

# CS 3204 Operating Systems

Lecture 19  
Godmar Back



## Announcements

- Project 3 milestone due **tonight**
- Project 3 due November 8
- Reading assignment:
  - Chapter 8 & 9
    - full read through is highly recommended!



CS 3204 Fall 2006

11/9/2006

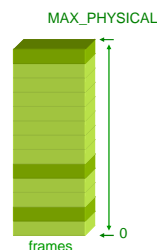
2

## Physical Memory Management



## Physical Memory Management

- Aka frame table management
- Task: keep efficiently track of which physical frames are used
- Allocate a frame when paging in, or eager loading
- Deallocate a frame when process exits or when page is evicted (later)



CS 3204 Fall 2006

11/9/2006

4

## Fragmentation

- Def: *The inability to use memory that is unused.*
- Internal fragmentation:
  - Not all memory inside an allocated unit is used; rest can't be allocated to other users
- External fragmentation:
  - Impossible to satisfy allocation request even though total amount of memory > size requested



CS 3204 Fall 2006

11/9/2006

5

## Internal Fragmentation

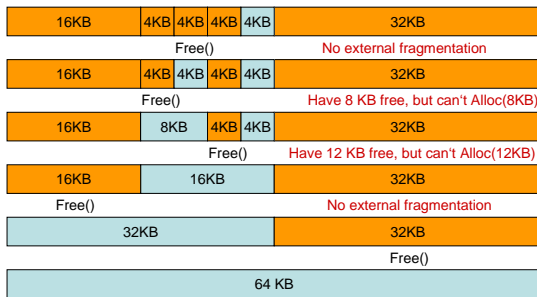


CS 3204 Fall 2006

11/9/2006

6

## External Fragmentation



## Buddy Allocator & Fragmentation

- Q.: what is the average internal fragmentation (per allocated object) for
  - buddy allocator with size  $2^n$ ?
  - in bitmap allocator for objects of size  $n \cdot s$ , where each bit represents a unit of size  $s$ ?
  - in first-fit allocator from project 0?
- Q.: what external fragmentation can you expect from buddy allocator scheme?
- Q.: what's a good way to measure fragmentation in general?

## Page Size & Fragmentation

- How should a system's architect choose the page size? – Trade-Off
- Large pages:
  - Larger internal fragmentation
  - (not an issue if most pages are full...)
  - Page-in & write-back cost larger
- Small pages:
  - Higher overhead to store page table (more entries to maintain)
- Modern architectures provide support for “super pages” – 2MB or 4MB

## Page Replacement

## Page Replacement Algorithms

- Goal: want to minimize number of page faults (situations where a page must be brought in from disk.)
  - Also: want to reduce their cost (ideally, evict those pages from their frames that are already on disk – save writeback time)
- Number of algorithms have been developed
  - Global replacement algorithms
    - Treat frames used by all processes equally
  - Local replacement
    - Pool frames according to user or process when considering replacement

## Replacement Algorithms

- Optimal:
  - “know the future”
  - Obviously impractical, just a benchmark for comparison/analysis
- FIFO – evict oldest page
- LRU – evict least recently used page
- Clock algorithm (“NRU”)
  - Enhanced versions of clock

## Optimal or MIN Replacement

- To analyze algorithms, consider stream of accesses; each access falls into a given page, e.g.  
2 3 2 1 5 2 4 5 3 2 5 2
- Optimal (also known as MIN, or Belady's algorithm)
  - Replace the page that is accessed the farthest in the future, e.g. that won't be accessed for the longest time
- Problem: don't know what the future holds

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

## FIFO

- Evict oldest page:
  - Problem: completely ignores usage pattern – first pages loaded are often frequently accessed

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

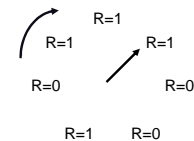
## LRU

- Evict least-recently-used page
- Great if past = future: becomes MIN!
- Major problem: would have to keep track of "recency" on every access, either timestamp, or move to front of a list
  - infeasible

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

## Clock

- Also known as NRU (Not Recently Used) or 2<sup>nd</sup> Chance
- Two ways to look at it:
  - Approximation of LRU
  - FIFO, but keep recently used pages
- Use access (or reference bit)
  - R=1 was accessed
  - R=0 was not accessed
- Hand moves & clears R
- Hand stops when it finds R=0



## Clock Example

- In this example, assume hand advances only on allocation
  - as you can do for Pintos P3
- To avoid running out of frames, use clock daemon that periodically scans pages and resets their access bits
  - Q.: what if clock daemon scans too fast?
  - Q.: what if too slow?

\* means R=1 (page was accessed since last scan)

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

## Variations on Clock Algorithm

- 2-handed Clock
  - If lots of frames, may need to scan many pages until one is found – so introduce second hand
    - Leading hand clears ref bits
    - Trailing hand evicts pages
- Enhanced Clock: exploit modified (or "dirty") bit
  - First find unreferenced & unmodified pages to evict
  - Only if out of those, consider unreferenced & modified pages
  - Clear reference bit as usual

## N-bit Clock Algorithm

- 1-bit says was recently used or wasn't
  - But how recently?
- Idea: associate n-bit counter with page
  - “age” or “act\_count”
  - have R-bit as before
- When hand passes page
  - $\text{act\_count} \gg= 2$  aging
  - $\text{act\_count} |= (R \ll (n-1))$  recent access
- Replace page with lowest act\_count

