

## TCCS 422: OPERATING SYSTEMS

### Address Spaces and the Memory API

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



## FEEDBACK – 5/2

- Can we go over mutual exclusion again?
- Q#4 - mock midterm: Perfect-multi-tasking operating system
- Every process of the same priority will always receive exactly 1/nth of the available CPU time. Important CPU improvements for multi-tasking include: (1) fast context switching to enable jobs to be swapped in-and-out of the CPU very quickly, and (2) the use of a timer interrupt to preempt running jobs without the user voluntarily yielding the CPU. These innovations have enabled major improvements taking positive steps towards achieving a coveted "Perfect Multi-Tasking System".
- What are two challenges complicating the realization of a Perfect Multi-Tasking Operating System?

May 9, 2017

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

L11.2

## FEEDBACK – 2

- How can we wake up a specific thread with broadcast?
  - When we call broadcast every thread waiting is guaranteed to be awoken
  - We just don't know the order
  - Set a state variable or thread ID (integer), to allow some to proceed, others to go back to sleep

May 9, 2017

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

L11.3

## FEEDBACK - 3

- Can you review compare and swap?
- How come we no longer need locks because of this?
- On a single-core CPU the non-atomic compare-and-swap implementation is correct
- On multi-core the other cores can interrupt the non-atomic compare-and-swap code, resulting in mutual exclusion violation
- Compare-and-swap extends test-and-set to only "swap" (assign the lock to 1), if it is the expected value of 0.

May 9, 2017

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

L11.4

## COMPARE AND SWAP

### ■ Compare and Swap

```
1 int CompareAndSwap(int *ptr, int expected, int new) {
2     int actual = *ptr;
3     if (actual == expected)
4         *ptr = new;
5     return actual;
6 }
```

### ■ Spin lock usage

```
1 void lock(lock_t *lock) {
2     while (compareAndSwap(&lock->flag, 0, 1) == 1)
3         ; // spin
4 }
```

May 9, 2017

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

L11.5

## UPDATES

- Check quiz and tutorial grades
- Under review
  - Program 1
  - Midterm
- Program 2:
  - Due Sunday May 14

May 9, 2017

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

L11.6

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

May 9, 2017

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

L11.7

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 9, 2017

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

L11.8

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 9, 2017

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

L11.9

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 9, 2017

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

L11.10

EARLY MEMORY MANAGEMENT

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

0KB

64KB

Operating System  
(code, data, etc.)

Current Program  
(code, data, etc.)

max

Physical Memory

May 9, 2017

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

L11.11

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

Operating System  
(code, data, etc.)

Free

Process C  
(code, data, etc.)

Process B  
(code, data, etc.)

Free

Process A  
(code, data, etc.)

Free

Free

Physical Memory

May 9, 2017

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

L11.12

ADDRESS SPACE

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

May 9, 2017

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

L11.13

ADDRESS SPACE - 2

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

May 9, 2017

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

L11.14

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 9, 2017

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

L11.15

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 9, 2017

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

L11.16

VIRTUAL ADDRESSING - 2

- Output from 64-bit Linux:

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

May 9, 2017

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

L11.17

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 9, 2017

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

L11.18

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

May 9, 2017

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

L11.19

CHAPTER 14: THE MEMORY API



May 9, 2017

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

L11.20

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

- size() often used to ask the system how large a given datatype or struct is

May 9, 2017

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

L11.21

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));

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

4

40

May 9, 2017

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

L11.22

FREE()

```
#include <stdlib.h>

void free(void* ptr)
```

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

- Returns: nothing

May 9, 2017

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

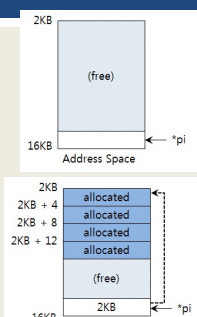
L11.23

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);
```



May 9, 2017

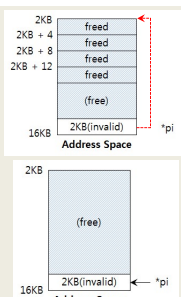
TCCS422: Operating Systems [Spring 2017]  
Institute of Technology, University of Washington - Tacoma

L11.24

VIRTUAL ADDRESS SPACE - 2

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

```
free(pi);
```



May 9, 2017

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

L11.25

COMMON MEMORY ERRORS

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

May 9, 2017

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

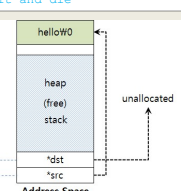
L11.26

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



May 9, 2017

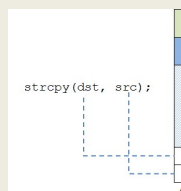
TCSS422: Operating Systems [Spring 2017]  
Institute of Technology, University of Washington - Tacoma

L11.27

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 ?



May 9, 2017

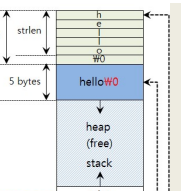
TCSS422: Operating Systems [Spring 2017]  
Institute of Technology, University of Washington - Tacoma

L11.28

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



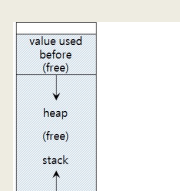
May 9, 2017

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

L11.29

FORGETTING TO INITIALIZE

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



May 9, 2017

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

L11.30

## MEMORY LEAK

May 9, 2017 TCS5422: Operating Systems [Spring 2017]  
Institute of Technology, University of Washington - Tacoma L11.31

## What will this code do?

```
#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;
}
```

32

## What will this code do?

