


TCSS 422: OPERATING SYSTEMS

Memory API, Address Translation, Segmentation

Wes J. Lloyd
School of Engineering and Technology
University of Washington - Tacoma

May 14, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma



Tacoma

OBJECTIVES – 5/14

- **Questions from 5/12**
- **Assignment 2 (based on Ch. 30)**
- **Chapter 13: Introduction to memory virtualization**
 - The address space
 - Goals of OS memory virtualization
- **Chapter 14: Memory API**
 - Common memory errors
- **Chapter 15: Address translation**
 - Base and bounds
 - HW and OS Support
- **Chapter 16: Segmentation**
- **Chapter 17: Free Space Management**
 - Fragmentation, Splitting, coalescing
 - The Free List
 - Memory Allocation Strategies

May 14, 2020	TCSS422: Operating Systems [Spring 2020] School of Engineering and Technology, University of Washington - Tacoma	L13.2
--------------	---	-------

MATERIAL / PACE

- Please classify your perspective on material covered in today’s class (25 respondents):
 - 1-mostly review, 5-equal new/review, 10-mostly new
 - **Average – 6.54 (↓ from 6.87)**
- Please rate the pace of today’s class:
 - 1-slow, 5-just right, 10-fast
 - **Average – 5.86 (↑ from 5.81)**

May 14, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.3

FEEDBACK FROM 5/12

May 14, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.4

OBJECTIVES – 5/14


- Questions from 5/12
- Assignment 2 (based on Ch. 30)
- Chapter 13: Introduction to memory virtualization
 - The address space
 - Goals of OS memory virtualization
- Chapter 14: Memory API
 - Common memory errors
- Chapter 15: Address translation
 - Base and bounds
 - HW and OS Support
- Chapter 16: Segmentation
- Chapter 17: Free Space Management
 - Fragmentation, Splitting, coalescing
 - The Free List
 - Memory Allocation Strategies

May 14, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.5

CHAPTER 13:
ADDRESS SPACES



May 14, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.6

OBJECTIVES – 5/14

- Questions from 5/12
- Assignment 2 (based on Ch. 30)
- **Chapter 13: Introduction to memory virtualization**
 - The address space
 - Goals of OS memory virtualization
- **Chapter 14: Memory API**
 - Common memory errors
- **Chapter 15: Address translation**
 - Base and bounds
 - HW and OS Support
- **Chapter 16: Segmentation**
- **Chapter 17: Free Space Management**
 - Fragmentation, Splitting, coalescing
 - The Free List
 - Memory Allocation Strategies

May 14, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.7

MEMORY VIRTUALIZATION

- What is memory virtualization?
- This is not “virtual” memory,
 - Classic use of disk space as additional RAM
 - When available RAM was low
 - Less common recently

May 14, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.8

MEMORY VIRTUALIZATION - 2

- Presentation of system memory to each process
- Appears as if each process can access the entire machine's address space
- Each process's view of memory is isolated from others
- Everyone has their own sandbox

Process A



Process B



Process C



May 14, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.9

MOTIVATION FOR MEMORY VIRTUALIZATION

- Easier to program
 - Programs don't need to understand special memory models
- Abstraction enables sophisticated approaches to manage and share memory among processes
- Isolation
 - From other processes: easier to code
- Protection
 - From other processes
 - From programmer error (segmentation fault)

May 14, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.10

EARLY MEMORY MANAGEMENT

- Load one process at a time into memory
- Poor memory utilization
- Little abstraction

0KB
64KB
max
Physical Memory

Operating System
(code, data, etc.)

Current Program
(code, data, etc.)

May 14, 2020

TCCS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.11

MULTIPROGRAMMING WITH SHARED MEMORY

- Later machines supported running multiple processes
- Swap out processes during I/O waits to increase system utilization and efficiency
- Swap entire memory of a process to disk for context switch
- Too slow, especially for large processes
- Solution→
 - Leave processes in memory
- Need to protect from errant memory accesses in a multiprocessing environment

0KB
64KB
128KB
192KB
256KB
320KB
384KB
448KB
512KB
Physical Memory

Operating System
(code, data, etc.)

Free

Process C
(code, data, etc.)

Process B
(code, data, etc.)

Free

Process A
(code, data, etc.)

