

# 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.
- ▶ Virtual memory can be implemented via:
  - Demand paging
  - Demand segmentation



# Demand Paging

- ▶ Bring a page into memory only when it is needed
  - Less I/O needed if not all pages are 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



# Valid-Invalid Bit

- ▶ With each page table entry a valid–invalid bit is associated (1  $\Rightarrow$  in-memory, 0  $\Rightarrow$  not-in-memory)
- ▶ Initially valid–invalid but is set to 0 on all entries
- ▶ Example of a page table snapshot:


Frame #	valid-invalid bit
	1
	1
	1
	1
	0
⋮	
	0
	0

page table

- ▶ During address translation, if valid–invalid bit in page table entry is 0  $\Rightarrow$  page fault




# Page Table When Some Pages Are Not in Main Memory



# Page Fault

- ▶ If there is ever a reference to a page, first reference will trap to OS  $\Rightarrow$  page fault
- ▶ OS looks at another table to decide:
  - Invalid reference  $\Rightarrow$  abort.
  - Just not in memory.
- ▶ Get empty frame.
- ▶ Swap page into frame.
- ▶ Reset tables, validation bit = 1.
- ▶ Restart instruction: Least Recently Used


- block move



- auto increment/decrement location



# Steps in Handling a Page Fault



# 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



# 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)

$$\begin{aligned} \text{EAT} = & (1 - p) \times \text{memory access} \\ & + p (\text{page fault overhead} \\ & + [\text{swap page out}] \\ & + \text{swap page in} \\ & + \text{restart overhead}) \end{aligned}$$





# Demand Paging Example

- ▶ Memory access time = 1 microsecond
- ▶ 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 msec

$$\text{EAT} = (1 - p) \times 1 + p (15000)$$
$$1 + 15000P \quad (\text{in msec})$$



# Process Creation

- ▶ Virtual memory allows other benefits during process creation:
  - Copy-on-Write
  - Memory-Mapped Files (later)



# Copy-on-Write

- ▶ Copy-on-Write (COW) allows both parent and child processes to initially *share* the same pages in memory

If either process modifies a shared page, only then is the page copied

- ▶ COW allows more efficient process creation as only modified pages are copied
- ▶ Free pages are allocated from a **pool** of zeroed-out pages




# 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



# Need For Page Replacement




# Basic Page Replacement

- ▶ Find the location of the desired page on disk
- ▶ 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
- ▶ Read the desired page into the (newly) free frame. Update the page and frame tables.
- ▶ Restart the process



# 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







# Graph of Page Faults Versus The Number of Frames



# 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)



- ▶ FIFO Replacement – Belady’s Anomaly
  - more frames  $\Rightarrow$  more page faults



# FIFO 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																	7	7	7	
	0	0	0																		1	0	0
		1	1																		2	2	1

page frames




# FIFO Illustrating Belady's Anomaly



# 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



6 page faults


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



# 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



page frames

