
7 0 1 2 0 3 0 4 2 3 0 3 2 1 2 0 1 7 0 1

7	7	7	2	2	2	2	2	2	7
0	0	0	0	0	4	0	0	0	0
1	1	1	3	3	3	3	1	1	1

page frames

เกิด 9 page faults

Answer!!

Operating System Concepts – 10th Edition 8.28 Silberschatz, Galvin and Gagne ©2018

28

Least Recently Used (LRU) Algorithm

(มองย้อนกลับไปเป็นอดีตว่า page ไตไม่ได้อีกใช้มานานที่สุดจะถูก replace)

- Reference string: 1, 2, 3, 4, 1, 2, 5, 1, 2, 3, 4, 5

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

เกิด 8 page faults

- Counter implementation
 - Every page entry has a counter; every time page is referenced through this entry, copy the clock into the counter
 - When a page needs to be changed, look at the counters to determine which are to change

Operating System Concepts – 10th Edition 8.29 Silberschatz, Galvin and Gagne ©2018

29

Practice: LRU Page Replacement

reference string
7 0 1 2 0 3 0 4 2 3 0 3 2 1 2 0 1 7 0 1

7	7	7	2	2	4	4	4	0	1	1	1
0	0	0	0	0	0	0	3	3	3	0	0
1	1	1	3	3	3	2	2	2	2	2	7

page frames

เกิดกี่ page fault ???

Operating System Concepts – 10th Edition 8.30 Silberschatz, Galvin and Gagne ©2018

30

Practice: LRU Page Replacement

Answer!!

reference string
7 0 1 2 0 3 0 4 2 3 0 3 2 1 2 0 1 7 0 1

7	7	7	2	2	4	4	4	0	1	1	1
	0	0	0	0	0	0	3	3	3	0	0
		1	1	3	3	2	2	2	2	2	7

page frames

เกิด 12 page faults

- Better than FIFO but worse than Optimal Algorithm

Operating System Concepts – 10th Edition 8.31 Silberschatz, Galvin and Gagne ©2018

31

Counting Algorithms

- Keep a counter of the number of references that have been made to each page
- LFU Algorithm:** replaces page with smallest count
- MFU Algorithm:** based on the argument that the page with the smallest count was probably just brought in and has yet to be used

LFU : Least Frequently Used (ความถี่ในการถูกใช้ที่น้อยที่สุด)
MFU: Most Frequently Used (ความถี่ในการถูกใช้มากที่สุด)

Operating System Concepts – 10th Edition 8.32 Silberschatz, Galvin and Gagne ©2018

32

Allocation of Frames

- Each process needs *minimum* number of pages
- Example: IBM 370 – 6 pages to handle SS MOVE instruction:
 - instruction is 6 bytes, might span 2 pages
 - 2 pages to handle *from*
 - 2 pages to handle *to*
- Two major allocation schemes
 - fixed allocation
 - priority allocation

Operating System Concepts – 10th Edition 8.33 Silberschatz, Galvin and Gagne ©2018

33

Fixed Allocation

- Equal allocation – For example, if there are 100 frames and 5 processes, give each process 20 frames. (ได้จาก (frame / process) = (100 / 5))
- Proportional allocation – Allocate according to the size of process
 - s_i = size of process p_i
 - $S = \sum s_i$
 - m = total number of frames
 - a_i = allocation for $p_i = \frac{s_i}{S} \times m$

$m = 64$

$s_1 = 10$

$s_2 = 127$

$a_1 = \frac{10}{137} \times 64 \approx 5$

$a_2 = \frac{127}{137} \times 64 \approx 59$

P₁ ได้ 5 frames

P₂ ได้ 59 frames

Operating System Concepts – 10th Edition 8.34 Silberschatz, Galvin and Gagne ©2018

34

Priority Allocation

- Use a proportional allocation scheme using priorities rather than size
- If process P_i generates a page fault,
 - select for replacement one of its frames
 - select for replacement a frame from a process with lower priority number

Operating System Concepts – 10th Edition 8.35 Silberschatz, Galvin and Gagne ©2018

35

Global vs. Local Allocation

- Global replacement** – process selects a replacement frame from the set of all frames; one process can take a frame from another
- Local replacement** – each process selects from only its own set of allocated frames

Operating System Concepts – 10th Edition 8.36 Silberschatz, Galvin and Gagne ©2018

36

End of Chapter 8