Free

Free

May 14, 2020

TCCS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.12


ADDRESS SPACE

- Easy-to-use abstraction of physical memory for a process

- Main elements:

- Program code
- Stack
- Heap

- Example: 16KB address space



May 14, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.13

ADDRESS SPACE - 2

- Code


- Program code

- Stack

- Program counter (PC)
- Local variables
- Parameter variables
- Return values (for functions)

- Heap

- Dynamic storage
- Malloc() new()




May 14, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.14

ADDRESS SPACE - 3

- Program code
 - Static size
- Heap and stack
 - Dynamic size
 - Grow and shrink during program execution
 - Placed at opposite ends
- Addresses are virtual
 - They must be physically mapped by the OS



May 14, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.15

VIRTUAL ADDRESSING

- Every address is virtual
 - OS translates virtual to physical addresses

```
#include <stdio.h>
#include <stdlib.h>

int main(int argc, char *argv[]){

    printf("location of code : %p\n", (void *) main);
    printf("location of heap : %p\n", (void *) malloc(1));
    int x = 3;
    printf("location of stack : %p\n", (void *) &x);

    return x;
}
```

- EXAMPLE: virtual.c

May 14, 2020

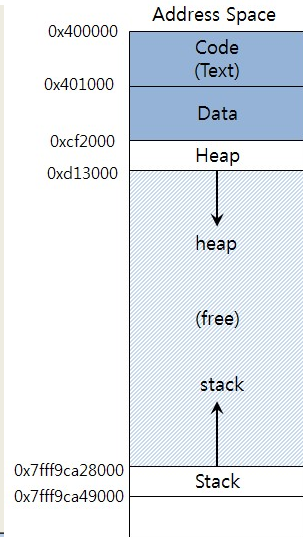
TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.16

VIRTUAL ADDRESSING - 2

■ Output from 64-bit Linux:

location of code: 0x400686
location of heap: 0x1129420
location of stack: 0x7ffe040d77e4



May 14, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.17

GOALS OF OS MEMORY VIRTUALIZATION

■ Transparency

- Memory shouldn't appear virtualized to the program
- OS multiplexes memory among different jobs behind the scenes

■ Protection

- Isolation among processes
- OS itself must be isolated
- One program should not be able to affect another (or the OS)

May 14, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.18

GOALS - 2

■ Efficiency

■ Time

- Performance: virtualization must be fast

■ Space

- Virtualization must not waste space
- Consider data structures for organizing memory
- Hardware support for virtual address translation:
TLB: Translation Lookaside Buffer

■ *Goals considered when evaluating memory virtualization schemes*

May 14, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.19

Which of the following DOES NOT challenge the efficiency of an OS memory virtualization scheme?

Virtual address translation

Data structures used to organize memory


Simultaneous sharing (e.g. multiplexing) of memory among multiple programs

Hardware support - translation lookaside buffer

None of the above

Start the presentation to see live content. For screen share software, share the entire screen. Get help at pollev.com/app

CHAPTER 14: THE MEMORY API



May 14, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.21

OBJECTIVES – 5/14

- **Questions from 5/12**
- **Assignment 2 (based on Ch. 30)**
- **Chapter 13: Introduction to memory virtualization**
 - The address space
 - Goals of OS memory virtualization
- **Chapter 14: Memory API**
 - Common memory errors
- **Chapter 15: Address translation**
 - Base and bounds
 - HW and OS Support
- **Chapter 16: Segmentation**
- **Chapter 17: Free Space Management**
 - Fragmentation, Splitting, coalescing
 - The Free List
 - Memory Allocation Strategies

May 14, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.22

MALLOC

```
#include <stdlib.h>

void* malloc(size_t size)
```

- Allocates memory on the heap
- `size_t` unsigned integer (must be +)
- `size` size of memory allocation in bytes
- Returns
 - SUCCESS: A void * to a memory address
 - FAIL: NULL
- `sizeof()` often used to ask the system how large a given datatype or struct is

May 14, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.23

SIZEOF()

- Not safe to assume data type sizes using different compilers, systems
- Dynamic array of 10 ints
- Static array of 10 ints

```
int *x = malloc(10 * sizeof(int));
printf("%d\n", sizeof(x));
```

