


## TCCS 422: OPERATING SYSTEMS

### Free Space Management

Wes J. Lloyd  
Institute of Technology  
University of Washington - Tacoma



## OBJECTIVES

- Midterm review
- Assignment 1 – grades posted, please double-check
- Assignment 2
- Assignment 3, coming soon
- Chapter 17
  - Memory header
  - Free list operations
- Chapter 18
  - Introduction to memory paging
- Chapter 19
  - Translation Look Ahead Buffer

May 16, 2017

TCCS422: Operating Systems [Spring 2017]  
Institute of Technology, University of Washington - Tacoma

L13.2

## FREE SPACE MANAGEMENT

- Management of memory using
  - Only fixed-sized units: “**pages**”
    - Easy: keep a list
    - Memory request → return first free entry
      - Simple search
  - With variable sized units: “**segments**”
    - More challenging
    - Results from variable sized malloc requests
    - Leads to fragmentation


May 16, 2017

TCCS422: Operating Systems [Spring 2017]  
Institute of Technology, University of Washington - Tacoma

L13.3

## FRAGMENTATION

- Internal vs. external



May 16, 2017

TCCS422: Operating Systems [Spring 2017]  
Institute of Technology, University of Washington - Tacoma

L13.4

## FRAGMENTATION - 2

- **External:** we can compact
  - Example: Client asks for 100 bytes: malloc(100)
  - OS: No 100byte contiguous chunk(s) available: returns NULL
  - Memory is externally fragmented - - Compaction can fix!
- **Internal:** lost space – can't compact
  - OS returns memory units that are too large
  - Example: Client asks for 100 bytes: malloc(100)
  - OS: Returns 125 byte chunk
  - Fragmentation is \*in\* the allocated chunk
  - Memory is lost, and unaccounted for – can't compact


May 16, 2017

TCCS422: Operating Systems [Spring 2017]  
Institute of Technology, University of Washington - Tacoma


L13.5

## ALLOCATION STRATEGY: SPLITTING

- Request for 1 byte of memory: malloc(1)



- OS locates a free chunk to satisfy request
- Splits chunk into two, returns first chunk



May 16, 2017

TCCS422: Operating Systems [Spring 2017]  
Institute of Technology, University of Washington - Tacoma

L13.6

## ALLOCATION STRATEGY: COALESCING

- Consider 30-byte heap
- Free() frees all 10 bytes segments (list of 3-free 10-byte chunks)



- Request arrives: malloc(30)
- No contiguous 30-byte chunk exists
- Coalescing regroups chunks into contiguous chunk



- Allocation can now proceed

May 16, 2017

TCS5422: Operating Systems [Spring 2017]  
Institute of Technology, University of Washington - Tacoma


L13.7

## MEMORY HEADERS

- free(void \*ptr): Does not require a size parameter
- How does the OS know how much memory to free?

- Header block

- Small descriptive block of memory at start of chunk




May 16, 2017

TCS5422: Operating Systems [Spring 2017]  
Institute of Technology, University of Washington - Tacoma

L13.8

## GROWING THE HEAP

- Start with small sized heap
- Request more memory when full
- sbrk(), brk()



May 16, 2017

TCS5422: Operating Systems [Spring 2017]  
Institute of Technology, University of Washington - Tacoma

L13.9

## MEMORY ALLOCATION STRATEGIES

- Best fit

- Traverse free list
- Identify all candidate free chunks
- Note which is smallest (has best fit)
- When splitting, "leftover" pieces are small (and potentially less useful -- fragmented)

- Worst fit

- Traverse free list
- Identify largest free chunk
- Split largest free chunk, leaving a still large free chunk

May 16, 2017

TCS5422: Operating Systems [Spring 2017]  
Institute of Technology, University of Washington - Tacoma

L13.10

## EXAMPLES

- Allocation request for 15 bytes



- Result of Best Fit



- Result of Worst Fit



May 16, 2017

TCS5422: Operating Systems [Spring 2017]  
Institute of Technology, University of Washington - Tacoma

L13.11

## MEMORY ALLOCATION STRATEGIES - 2

- First fit

- Start search at beginning of free list
- Find first chunk large enough for request
- Split chunk, returning a "fit" chunk, saving the remainder
- Avoids full free list traversal of best and worst fit

- Next fit

- Similar to first fit, but start search at last search location
- Maintain a pointer that "cycles" through the list
- Helps balance chunk distribution vs. first fit
- Find first chunk, that is large enough for the request, and split
- Avoids full free list traversal

May 16, 2017

TCS5422: Operating Systems [Spring 2017]  
Institute of Technology, University of Washington - Tacoma

L13.12

## SEGREGATED LISTS

- For popular sized requests  
e.g. for kernel objects such as locks, inodes, etc.
- Manage as segregated free lists
- Provide object caches: stores pre-initialized objects
- How much memory should be dedicated for specialized requests (object caches)?
- If a given cache is low in memory, can request “slabs” of memory from the general allocator for caches.
- General allocator will reclaim slabs when not used


May 16, 2017

TCS422: Operating Systems [Spring 2017]  
Institute of Technology, University of Washington - Tacoma

