


TCCS 422: OPERATING SYSTEMS

Address Spaces
and the Memory API

Wes J. Lloyd

Institute of Technology

University of Washington - Tacoma



UPDATES

C Help: Viveret, Ian - recommended mentors

Winter 2017 Mentor Lab Schedule (SCI 108)

Jan. 9 - Mar. 3*

Monday	Tuesday	Wednesday	Thursday	Friday	additional 305 mentoring (Adam) in DOU 110 MW 1:30 - 4 F 3:00 - 5
Chinh (10:00 - 12:30)	Chinh (10:00 - 12:30)	Chinh (10:00 - 12:30)	Chinh (10:00 - 12:30)		
Ian (12:30 - 3:30)	Viveret (12:30 - 3:30)	Ian (12:30 - 3:30)	Mike (1:00 - 3:00)	Ian (12:30 - 4:30)	
Viveret (3:30 - 6:50)	Mike (3:30 - 7:30)	Viveret (3:30 - 6:50)	Mike (3:30 - 7:30)		

February 15, 2017

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

L12.2

UPDATES

Program 1:

To create a tar gzip archive:

From your source directory with mash.c

cd ..

tar czf mash.tar.gz mash

Creates a tar archive file which is automatically gzipped

February 15, 2017

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

L12.3

UPDATES

Programs:

Please start early

Work as if though the deadline is 2 days earlier

When do students start working?

Due Date

90% 90% 90%

0% 0% 0%

60% 40% 20% 0%

0% 20% 40% 60%

>8 8 7 6 5 4 3 2 1 -1 -2 <-2

Days before due date

A/B Grades

C/D/F Grades

Better than 50% chance of A/B

Less than 50% chance of A/B

From Virginia Tech Department of Computer Science - 2011

February 15, 2017

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

L12.4

OBJECTIVES

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

Memory segments, fragmentation

February 15, 2017

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

L12.5

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

February 15, 2017

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

L12.6

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

February 15, 2017

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

L12.7

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)

February 15, 2017

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

L12.8

EARLY MEMORY MANAGEMENT

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

February 15, 2017

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

L12.9

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

February 15, 2017

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

L12.10

ADDRESS SPACE

- Easy-to-use abstraction of physical memory for a process
- Main elements:
 - Program code
 - Stack
 - Heap
- Example: 16KB address space

February 15, 2017

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

L12.11

ADDRESS SPACE - 2

- Code
 - Program code
- Stack
 - Program counter (PC)
 - Local variables
 - Parameter variables
 - Return values (for functions)
- Heap
 - Dynamic storage
 - Malloc() new()

February 15, 2017

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

L12.12

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

Address Space

February 15, 2017 TCS5422: Operating Systems [Winter 2017]
Institute of Technology, University of Washington - Tacoma L12.13

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

February 15, 2017 TCS5422: Operating Systems [Winter 2017]
Institute of Technology, University of Washington - Tacoma L12.14

VIRTUAL ADDRESSING - 2

- Output from 64-bit Linux:

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

Address Space

February 15, 2017 TCS5422: Operating Systems [Winter 2017]
Institute of Technology, University of Washington - Tacoma L12.15

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)

February 15, 2017 TCS5422: Operating Systems [Winter 2017]
Institute of Technology, University of Washington - Tacoma L12.16

GOALS - 2

- Efficiency
 - Time
 - Performance: virtualization must be fast
 - Space
 - Virtualization must not waste space
 - Consider data structures for organizing memory
 - Hardware support TLB: Translation Lookaside Buffer
- Goals considered when evaluation memory virtualization schemes

February 15, 2017 TCS5422: Operating Systems [Winter 2017]
Institute of Technology, University of Washington - Tacoma L12.17

CHAPTER 14: THE MEMORY API

February 15, 2017 TCS5422: Operating Systems [Winter 2017]
Institute of Technology, University of Washington - Tacoma L12.18

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

February 15, 2017

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

L12.19

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

February 15, 2017

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

L12.20

FREE()

```
#include <stdlib.h>
void free(void* ptr)
```

- Free memory allocated with malloc()
- Provide: (void *) ptr to malloc'd memory

- Returns: nothing

February 15, 2017

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

L12.21

VIRTUAL ADDRESS SPACE

```
int *pi; // local variable
```

- Pointer is a local variable on the stack
- Malloc returns space on the heap

```
pi = (int *)malloc(sizeof(int) * 4);
```

February 15, 2017

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

L12.22

VIRTUAL ADDRESS SPACE - 2

- Releases heap space pointed to by the pointer on the stack

```
free(pi);
```

February 15, 2017

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

L12.23

COMMON MEMORY ERRORS

- Forgetting to malloc memory
- Unterminated string
- Uninitialized memory
- Memory leak
- Dangling pointer