4

```
int x[10];
printf("%d\n", sizeof(x));
```

40

May 14, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.24

FREE()

```
#include <stdlib.h>

void free(void* ptr)
```

- Free memory allocated with malloc()
- Provide: (void *) ptr to malloc'd memory
- Returns: nothing

May 14, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.25

```
#include<stdio.h>
```

What will this code do?

```
int * set_magic_number_a()
{
    int a =53247;
    return &a;
}

void set_magic_number_b()
{
    int b = 11111;
}

int main()
{
    int * x = NULL;
    x = set_magic_number_a();
    printf("The magic number is=%d\n",*x);
    set_magic_number_b();
    printf("The magic number is=%d\n",*x);
    return 0;
}
```

26

```
#include<stdio.h>

int * set_magic_number_a()
{
    int a =53247;
    return &a;
}

void set_magic_number_b()
{
    int b = 11111;
}

int main()
{
    int * x = NULL;
    x = set_magic_number_a();
    printf("The magic number is=%d\n",*x);
    set_magic_number_b();
    printf("The magic number is=%d\n",*x);
    return 0;
}
```

What will this code do?

Output:
\$./pointer_error
The magic number is=53247
The magic number is=11111

We have not changed *x but
the value has changed!!

Why?

27

DANGLING POINTER (1/2)

- Dangling pointers arise when a variable referred (a) goes “out of scope”, and it’s memory is destroyed/overwritten (by b) without modifying the value of the pointer (*x).
- The pointer still points to the original memory location of the deallocated memory (a), which has now been reclaimed for (b).

DANGLING POINTER (2/2)

- Fortunately in the case, a compiler warning is generated:

```
$ g++ -o pointer_error -std=c++0x pointer_error.cpp
```

```
pointer_error.cpp: In function 'int*  
set_magic_number_a()':  
pointer_error.cpp:6:7: warning: address of local  
variable 'a' returned [enabled by default]
```

- This is a common mistake - - -
accidentally referring to addresses that have
gone “out of scope”

May 14, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.29

CALLOC()

```
#include <stdlib.h>  
  
void *calloc(size_t num, size_t size)
```

- Allocate “C”lear memory on the heap
- Calloc wipes memory in advance of use...
- `size_t num` : number of blocks to allocate
- `size_t size` : size of each block(in bytes)
- Calloc() prevents...

```
char *dest = malloc(20);  
printf("dest string=%s\n", dest);
```

dest string=??F

May 14, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.30

REALLOC()

```
#include <stdlib.h>

void *realloc(void *ptr, size_t size)
```

- Resize an existing memory allocation
- Returned pointer may be same address, or a new address
 - New if memory allocation must move
- `void *ptr`: Pointer to memory block allocated with `malloc`, `calloc`, or `realloc`
- `size_t size`: New size for the memory block(in bytes)
- EXAMPLE: `realloc.c`
- EXAMPLE: `nom.c`

May 14, 2020

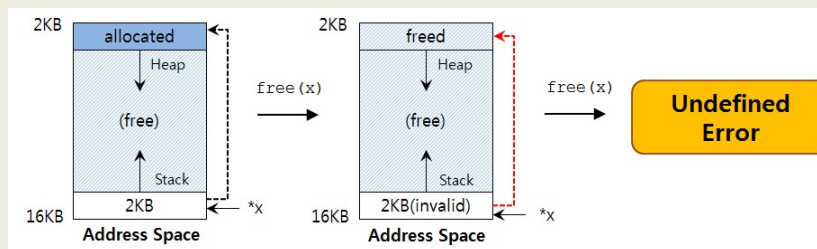
TCCS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.31

DOUBLE FREE

```
int *x = (int *)malloc(sizeof(int)); // allocated
free(x); // free memory
free(x); // free repeatedly
```

- Can't deallocate twice
- Second call core dumps



May 14, 2020

TCCS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.32

SYSTEM CALLS

- **brk(), sbrk()**
 - Used to change data segment size (the end of the heap)
 - Don't use these
- **Mmap(), munmap()**
 - Can be used to create an extra independent “heap” of memory for a user program
- See man page

May 14, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.33

Which of the following is the most performant (e.g. fast) memory API?

calloc()

malloc()

realloc()

None of
the above

Start the presentation to see live content. For screen share software, share the entire screen. Get help at pollev.com/app