L13.13

## BUDDY ALLOCATION

- Binary buddy allocation
  - Divides free space by two to find a block that is big enough to accommodate the request; the next split is too small...
- Consider a 7KB request



May 16, 2017

TCS422: Operating Systems [Spring 2017]  
Institute of Technology, University of Washington - Tacoma

L13.14

## BUDDY ALLOCATION - 2

- Buddy allocation: suffers from internal fragmentation
- Allocated fragments, typically too large
- Coalescing is simple
  - Two adjacent blocks are promoted up

May 16, 2017

TCS422: Operating Systems [Spring 2017]  
Institute of Technology, University of Washington - Tacoma

L13.15

## INTRODUCTION TO PAGING

May 16, 2017

TCS422: Operating Systems [Spring 2017]  
Institute of Technology, University of Washington - Tacoma

L13.16

## OBJECTIVES

- Chapter 18
  - Paging
  - Address translation
  - Paging questions

May 16, 2017

TCS422: Operating Systems [Spring 2017]  
Institute of Technology, University of Washington - Tacoma

L13.17

## PAGING

- Split up address space of process into fixed sized pieces called **pages**
- Alternative to variable sized pieces (Segmentation) which suffers from significant fragmentation
- Physical memory is split up into an array of fixed-size slots called **page frames**.
- Each process has a **page table** which translates virtual addresses to physical addresses

May 16, 2017

TCS422: Operating Systems [Spring 2017]  
Institute of Technology, University of Washington - Tacoma

L13.18

ADVANTAGES OF PAGING

- Flexibility
  - Abstracts the process address space into pages
  - No need to track direction of HEAP / STACK growth
  - No need to store unused space
- Simplicity
  - Pages and page frames are the same size
  - Easy to allocate and keep a free list of pages

May 16, 2017

TCCS422: Operating Systems [Spring 2017]  
Institute of Technology, University of Washington - Tacoma

L13.19

PAGING: EXAMPLE

Page Table:  
VP0 → PF3  
VP1 → PF7  
VP2 → PF5  
VP3 → PF2

- Consider a 128 byte address space with 16-byte pages
- Consider a 64-byte program address space

A Simple 64-byte Address Space

64-Byte Address Space Placed In Physical Memory

May 16, 2017

TCCS422: Operating Systems [Spring 2017]  
Institute of Technology, University of Washington - Tacoma

L13.20

PAGING: ADDRESS TRANSLATION

- PAGE: Has two address components
  - VPN: Virtual Page Number
  - Offset: Offset within a Page

Example:  
Page Size: 16-bytes, Address Space: 64-bytes

Here there are just four pages...

May 16, 2017

TCCS422: Operating Systems [Spring 2017]  
Institute of Technology, University of Washington - Tacoma

L13.21

EXAMPLE:  
PAGING ADDRESS TRANSLATION

- Consider a 64-byte program address space (4 pages)
- Stored in 128-byte physical memory (8 frames)
- Offset is preserved
- VPN is looked up

Page Table:  
VP0 → PF3  
VP1 → PF7  
VP2 → PF5  
VP3 → PF2

Virtual Address

Physical Address

VPN

offset

PFN

offset

May 16, 2017

TCCS422: Operating Systems [Spring 2017]  
Institute of Technology, University of Washington - Tacoma

L13.22

PAGING DESIGN QUESTIONS

- Where are page tables stored?
- What are the typical contents of the page table?
- How big are page tables?
- Does paging make the system too slow?

May 16, 2017

TCCS422: Operating Systems [Spring 2017]  
Institute of Technology, University of Washington - Tacoma

L13.23

WHERE ARE PAGE TABLES STORED?

- Example:
  - Consider a 32-bit process address space (up to 4GB)
  - With 4 KB pages
  - 20 bits for VPN ( $2^{20}$  pages)
  - 12 bits for the page offset ( $2^{12}$  unique bytes in a page)
- Page tables for each process are stored in RAM
  - Support potential storage of  $2^{20}$  translations = 1,048,576 pages per process
  - Each page has a page table entry size of 4 bytes

May 16, 2017

TCCS422: Operating Systems [Spring 2017]  
Institute of Technology, University of Washington - Tacoma

L13.24

PAGE TABLE EXAMPLE

- With  $2^{20}$  slots in our page table for a single process
- Each slot dereferences a VPN
- Provides physical frame number
- Each slot requires 4 bytes (32 bits)
  - 20 for the PFN on a 4GB system with 4KB pages
  - 12 for the offset which is preserved
  - (note we have no status bits, so this is unrealistically small)
- How much memory to store page table for 1 process?
  - 4,194,304 bytes (or 4MB) to index one process

VPN <sub>0</sub>
VPN <sub>1</sub>
VPN <sub>2</sub>
...
...
VPN <sub>1048576</sub>

May 16, 2017

TCCS422: Operating Systems [Spring 2017]  
Institute of Technology, University of Washington - Tacoma

L13.25

NOW FOR AN ENTIRE OS

- If 4 MB is required to store one process
- Consider how much memory is required for an entire OS?
  - With for example 100 processes...
- Page table memory requirement is now 4MB x 100 = 400MB
- If computer has 4GB memory (maximum for 32-bits), the page table consumes 10% of memory

400 MB / 4000 GB

- Is this efficient?