February 15, 2017

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

L12.24

FORGETTING TO MALLOC

■ C is not Java

■ When forgetting to malloc:

```
char *src = "hello"; //character string constant
char *dst;           //unallocated
strcpy(dst, src);     //segfault and die
```

dst has not been initialized. It has no place to store anything

strcpy(dst, src);

Address Space

heap (free) stack

unallocated

helloW0

*dst

*src

Segmentation fault (core dumped)

February 15, 2017

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

L12.25

CORRECTION

```
char *src = "hello"; //character string constant
char *dst (char *)malloc(strlen(src) + 1 ); // allocated
strcpy(dst, src);     //work properly
```

■ Why do we malloc length + 1 ?

strcpy(dst, src);

Address Space

heap (free) stack

helloW0

allocated

*dst

*src

February 15, 2017

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

L12.26

UNTERMINATED STRING

```
char *src = "hello"; //character string constant
char *dst (char *)malloc(strlen(src)); // too small
strcpy(dst, src);     //work properly
```

Malloc too little memory

strcpy(dst, src);

Address Space

heap (free) stack

helloW0

6 bytes

strlen

5 bytes

*dst

*src

February 15, 2017

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

L12.27

FORGETTING TO INITIALIZE

```
int *x = (int *)malloc(sizeof(int)); // allocated
printf("x = %d\n", *x); // uninitialized memory access
```

Address Space

heap (free) stack

value used before (free)

*x

Address Space

heap (free) stack

allocated with value used before

*x

February 15, 2017

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

L12.28

MEMORY LEAK

unused : unused, but not freed

Program runs out of memory and eventually dies...

Address Space

heap (free) stack

allocated

*a

Address Space

heap (free) stack

allocated

unused

*b

*a

Address Space

heap (free) stack

allocated

unused

unused

unused

*d

*c

*b

*a

run out of memory

February 15, 2017

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

L12.29

#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;
}
```

30

```
#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?

31

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).

February 15, 2017 TCS5422: Operating Systems [Winter 2017]
Institute of Technology, University of Washington - Tacoma L12.32

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"

February 15, 2017 TCS5422: Operating Systems [Winter 2017]
Institute of Technology, University of Washington - Tacoma L12.33

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
```

February 15, 2017 TCS5422: Operating Systems [Winter 2017]
Institute of Technology, University of Washington - Tacoma L12.34

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

February 15, 2017 TCS5422: Operating Systems [Winter 2017]
Institute of Technology, University of Washington - Tacoma L12.35

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

February 15, 2017 TCS5422: Operating Systems [Winter 2017]
Institute of Technology, University of Washington - Tacoma L12.36

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

February 15, 2017

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

L12.37

CHAPTER 15: ADDRESS TRANSLATION



February 15, 2017

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

L12.38

OBJECTIVES

- Address translation
- Base and bounds
- HW and OS Support
- Memory segments
- Memory fragmentation

February 15, 2017

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

L14.39

MEMORY VIRTUALIZATION

- Using hardware support provide virtualization that is:
 - Efficient
 - Flexible
 - Secure and isolated

February 15, 2017

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

L14.40

HARDWARE BASED ADDRESS TRANSLATION

- For each and every memory reference... address translation is performed
- Hardware transforms
 - Virtual address → physical address
- OS tracks which memory locations are free / in-use

February 15, 2017

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

L14.41

EXAMPLE: ADDRESS TRANSLATION

