
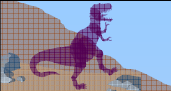


Page Replacement

- 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.




Operating System Concepts 10.1 Silberschatz, Galvin and Gagne ©2002




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. Read the desired page into the (newly) free frame.
Update the page and frame tables.
4. Restart the process.




Operating System Concepts 10.2 Silberschatz, Galvin and Gagne ©2002

Page Replacement



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.



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

9 page faults

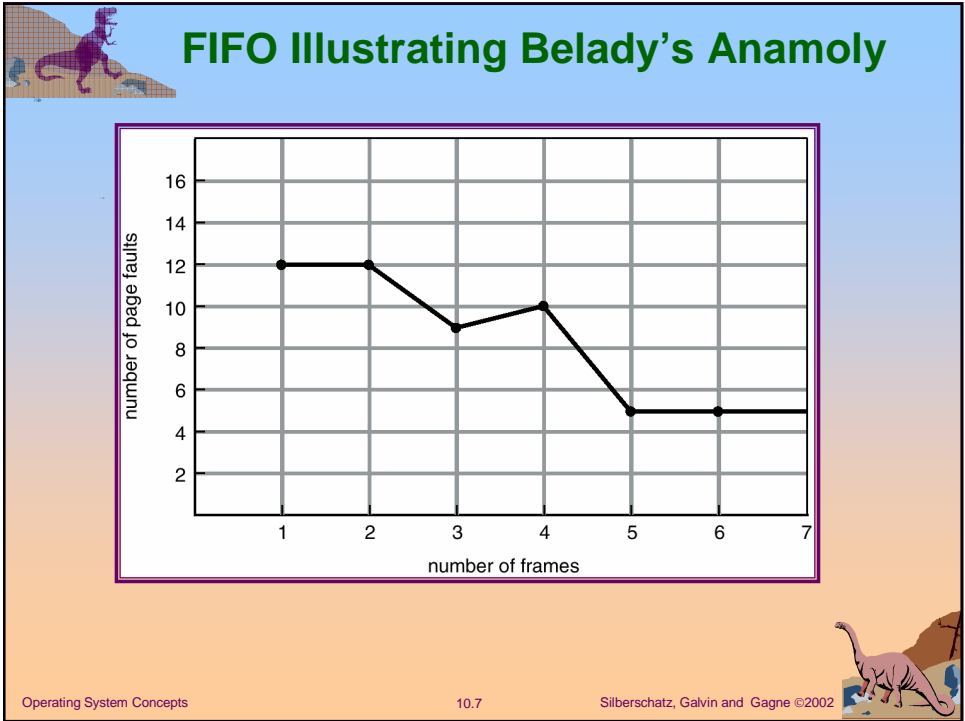
- 4 frames

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

10 page faults

- FIFO Replacement – Belady’s Anomaly
- ☞ more frames ⇒ less page faults

Operating System Concepts 10.6 Silberschatz, Galvin and Gagne ©2002



Optimal Algorithm

- Replace page that will not be used for longest period of time.
- 4 frames example
 - 1, 2, 3, 4, 1, 2, 5, 1, 2, 3, 4, 5

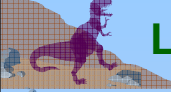
1
2
3
4

4

6 page faults

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

Operating System Concepts 10.8 Silberschatz, Galvin and Gagne ©2002




Least Recently Used (LRU) Algorithm


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

1	5
2	
3	5 4
4	3

- 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 10.9 Silberschatz, Galvin and Gagne ©2002



LRU Algorithm (Cont.)


- Stack implementation – keep a stack of page numbers in a double link form:
 - ☞ Page referenced:
 - ☐ move it to the top
 - ☐ requires 6 pointers to be changed
 - ☞ No search for replacement



Operating System Concepts 10.10 Silberschatz, Galvin and Gagne ©2002

LRU Approximation Algorithms

- Reference bit
 - With each page associate a bit, initially = 0
 - When page is referenced bit set to 1.
 - Replace the one which is 0 (if one exists). We do not know the order, however.
- Second chance
 - Need reference bit.
 - Clock replacement.
 - If page to be replaced (in clock order) has reference bit = 1. then:
 - 📄 set reference bit 0.
 - 📄 leave page in memory.
 - 📄 replace next page (in clock order), subject to same rules.




Operating System Concepts
10.11
Silberschatz, Galvin and Gagne ©2002

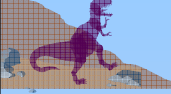
Second-Chance (clock) Page-Replacement Algorithm

	reference bits	pages		reference bits	pages
	0	□		0	□
	0	↓		0	↓
next victim →	1	□		0	↓
	1	↓		0	↓
	0	□		0	↓
	⋮	⋮		⋮	⋮
	1	□		1	↓
	1	↓		1	□

(a)
(b)




Operating System Concepts
10.12
Silberschatz, Galvin and Gagne ©2002

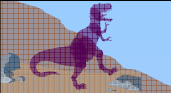


Counting Algorithms

- 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.