L13.34

4

TCSS 422 WILL RETURN AT ~2:40PM

May 14, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma




L13.3
5

CHAPTER 15: ADDRESS TRANSLATION

May 14, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma



L13.36

OBJECTIVES – 5/14

- Questions from 5/12
- Assignment 2 (based on Ch. 30)
- Chapter 13: Introduction to memory virtualization
 - The address space
 - Goals of OS memory virtualization
- Chapter 14: Memory API
 - Common memory errors
- Chapter 15: Address translation
 - Base and bounds
 - HW and OS Support
- Chapter 16: Segmentation
- Chapter 17: Free Space Management
 - Fragmentation, Splitting, coalescing
 - The Free List
 - Memory Allocation Strategies

May 14, 2020

TCCS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.37

ADDRESS TRANSLATION

- 64KB Address space example
- Translation: mapping virtual to physical

The diagram illustrates the mapping of a 64KB virtual address space to physical memory. On the left, the 'Address Space' is shown with segments: Program Code (0KB-16KB), Heap (16KB-32KB), heap (free) (32KB-48KB), stack (48KB-64KB), and Stack (64KB-16KB). On the right, 'Physical Memory' is shown with segments: Operating System (0KB-16KB), (not in use) (16KB-32KB), Code Heap (32KB-48KB), (allocated but not in use) (48KB-64KB), and Stack (64KB-16KB). A 'Relocated Process' is shown in the physical memory, with its code and heap mapped to the physical memory. Dashed lines indicate the mapping from the virtual address space to the physical memory.

May 14, 2020

TCCS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.38

BASE AND BOUNDS

- Dynamic relocation
- Two registers base & bounds: on the CPU
- OS places program in memory
- Sets base register

$$\text{physical address} = \text{virtual address} + \text{base}$$

- Bounds register
 - Stores size of program address space (16KB)
- OS verifies that every address:

$$0 \leq \text{virtual address} < \text{bounds}$$

May 14, 2020

TCCS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.39

INSTRUCTION EXAMPLE

128 : movl 0x0(%ebx), %eax

- Base = 32768
- Bounds =16384
- Fetch instruction at 128 (virt addr) ↑
 - Phy addr = virt addr + base reg
 - 32896 = 128 + 32768 (base)
- Execute instruction
 - Load from address (var x is @ 15kb=15360)
 - 48128 = 15360 + 32768 (base) -- found x...
- Bounds register: terminate process if
 - ACCESS VIOLATION: Virtual address > bounds reg

$$\text{physical address} = \text{virtual address} + \text{base}$$

0KB 128
1KB 132
2KB 135

movl 0x0(%ebx), %eax
Addl 0x03, %eax
movl %eax, 0x0(%ebx)

Program Code

Heap

↓ heap

(free)

↑ stack

14KB
15KB 3000
16KB

Int x
Stack

May 14, 2020

TCCS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.40

MEMORY MANAGEMENT UNIT

- MMU

- Portion of the CPU dedicated to address translation
- Contains base & bounds registers

- Base & Bounds Example:

- Consider address translation
- 4 KB (4096 bytes) address space, loaded at 16 KB physical location

Virtual Address	Physical Address
0	16384
1024	17408
3000	19384
4400	20784 (out of bounds)

FAULT

May 14, 2020

TCCS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.41

DYNAMIC RELOCATION OF PROGRAMS

- Hardware requirements:

Requirements	HW support
Privileged mode	CPU modes: kernel, user
Base / bounds registers	Registers to support address translation
Translate virtual addr; check if in bounds	Translation circuitry, check limits
Privileged instruction(s) to update base / bounds regs	Instructions for modifying base/bound registers
Privileged instruction(s) to register exception handlers	Set code pointers to OS code to handle faults
Ability to raise exceptions	For out-of-bounds memory access, or attempts to access privileged instr.