```
void func()
{
    int x=0;
    ...
    x = x + 3; // this is the line of code we are interested in
}
```

- Load value from memory
- Increment by three
- Store value back in memory
- In assembly...

February 15, 2017

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

L14.42

EXAMPLE: ADDRESS TRANSLATION - 2

128 : movl 0x0(%ebx), %eax ; load 0+ebx into eax
132 : addl \$0x03, %eax ; add 3 to eax register
135 : movl %eax, 0x0(%ebx) ; store eax back to mem

- Load value at address into register (eax)
- Add (3) to eax register
- Store the value of eax back into memory

February 15, 2017

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

L14.43

EXAMPLE: ADDRESS TRANSLATION - 3

- Fetch instruction at address 128
- Execute this instruction (load from address 15KB)
- Fetch instruction at address 132
- Execute this instruction (no memory reference)
- Fetch the instruction at address 135
- Execute this instruction (store to address 15 KB)

- Program's perspective:
 - Address space starts at 0
- Machine's perspective:
 - Program is located somewhere, not at 0

0KB 128
1KB 132
2KB 135
3KB
4KB
14KB
15KB 3000
16KB

movl 0x0(%ebx), %eax
addl 0x03, %eax
movl %eax, 0x0(%ebx)
↓
Heap
↓
(free)
↑
stack
↑
Int x Stack

February 15, 2017

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

L14.44

PLACEMENT IN PHYSICAL RAM

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

0KB
↓
Program Code
↓
Heap
↓
heap (free)
↓
stack
↑
Stack
16KB
Address Space

0KB
↓
Operating System
↓
16KB
↓
Code
↓
Heap
↓
(allocated but not in use)
↓
Stack
↑
48KB
↓
Stack
64KB
Physical Memory

Relocated Process

February 15, 2017

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

L14.45

BASE AND BOUNDS

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

physical address = virtual address + base

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

0 ≤ virtual address < bounds

February 15, 2017

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

L14.46

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

physical address = virtual address + base

0KB 128
1KB 132
2KB 135
3KB
4KB
14KB
15KB 3000
16KB

movl 0x0(%ebx), %eax
addl 0x03, %eax
movl %eax, 0x0(%ebx)
↓
Program Code
↓
Heap
↓
heap
↓
(free)
↑
stack
↑
Int x Stack

February 15, 2017

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

L14.47

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

February 15, 2017

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

L14.48

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.

February 15, 2017

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

L14.49

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

February 15, 2017

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

L14.50

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

February 15, 2017

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

L14.51

OS: WHEN PROCESS IS TERMINATED

- OS places memory back on the free list

February 15, 2017

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

L14.52

OS: WHEN CONTEXT SWITCH OCCURS

- OS must save base and bounds registers
 - Saved to the Process Control Block PCB (task_struct in Linux)

February 15, 2017

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

L14.53

DYNAMIC RELOCATION

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

February 15, 2017

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

L14.54

CHAPTER 16:
SEGMENTATION

February 22, 2017

TCSS422: Operating Systems [Winter 2017]
Institute of Technology, University of Washington - Tacoma

L15.55

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?

February 22, 2017

TCSS422: Operating Systems [Winter 2017]
Institute of Technology, University of Washington - Tacoma

L15.56

MULTIPLE SEGMENTS

- Memory segmentation
- Address space has (3) segments
 - Contiguous portions of address space
 - Logically separate segments for: code, stack, heap
- Each segment can placed separately
- Track base and bounds for each segment (registers)

February 22, 2017

TCSS422: Operating Systems [Winter 2017]
Institute of Technology, University of Washington - Tacoma

L15.57

SEGMENTS IN MEMORY

- Consider 3 segments:

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

February 22, 2017

TCSS422: Operating Systems [Winter 2017]
Institute of Technology, University of Washington - Tacoma

L15.58

ADDRESS TRANSLATION: CODE SEGMENT

$physical\ address = offset + base$

- Code segment - physically starts at 32KB (base)
- Starts at "0" in virtual address space

February 22, 2017

TCSS422: Operating Systems [Winter 2017]
Institute of Technology, University of Washington - Tacoma

L15.59

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)

February 22, 2017

TCSS422: Operating Systems [Winter 2017]
Institute of Technology, University of Washington - Tacoma

L15.60

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)

4KB
6KB
7KB
8KB

↓
Heap
(not in use)

Address Space

February 22, 2017

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

L15.61

SEGMENT REGISTERS

- Used to dereference memory during translation

13 12 11 10 9 8 7 6 5 4 3 2 1 0

SegmentOffset

■ First two bits identify segment type

■ Remaining bits identify memory offset

■ Example: virtual heap address 4200 (01000001101000)

13 12 11 10 9 8 7 6 5 4 3 2 1 0

0 1 0 0 0 0 0 1 1 0 1 0 0 0

SegmentOffset

Segment	bits
Code	00
Heap	01
Stack	10
-	11

February 22, 2017

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

L15.62

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

February 22, 2017

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

L15.63

STACK SEGMENT

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

26KB
28KB

(not in use)
↑
Stack
(not in use)

Physical Memory

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

February 22, 2017

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

L15.64

SHARED CODE SEGMENTS

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

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

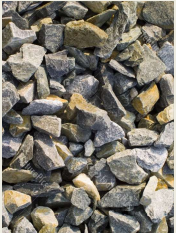
February 22, 2017

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

L15.65

SEGMENTATION GRANULARITY

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




February 22, 2017

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

L15.66

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



February 22, 2017

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

L15.67

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

February 22, 2017

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

L15.68

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

February 22, 2017

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

L15.69

QUESTIONS



February 15, 2017

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

L12.70