May 16, 2017

TCCS422: Operating Systems [Spring 2017]  
Institute of Technology, University of Washington - Tacoma

L13.26

WHAT'S ACTUALLY IN THE PAGE TABLE

- Page table is data structure used to map virtual page numbers (VPN) to the physical address (Physical Frame Number PFN)
  - Linear page table → simple array
- Page-table entry
  - 32 bits for capturing state

31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0		
PFN																				G		R/W		D		A		P		U/S		R	

An x86 Page Table Entry(PTE)

May 16, 2017

TCCS422: Operating Systems [Spring 2017]  
Institute of Technology, University of Washington - Tacoma

L13.27

PAGE TABLE ENTRY

- P: present
- R/W: read/write bit
- U/S: supervisor
- A: accessed bit
- D: dirty bit
- PFN: the page frame number

31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0		
PFN																				G		R/W		D		A		P		U/S		R	

An x86 Page Table Entry(PTE)

May 16, 2017

TCCS422: Operating Systems [Spring 2017]  
Institute of Technology, University of Washington - Tacoma

L13.28

PAGE TABLE ENTRY - 2

- Common flags:
- Valid Bit:** Indicating whether the particular translation is valid.
- Protection Bit:** Indicating whether the page could be read from, written to, or executed from
- Present Bit:** Indicating whether this page is in physical memory or on disk(swapped out)
- Dirty Bit:** Indicating whether the page has been modified since it was brought into memory
- Reference Bit(Accessed Bit):** Indicating that a page has been accessed

May 16, 2017

TCCS422: Operating Systems [Spring 2017]  
Institute of Technology, University of Washington - Tacoma

L13.29

HOW BIG ARE PAGE TABLES?

- Page tables are too big to store on the CPU
- Page tables are stored using physical memory
- Paging supports efficiently storing a sparsely populated address space
  - Reduced memory requirement
  - Compared to base and bounds, and segments

May 16, 2017

TCCS422: Operating Systems [Spring 2017]  
Institute of Technology, University of Washington - Tacoma

L13.30

DOES PAGING MAKE THE SYSTEM TOO SLOW?

■ Translation

■ **Issue #1:** Starting location of the page table is needed

■ HW Support: Page-table base register

■ stores active process

■ Facilitates translation

■ **Issue #2:** Each memory address translation for paging requires an extra memory reference

■ HW Support: TLBs (Chapter 19)

Page Table:  
VP0 → PF3  
VP1 → PF7  
VP2 → PF5  
VP3 → PF2

Stored in RAM →

May 16, 2017

TCCS422: Operating Systems [Spring 2017]  
Institute of Technology, University of Washington - Tacoma

L13.31

PAGING MEMORY ACCESS

```
1. // Extract the VPN from the virtual address
2. VPN = (VirtualAddress & VPN_MASK) >> SHIFT
3.
4. // Form the address of the page-table entry (PTE)
5. PTEAddr = PTBR + (VPN * sizeof(PTE))
6.
7. // Fetch the PTE
8. PTE = AccessMemory(PTEAddr)
9.
10. // Check if process can access the page
11. if (PTE.Valid == False)
12.     RaiseException(SEGMENTATION_FAULT)
13. else if (CanAccess(PTE.ProtectBits) == False)
14.     RaiseException(PROTECTION_FAULT)
15. else
16.     // Access is OK: form physical address and fetch it
17.     offset = VirtualAddress & OFFSET_MASK
18.     PhysAddr = (PTE.PFN << PFN_SHIFT) | offset
19.     Register = AccessMemory(PhysAddr)
```

May 16, 2017

TCCS422: Operating Systems [Spring 2017]  
Institute of Technology, University of Washington - Tacoma

L13.32

COUNTING MEMORY ACCESSES

■ Example: Use this Array initialization Code

```
int array[1000];
...
for (i = 0; i < 1000; i++)
    array[i] = 0;
```

■ Assembly equivalent:

```
0x1024 movl $0x0, (%edi,%eax,4)
0x1028 incl %eax
0x102c cmpl $0x03e8,%eax
0x1030 jne 0x1024
```

May 16, 2017

TCCS422: Operating Systems [Spring 2017]  
Institute of Technology, University of Washington - Tacoma

L13.33

WHERE THE MEMORY ACCESSES ARE:  
FOR THE FIRST 5 LOOP ITERATIONS

■ Locations:

■ Page table

■ Array

■ Code

■ 50 accesses for 5 loop iterations

Iteration	Page Table (PA)	Array (PA)	Code (PA)
1	1024	7132	4096
2	1174	7282	4146
3	1124	7232	4096
4	1074	7132	4146
5	1024	7282	4096

May 16, 2017

TCCS422: Operating Systems [Spring 2017]  
Institute of Technology, University of Washington - Tacoma

L13.34

QUESTIONS

May 16, 2017

TCCS422: Operating Systems [Spring 2017]  
Institute of Technology, University of Washington - Tacoma

L13.35