May 14, 2020

TCCS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.42

OS SUPPORT FOR MEMORY VIRTUALIZATION

- For base and bounds OS support required
 - When process starts running
 - Allocate address space in physical memory
 - When a process is terminated
 - Reclaiming memory for use
 - When context switch occurs
 - Saving and storing the base-bounds pair
 - Exception handlers
 - Function pointers set at OS boot time

May 14, 2020

TCCS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.43

OS: WHEN PROCESS STARTS RUNNING

- OS searches for free space for new process
 - Free list: data structure that tracks available memory slots

The OS lookup the free list

Free list

16KB

48KB

0KB

16KB

32KB

48KB

64KB

Operating System

(not in use)

Code

Heap

(allocated but not in use)

Stack

(not in use)

Physical Memory

May 14, 2020

TCCS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.44

OS: WHEN PROCESS IS TERMINATED

■ OS places memory back on the free list

0KB

16KB

32KB

48KB

64KB

Operating System

(not in use)

Process A

(not in use)

16KB

48KB

Physical Memory

Free list

0KB

16KB

32KB

48KB

64KB

Operating System

(not in use)

(not in use)

(not in use)

16KB

32KB

48KB

Physical Memory

May 14, 2020

TCCS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.45

OS: WHEN CONTEXT SWITCH OCCURS

■ OS must save base and bounds registers

■ Saved to the Process Control Block PCB (task_struct in Linux)

0KB

16KB

32KB

48KB

64KB

Operating System

(not in use)

Process A
Currently Running

Process B

base

32KB

bounds

48KB

Physical Memory

Context Switching

0KB

16KB

32KB

48KB

64KB

Operating System

(not in use)

Process A

Process B
Currently Running

base

48KB

bounds

64KB

Physical Memory

Process A PCB

...
base : 32KB
bounds : 48KB
...

May 14, 2020

TCCS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.46

DYNAMIC RELOCATION

- OS can move process data when not running
 1. OS deschedules process from scheduler
 2. OS copies address space from current to new location
 3. OS updates PCB (base and bounds registers)
 4. OS reschedules process
- When process runs new base register is restored to CPU
- **Process doesn't know it was even moved!**

May 14, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.47

Consider a 64KB computer the loads a program. The BASE register is set to 32768, and the BOUNDS register is set to 4096. What is the physical memory address translation for a virtual address of 6000 ?

34768

38768

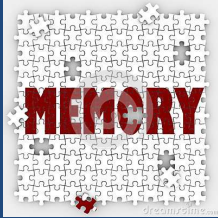
32769

36864

Out of bounds

Start the presentation to see live content. For screen share software, share the entire screen. Get help at pollev.com/app

CHAPTER 16: SEGMENTATION



May 14, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.49

OBJECTIVES – 5/14

- **Questions from 5/12**
- **Assignment 2 (based on Ch. 30)**
- **Chapter 13: Introduction to memory virtualization**
 - The address space
 - Goals of OS memory virtualization
- **Chapter 14: Memory API**
 - Common memory errors
- **Chapter 15: Address translation**
 - Base and bounds
 - HW and OS Support
- **Chapter 16: Segmentation**
- **Chapter 17: Free Space Management**
 - Fragmentation, Splitting, coalescing
 - The Free List
 - Memory Allocation Strategies


May 14, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.50

BASE AND BOUNDS INEFFICIENCIES

- **Address space**
 - Contains significant unused memory
 - Is relatively large
 - Preallocates space to handle stack/heap growth
- **Large address spaces**
 - Hard to fit in memory
- **How can these issues be addressed?**



May 14, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.51

MULTIPLE SEGMENTS

- **Memory segmentation**
- **Manage the address space as (3) separate segments**
 - Each is a contiguous address space
 - Provides logically separate segments for: code, stack, heap
- **Each segment can be placed separately**
- **Track base and bounds for each segment (registers)**

May 14, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.52

SEGMENTS IN MEMORY

■ Consider 3 segments:

0KB

Operating System

16KB

(not in use)

Stack

(not in use)

32KB

Code

Heap

48KB

(not in use)

64KB

Physical Memory

Much smaller

↓

Segment	Base	Size
Code	32K	2K
Heap	34K	2K
Stack	28K	2K

May 14, 2020

TCCS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.53

ADDRESS TRANSLATION: CODE SEGMENT

$$physical\ address = offset + base$$

■ Code segment - physically starts at 32KB (base)

■ Starts at "0" in virtual address space

Segment	Base	Size
Code	32KB	2KB

0KB

100 instruction

2KB

Program

4KB

Virtual Address Space

Physical Address Space

(not in use)

Bounds check:

Is virtual address within 2KB address space?

or 32868

desired

address

May 14, 2020

TCCS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.54

Slides by Wes J. Lloyd

L13.27

ADDRESS TRANSLATION: HEAP

Virtual address + base is not the correct physical address.

