

Operating Systems Design (CS 423)

Elsa L Gunter
2112 SC, UIUC

<http://www.cs.illinois.edu/class/cs423/>

Based on slides by Roy Campbell, Sam King, and
Andrew S Tanenbaum


3/7/11

1

Base and Limit

- Load each process into contiguous regions of physical memory

```
If(virt addr > bound){
    trap to kern
} else {
    phys addr =
        virt addr + base
}
```



3/9/11

2

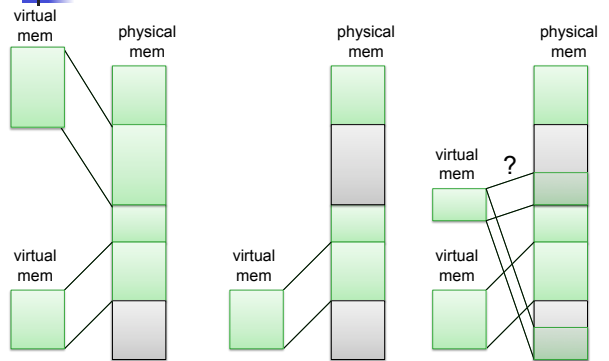
Base and limit

- What must change during a context switch?
- Can a process change its own base and limit?
- Can you share memory with another process?
- How does the kernel handle the address space growing
 - You are the OS designer, come up with an algorithm for allowing processes to grow

3/9/11

3


Memory Allocation in the small



3/9/11

4

Memory Management w/ Linked Lists




- Info for each hole
 - Both ends:
 - Size of hole
 - Bit saying free or allocated
 - Together
 - Next and Previous pointers to other holes
- Info for each allocated "chunk":
 - Size and allocation bit, together, on each end

3/9/11


5

Memory Allocation Algorithms




Allocated 2 units:


- First fit



- Best fit




- Worst fit



3/9/11

6

Memory Deallocation – Linked Lists




Buddy System

- Blocks of memory allocated in sizes that are 2^x
- Determine the largest size block 2^u that fits in available memory
- Fix a smallest size to be allocated 2^l
- When y memory needed, determine smallest k such that $2^k \geq y$
- Find free block of that size
 - Divide it in half repeatedly until small enough
 - Allocate in first piece, marking all others as free

Buddy System - Deallocating

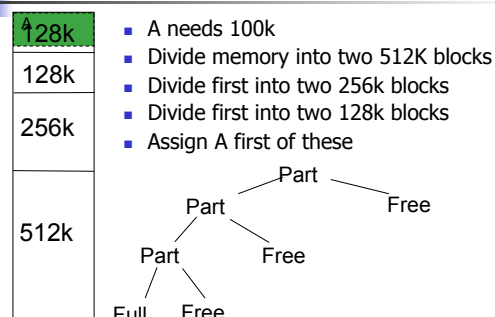
- Each block has a "buddy" to its left or right
- When block is released, its buddy is checked
- When two buddies are free, they are merged
- Minimizes external fragmentation - gaps in allocated memory
- May still have much internal fragmentation - gaps between allocation and actual use

Buddy System

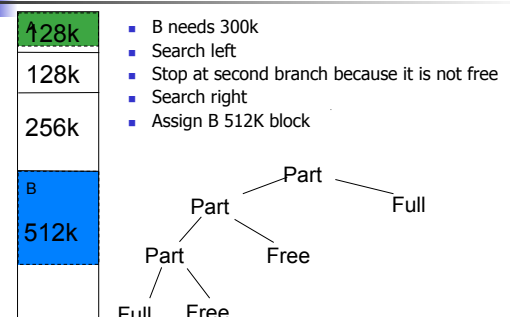


- System Initialization
- $u = 10$; $2^u k = 1024k = 1m$
- $l = 2$; $2^l = 4k$
- Initialize search tree


Buddy System




Buddy System




Buddy System




- C needs 200k
- Find leftmost block that fits
- Assign C 256K block




Buddy System




- D needs 50k
- Split 128K block
- Assign C first 64K block




Buddy System




- Deallocate A
- Buddy is part full; just free A




Buddy System




- Deallocate D
- Immediate buddy is free; merge
- One level up, buddy is free; merge
- Mark 256K free




Buddy System



- Deallocate B
- Immediate buddy is not free
- Just free 512K



Buddy System



- Deallocate C
- Immediate buddy is free; merge
- Next buddy is free; merge
- Mark all 1024k free

Free


Paging

- Allocate physical memory in terms of fixed-size chunk
 - Fixed unit easier to allocate
 - Any free physical page (page frame) can store any virtual page
- Virtual address
 - Virtual page # (high bits of address)
 - Offset (low bits of address, e.g., bits 11-0 for 4k page)

3/11/11

19

Translations table



■ Fig 3.9 from Tanenbaum

3/11/11

20

Translation Process

```

If(virtual page is invalid or non-
resident or protected) {
    trap to OS fault handler
} else {
    physical page # =
    pageTable[virtpage#].physPageNum
}
  
```

3/11/11

21

Translation Process

- What must change on a context switch?

3/11/11

22

Translation Process

- What must change on a context switch?
 - Page table must be replaced
- Each virtual page can be in physical memory or swapped out to disk (called paged)

3/11/11

23

Resident Pages

- How does the processor know that a virtual page is not in memory?

3/11/11

24



Resident Pages

- How does the processor know that a virtual page is not in memory?
 - A bit in the page table entry
- **Resident** means a virtual page is in memory
- NOT an error for a program to access non-resident page

3/11/11

25



Valid Pages Accesses

- Pages can have different protections
 - Read, write, execute
- **Valid** means that a virtual page is legal for the program to access
 - E.g. page not part of the address space is invalid page
- IS an error to try to access an invalid page

3/11/11

26



Valid vs Resident

- Who makes a page resident / non-resident?
- Who makes a virtual page valid / invalid?
- Why would a process want one of its virtual pages to be invalid?

3/11/11

27



Valid vs Resident

- Who makes a page resident / non-resident?
 - OS
- Who makes a virtual page valid / invalid?
 - (user) program
- Why would a process want one of its virtual pages to be invalid?
 - Security, and general protection from self

3/11/11

28