Operating System Concepts 10.13 Silberschatz, Galvin and Gagne ©2002



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 10.14 Silberschatz, Galvin and Gagne ©2002



Fixed Allocation

- Equal allocation – e.g., if 100 frames and 5 processes, give each 20 pages.
- 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 $p_i = \frac{s_i}{S} \times m$

$$m = 64$$

$$s_1 = 10$$

$$s_2 = 127$$

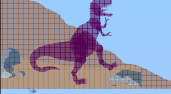
$$a_1 = \frac{10}{137} \times 64 \approx 5$$

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




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.

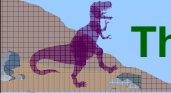


Global vs. Local Allocation

- **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 10.17 Silberschatz, Galvin and Gagne ©2002




Thrashing

- If a process does not have “enough” pages, the page-fault rate is very high. This leads to:
 - ☞ low CPU utilization.
 - ☞ operating system thinks that it needs to increase the degree of multiprogramming.
 - ☞ another process added to the system.
- **Thrashing** ≡ a process is busy swapping pages in and out.




Operating System Concepts 10.18 Silberschatz, Galvin and Gagne ©2002

Thrashing




- Why does paging work?
Locality model
 - Process migrates from one locality to another.
 - Localities may overlap.
- Why does thrashing occur?
 Σ size of locality > total memory size

Locality In A Memory-Reference Pattern



Working-Set Model

- $\Delta \equiv$ working-set window \equiv a fixed number of page references
Example: 10,000 instruction
- WSS_i (working set of Process P_i) = total number of pages referenced in the most recent Δ (varies in time)
 - ☞ if Δ too small will not encompass entire locality.
 - ☞ if Δ too large will encompass several localities.
 - ☞ if $\Delta = \infty \Rightarrow$ will encompass entire program.
- $D = \sum WSS_i \equiv$ total demand frames
- if $D > m \Rightarrow$ Thrashing
- Policy if $D > m$, then suspend one of the processes.



Operating System Concepts 10.21 Silberschatz, Galvin and Gagne ©2002

Working-set model

page reference table

... 2 6 1 5 7 7 7 5 1 6 2 3 4 1 2 3 4 4 4 3 4 3 4 4 4 1 3 2 3 4 4 4 3 4 4 4 ...

Δ

←—————↑

t_1


$WS(t_1) = \{1,2,5,6,7\}$

Δ

←—————↑

t_2

$WS(t_2) = \{3,4\}$



Operating System Concepts 10.22 Silberschatz, Galvin and Gagne ©2002

Keeping Track of the Working Set

- Approximate with interval timer + a reference bit
- Example: $\Delta = 10,000$
 - ☞ Timer interrupts after every 5000 time units.
 - ☞ Keep in memory 2 bits for each page.
 - ☞ Whenever a timer interrupts copy and sets the values of all reference bits to 0.
 - ☞ If one of the bits in memory = 1 \Rightarrow page in working set.
- Why is this not completely accurate?
- Improvement = 10 bits and interrupt every 1000 time units.

Operating System Concepts
10.23
Silberschatz, Galvin and Gagne ©2002

Page-Fault Frequency Scheme

- Establish “acceptable” page-fault rate.
 - ☞ If actual rate too low, process loses frame.
 - ☞ If actual rate too high, process gains frame.

Operating System Concepts
10.24
Silberschatz, Galvin and Gagne ©2002



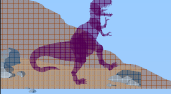
Other Considerations

- Prepaging
- Page size selection
 - ☞ fragmentation
 - ☞ table size
 - ☞ I/O overhead
 - ☞ locality




Other Considerations (Cont.)

- **TLB Reach** - The amount of memory accessible from the TLB.
- $\text{TLB Reach} = (\text{TLB Size}) \times (\text{Page Size})$
- Ideally, the working set of each process is stored in the TLB. Otherwise there is a high degree of page faults.

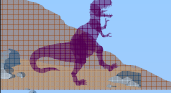


Increasing the Size of the TLB

- **Increase the Page Size.** This may lead to an increase in fragmentation as not all applications require a large page size.
- **Provide Multiple Page Sizes.** This allows applications that require larger page sizes the opportunity to use them without an increase in fragmentation.



Operating System Concepts 10.27 Silberschatz, Galvin and Gagne ©2002



Other Considerations (Cont.)

- Program structure
 - ☞ `int A[][] = new int[1024][1024];`
 - ☞ Each row is stored in one page
 - ☞ Program 1


```


          for (j = 0; j < A.length; j++)
            for (i = 0; i < A.length; i++)
              A[i,j] = 0;
          
```

1024 x 1024 page faults
 - ☞ Program 2



```

          for (i = 0; i < A.length; i++)
            for (j = 0; j < A.length; j++)
              A[i,j] = 0;
          
```

1024 page faults



Operating System Concepts 10.28 Silberschatz, Galvin and Gagne ©2002



Other Considerations (Cont.)

- **I/O Interlock** – Pages must sometimes be locked into memory.
- Consider I/O. Pages that are used for copying a file from a device must be locked from being selected for eviction by a page replacement algorithm.