- Heap starts at virtual address 4096
- The data is at 4200
- Offset= 4200 - 4096 = 104 (virt addr - virt heap start)
- Physical address = 104 + 34816 (offset + heap base)

Segment	Base	Size
Heap	34K	2K

Address Space

Physical Memory

May 14, 2020

TCCS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.55

SEGMENTATION FAULT

- Access beyond the address space
- Heap starts at virtual address: 4096
- Data pointer is to 7KB (7168)
- Is data pointer valid?

- Heap starts at 4096 + 2 KB seg size = 6144
- Offset= 7168 > 4096 + 2048 (6144)

Address Space


May 14, 2020

TCCS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma


L13.56

SEGMENT REGISTERS

- Used to dereference memory during translation



- First two bits identify segment type
- Remaining bits identify memory offset
- Example: virtual heap address 4200 (01000001101000)



Segment	bits
Code	00
Heap	01
Stack	10
-	11

May 14, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.57

SEGMENTATION DEREFERENCE

```

1 // get top 2 bits of 14-bit VA
2 Segment = (VirtualAddress & SEG_MASK) >> SEG_SHIFT
3 // now get offset
4 Offset = VirtualAddress & OFFSET_MASK
5 if (Offset >= Bounds[Segment])
6     RaiseException(PROTECTION_FAULT)
7 else
8     PhysAddr = Base[Segment] + Offset
9     Register = AccessMemory(PhysAddr)

```

- VIRTUAL ADDRESS = 01000001101000 (on heap)
- SEG_MASK = 0x3000 (11000000000000)
- SEG_SHIFT = 01 → *heap* (mask gives us segment code)
- OFFSET_MASK = 0xFFF (00111111111111)
- OFFSET = 000001101000 = 104 (isolates segment offset)
- OFFSET < BOUNDS : 104 < 2048


May 14, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.58

STACK SEGMENT

- Stack grows backwards (FILO)
- Requires hardware support:
- Direction bit: tracks direction segment grows



Segment Register(with Negative-Growth Support)

Segment	Base	Size	Grows	Positive?
Code	32K	2K	1	
Heap	34K	2K	1	
Stack	28K	2K	0	

May 14, 2020

TCCS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.59

SHARED CODE SEGMENTS

- Code sharing: enabled with HW support
- Supports storing shared libraries in memory only once
- DLL: dynamic linked library
- .so (linux): shared object in Linux (under /usr/lib)
- Many programs can access them
- Protection bits: track permissions to segment

Segment Register Values(with Protection)

Segment	Base	Size	Grows	Positive?	Protection
Code	32K	2K	1		Read-Execute
Heap	34K	2K	1		Read-Write
Stack	28K	2K	0		Read-Write

May 14, 2020

TCCS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.60


Consider a program with 2KB of code, a 1 KB stack, and a 2 KB heap. This program runs on a 64 KB computer that manages memory with 4 kb segments. If the computer is empty and segments were allocated as: code, stack, heap, how large can the heap grow to?

32 KB
56 KB
24 KB
4 KB
0 KB

Start the presentation to see live content. For screen share software, share the entire screen. Get help at pollev.com/app

SEGMENTATION GRANULARITY

- Coarse-grained
- Manage memory as large purpose based segments:
 - Code segment
 - Heap segment
 - Stack segment




May 14, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.62

SEGMENTATION GRANULARITY - 2

- Fine-grained
- Manage memory as list of segments
- Code, heap, stack segments composed of multiple smaller segments
- Segment table
 - On early systems
 - Stored in memory
 - Tracked large number of segments



May 14, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.63

MEMORY FRAGMENTATION

- Consider how much free space?
- We'll say about 24 KB
- Request arrives to allocate a 20 KB heap segment
- Can we fulfil the request for 20 KB of contiguous memory?