```
#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;
}
```

**Output:**  
\$ ./pointer\_error  
The magic number is=53247  
The magic number is=11111

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

**Why?**

33

## 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).

May 9, 2017 TCS5422: Operating Systems [Spring 2017]  
Institute of Technology, University of Washington - Tacoma L11.34

## 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 9, 2017 TCS5422: Operating Systems [Spring 2017]  
Institute of Technology, University of Washington - Tacoma L11.35

## 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 9, 2017 TCS5422: Operating Systems [Spring 2017]  
Institute of Technology, University of Washington - Tacoma L11.36

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 9, 2017

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

L11.37

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 9, 2017

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

L11.38

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 9, 2017

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

L11.39

CHAPTER 15: ADDRESS TRANSLATION

May 9, 2017

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

L11.40

OBJECTIVES

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

May 9, 2017

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

L11.41

MEMORY VIRTUALIZATION

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

May 9, 2017

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

L11.42

### 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

May 9, 2017

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

L11.43

### 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...

May 9, 2017

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

L11.44

### 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

May 9, 2017

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

L11.45

### 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 1KB 2KB 3KB 4KB 14KB 15KB 16KB

Program Code  
Heap  
↓  
heap  
(free)  
↑  
stack  
Int x  
Stack

Relocated Process

0KB 16KB 32KB 48KB 64KB

Operating System  
(not in use)  
Code  
Heap  
(allocated but not in use)  
Stack  
(not in use)

Virtual mapping

Address Space

Physical Memory

May 9, 2017

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

L11.46

### PLACEMENT IN PHYSICAL RAM

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

0KB 16KB 32KB 48KB 64KB

Operating System  
(not in use)  
Code  
Heap  
(allocated but not in use)  
Stack  
(not in use)

Relocated Process

0KB 1KB 2KB 3KB 4KB 14KB 15KB 16KB

Program Code  
Heap  
↓  
heap  
(free)  
↑  
stack  
Stack

Virtual mapping

Address Space

Physical Memory

May 9, 2017

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

L11.47

### 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:

$0 \leq virtual\ address < bounds$

May 9, 2017

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

L11.48



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

0KB128132135

1KB

2KB

3KB

4KB

Program Code

Heap

(free)

stack

14KB

15KB

16KB

Int x

Stack

May 9, 2017

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

L11.49

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 9, 2017

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

L11.50

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 9, 2017

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

L11.51

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 9, 2017

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

L11.52

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 9, 2017

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

L11.53

OS: WHEN PROCESS IS TERMINATED

- OS places memory back on the free list

16KB

48KB

0KB

16KB

32KB

48KB

64KB

Operating System

(not in use)

Process A

(not in use)

Physical Memory

Free list

16KB

32KB

48KB

0KB

16KB

32KB

48KB

64KB

Operating System

(not in use)

(not in use)

(not in use)

Physical Memory

May 9, 2017

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

L11.54

Slides by Wes J. Lloyd

L11.9

OS: WHEN CONTEXT SWITCH OCCURS

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

The diagram illustrates the state of physical memory during a context switch. On the left, Process A is running, with its base register at 32KB and bounds at 48KB. On the right, after a context switch, Process B is running, with its base register at 48KB and bounds at 64KB. The OS is shown in the 0KB to 16KB range, and there is 16KB of unused space between the OS and the processes. The Process A PCB is shown with its base register at 32KB and bounds at 48KB.

May 9, 2017

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

L11.55

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!**

May 9, 2017

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

L11.56

CHAPTER 16:  
SEGMENTATION

May 9, 2017

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

L11.57

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?

The diagram shows the memory layout with Program Code at the top, followed by Heap, (free) space, and Stack. The stack grows downwards from 16KB, and the heap grows upwards from 6KB. The (free) space is between the heap and the stack.

May 9, 2017

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

L11.58

MULTIPLE SEGMENTS

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

May 9, 2017

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

L11.59

SEGMENTS IN MEMORY

- Consider 3 segments:

The diagram shows the memory layout with Operating System at the top, followed by (not in use) space, Stack, (not in use) space, Code, Heap, and (not in use) space. The stack grows downwards from 16KB, and the heap grows upwards from 6KB. The (not in use) space is between the OS and the stack, between the stack and the code, and between the code and the heap.

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

May 9, 2017

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

L11.60

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		16KB	

Bounds check:  
Is virtual address within 2KB address space?

or 32868 desired address

May 9, 2017TCS5422: Operating Systems [Spring 2017]  
Institute of Technology, University of Washington - TacomaL11.61

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	

104 + 34K or 34920 is the desired physical address

May 9, 2017TCS5422: Operating Systems [Spring 2017]  
Institute of Technology, University of Washington - TacomaL11.62

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)

May 9, 2017TCS5422: Operating Systems [Spring 2017]  
Institute of Technology, University of Washington - TacomaL11.63

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 9, 2017TCS5422: Operating Systems [Spring 2017]  
Institute of Technology, University of Washington - TacomaL11.64

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 9, 2017TCS5422: Operating Systems [Spring 2017]  
Institute of Technology, University of Washington - TacomaL11.65

STACK SEGMENT

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

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

May 9, 2017TCS5422: Operating Systems [Spring 2017]  
Institute of Technology, University of Washington - TacomaL11.66

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 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

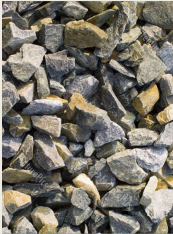
May 9, 2017

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

L11.67

SEGMENTATION GRANULARITY

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




May 9, 2017

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

L11.68

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 9, 2017

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

L11.69

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 9, 2017

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

L11.70

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	

May 9, 2017

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

L11.71

QUESTIONS



May 9, 2017

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

L11.72