	Not compacted
0KB	
8KB	Operating System
16KB	
24KB	(not in use)
32KB	Allocated
40KB	(not in use)
48KB	Allocated
56KB	(not in use)
64KB	Allocated

May 14, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.64

COMPACTION

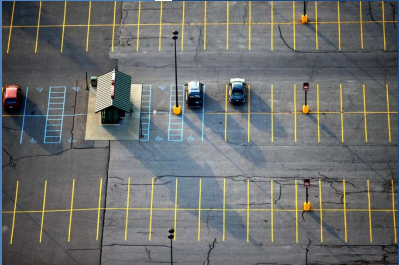
- Supports rearranging memory
- Can we fulfil the request for 20 KB of contiguous memory?
- **Drawback: Compaction is slow**
 - Rearranging memory is time consuming
 - 64KB is fast
 - 4GB+ ... slow
- **Algorithms:**
 - Best fit: keep list of free spaces, allocate the most snug segment for the request
 - Others: worst fit, first fit... (in future chapters)

Compacted	
0KB	Operating System
8KB	
16KB	
24KB	Allocated
32KB	
40KB	
48KB	(not in use)
56KB	
64KB	

May 14, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.65



CHAPTER 17: FREE SPACE MANAGEMENT

May 14, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.66

OBJECTIVES – 5/14

- Questions from 5/12
- Assignment 2 (based on Ch. 30)
- Chapter 13: Introduction to memory virtualization
 - The address space
 - Goals of OS memory virtualization
- Chapter 14: Memory API
 - Common memory errors
- Chapter 15: Address translation
 - Base and bounds
 - HW and OS Support
- Chapter 16: Segmentation
- Chapter 17: Free Space Management
 - Fragmentation, Splitting, coalescing
 - The Free List
 - Memory Allocation Strategies

May 14, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.67

FREE SPACE MANAGEMENT

- How should free space be managed, when satisfying variable-sized requests?
- What strategies can be used to minimize fragmentation?
- What are the time and space overheads of alternate approaches?

May 14, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.68

FREE SPACE MANAGEMENT

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

May 14, 2020

TCCS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.69

- Consider a 30-byte heap

30-byte heap:

free	used	free	
0	10	20	30

- Request for 15-bytes

free list: head

addr:0
len:10

addr:20
len:10

→ NULL

- Free space: 20 bytes
- No available contiguous chunk → return NULL

May 14, 2020

TCCS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.70

FRAGMENTATION - 2

- **External:** *OS can compact*
 - Example: Client asks for 100 bytes: malloc(100)
 - OS: No 100 byte contiguous chunk is available: returns NULL
 - Memory is externally fragmented - - Compaction can fix!
- **Internal:** *lost space – OS 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 14, 2020

TCCS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.71

ALLOCATION STRATEGY: SPLITTING

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

30-byte heap:

free	used	free	
0	10	20	30

free list: head →

addr:0	len:10
--------	--------

 →

addr:20	len:10
---------	--------

 → NULL

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

30-byte heap:

free	used	free	
0	10	21	30

free list: head →

addr:0	len:10
--------	--------

 →

addr:21	len:9
---------	-------

 → NULL

May 14, 2020

TCCS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.72

ALLOCATION STRATEGY: COALESCING

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

head →

addr:10
len:10

addr:0
len:10

addr:20
len:10

→ NULL

- Request arrives: malloc(30)
- **SPLIT DOES NOT WORK** - no contiguous 30-byte chunk exists!
- Coalescing regroups chunks into contiguous chunk

head →

addr:0
len:30

→ NULL

- Allocation can now proceed
- Coalescing is defragmentation of the free space list

May 14, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.73

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

ptr →

The header used by malloc library

The 20 bytes returned to caller


An Allocated Region Plus Header

May 14, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.74

MEMORY HEADERS - 2



- Contains size
- Pointers: for faster memory access
- Magic number: integrity checking

May 14, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.75

MEMORY HEADERS - 3

- Size of memory chunk is:
 - Header size + user malloc size
 - N bytes + sizeof(header)
- Easy to determine address of header

```
void free(void *ptr) {  
    header_t *hptr = (void *)ptr - sizeof(header_t);  
}
```

May 14, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.76

THE FREE LIST

- Simple free list struct

```
typedef struct __node_t {  
    int size;  
    struct __node_t *next;  
} node_t;
```

- Use mmap to create free list
- 4kb heap, 4 byte header, one contiguous free chunk

```
// mmap() returns a pointer to a chunk of free space  
node_t *head = mmap(NULL, 4096, PROT_READ|PROT_WRITE,  
                    MAP_ANON|MAP_PRIVATE, -1, 0);  
head->size = 4096 - sizeof(node_t);  
head->next = NULL;
```

May 14, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.77

FREE LIST - 2

- Create and initialize free-list “heap”

```
// mmap() returns a pointer to a chunk of free space  
node_t *head = mmap(NULL, 4096, PROT_READ|PROT_WRITE,  
                    MAP_ANON|MAP_PRIVATE, -1, 0);  
head->size = 4096 - sizeof(node_t);  
head->next = NULL;
```

- Heap layout:

The diagram illustrates the memory layout of the free list. A pointer labeled 'head' points to a node structure. This node has two fields: 'size' with the value 4088, and 'next' with the value 0. To the right of the node, text indicates '[virtual address: 16KB] header: size field' for the size field and 'header: next field(NULL is 0)' for the next field. Below the node, a bracket indicates 'the rest of the 4KB chunk' which contains three dots '...', representing the remaining memory space.

May 14, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.78

FREE LIST: MALLOC() CALL

- Consider a request for a 100 bytes: `malloc(100)`
- Header block requires 8 bytes
 - 4 bytes for size, 4 bytes for magic number
- Split the heap – header goes with each block

A 4KB Heap With One Free Chunk

head →

size:	4088
next:	0
...	

the rest of the 4KB chunk

A Heap : After One Allocation

ptr →

size:	100
magic:	1234567
First block is used	

the 100 bytes now allocated

head →

size:	3980
next:	0
...	

the free 3980 byte chunk

May 14, 2020

TCCS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.79

FREE LIST: FREE() CALL

- Addresses of chunks
- $\text{Start}=16384$
 - + 108 (end of 1st chunk)
 - + 108 (end of 2nd chunk)
 - + 108 (end of 3rd chunk)= 16708

8 bytes header

ptr →

size:	100
magic:	1234567
...	

100 bytes still allocated

sptr →

size:	100
magic:	1234567
Free this block	

100 bytes still allocated (but about to be freed)

head →

size:	100
magic:	1234567
...	

100 bytes still allocated

head →

size:	3764
next:	0
...	

The free 3764-byte chunk

Free Space With Three Chunks Allocated

May 14, 2020

TCCS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.80

FREE LIST:
FREE() CHUNK #2

- Free(sptr)
- Our 3 chunks start at 16 KB (@ 16,384 bytes)
- Free chunk #2 - sptr
- Sptr = 16500
 - addr – sizeof(node_t)
- Actual start of chunk #2
 - 16492

[virtual address: 16KB]

size: 100
magic: 1234567
...
size: 100
next: 16708
Block Now Free
size: 100
magic: 1234567
...
size: 3764
next: 0
...

head
sptr

100 bytes still allocated
(now a free chunk of memory)
100 bytes still allocated
The free 3764-byte chunk

May 14, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.81

FREE LIST- FREE ALL CHUNKS

- Now free remaining chunks:
- Free(16392)
- Free(16608)
- Walk back 8 bytes for actual start of chunk
- External fragmentation
- Free chunk pointers out of order
- Coalescing of next pointers is needed

[virtual address: 16KB]

size: 100
next: 16492
...
size: 100
next: 16708
...
size: 100
next: 16384
...
size: 3764
next: 0
...

head

(now free)
(now free)
(now free)
The free 3764-byte chunk


May 14, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.82

GROWING THE HEAP

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



May 14, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.83

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 14, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.84

EXAMPLES

■ Allocation request for 15 bytes



■ Result of Best Fit



■ Result of Worst Fit



May 14, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.85

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 14, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.86

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 14, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.87

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

The diagram illustrates the binary buddy allocation process. It starts with a single 64 KB block. This block is split into two 32 KB blocks. One 32 KB block is further split into two 16 KB blocks. One 16 KB block is split into two 8 KB blocks. The final state shows a 64 KB free space for a 7 KB request, indicating that the 8 KB blocks are not suitable for the request because the next split would result in blocks too small to accommodate the request.

May 14, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.88

BUDDY ALLOCATION - 2

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

May 14, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L13.89

QUESTIONS



**WILL RETURN IN A FEW
MINUTES**

