


# TCSS 422: OPERATING SYSTEMS

## Condition Variables, Producer/Consumer

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

February 27, 2019

TCSS422: Operating Systems [Winter 2019]  
School of Engineering and Technology, University of Washington - Tacoma



## OBJECTIVES

- Assignment 2
- Assignment 3
- Parallel programming with P-threads cont'd
- Chapter 32 – Concurrency Problems
- Memory Virtualization
- Chapter 13 – Address Spaces
- Chapter 14 – Memory API
- Chapter 15 – Address Translation
- Chapter 16 – Segmentation
- Chapter 17 – Free Space Management
- Chapter 18 – Introduction to Paging
- Chapter 19 – Translation Lookaside Buffer (TLB)

February 27, 2019

TCSS422: Operating Systems [Winter 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L12.2

## FEEDBACK FROM 2/25

- Do producers create work assignments and consumers complete them?
- Are atomicity violation bugs just sharing variables without (forgetting) to use locks?
  - **YES**, atomicity violation bugs result when a program is multithreaded, and a variable has **many** uses, but a programmer forgets to apply locks to **all** of the uses
  - Atomicity violations result from the difficulty in realizing whether each use should be atomic or not
    - Some uses may not be acted upon by multiple threads
  - Locks may be added to the program, after variables are already used/defined

February 27, 2019

TCSS422: Operating Systems [Winter 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L12.3

## FEEDBACK - 2

- Confusion on locking producers and consumers correctly
  - Bounded buffer has global variable for capacity:
  - `int BOUNDED_BUFFER_SIZE;`
  - Producer: count must be less than `BOUNDED_BUFFER_SIZE`
    - Otherwise – call `pthread_cond_wait` on FILL signal
  - Consumer: count must be greater than 0
    - Otherwise – call `pthread_cond_wait` on EMPTY signal
    - Tells producer(s) to produce

February 27, 2019

TCSS422: Operating Systems [Winter 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L12.4

## FEEDBACK - 3

- End of program
  - Need to coordinate that we've produced/consumed `NUMBER_OF_MATRICIES` total
  - For loop is insufficient if multiple producers/consumers
  - Need a counter – (good use case for synchronized counter data structure)
  - 3 producers: once p1 produces matrix 1200, need to stop p2 and p3 from producing more...

February 27, 2019

TCSS422: Operating Systems [Winter 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L12.5



## CHAPTER 32 – CONCURRENCY PROBLEMS

February 27, 2019

TCSS422: Operating Systems [Winter 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L12.6

## OBJECTIVES

- **Chapter 32:**
  - Non-deadlock concurrency bugs
  - Deadlock causes
  - Deadlock prevention

February 27, 2019

TCSS422: Operating Systems [Winter 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L12.7

## CONCURRENCY BUGS IN OPEN SOURCE SOFTWARE

- **“Learning from Mistakes – A Comprehensive Study on Real World Concurrency Bug Characteristics”**
  - Shan Lu et al.
  - Architectural Support For Programming Languages and Operating Systems (ASPLOS 2008), Seattle WA

Application	What it does	Non-Deadlock	Deadlock
MySQL	Database Server	14	9
Apache	Web Server	13	4
Mozilla	Web Browser	41	16
Open Office	Office Suite	6	2
<b>Total</b>		<b>74</b>	<b>31</b>

February 27, 2019

TCSS422: Operating Systems [Winter 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L12.8

## NON-DEADLOCK BUGS

- Majority of concurrency bugs
- Most common:
  - Atomicity violation: forget to use locks
  - Order violation: failure to initialize lock/condition before use

February 27, 2019

TCSS422: Operating Systems [Winter 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L12.9

## ATOMICITY VIOLATION - MYSQL

- Two threads access the `proc_info` field in `struct thd`
- `NULL` is 0 in C
- Serialized access to shared memory among separate threads is not enforced (e.g. non-atomic)
- Simple example:

Programmer intended  
variable to be accessed  
atomically...



```
1  Thread1::  
2  if(thd->proc_info){  
3      ...  
4      fputs(thd->proc_info , ...);  
5      ...  
6  }  
7  
8  Thread2::  
9  thd->proc_info = NULL;
```

February 27, 2019

TCSS422: Operating Systems [Winter 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L12.10

## ATOMICITY VIOLATION - SOLUTION

- Add locks for all uses of: `thd->proc_info`

```
1 pthread_mutex_t lock = PTHREAD_MUTEX_INITIALIZER;
2
3 Thread1::
4 pthread_mutex_lock(&lock);
5 if(thd->proc_info){
6     ...
7     fputs(thd->proc_info , ...);
8     ...
9 }
10 pthread_mutex_unlock(&lock);
11
12 Thread2::
13 pthread_mutex_lock(&lock);
14 thd->proc_info = NULL;
15 pthread_mutex_unlock(&lock);
```

February 27, 2019

TCSS422: Operating Systems [Winter 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L12.11

## ORDER VIOLATION BUGS

- Desired order between memory accesses is flipped
- E.g. something is checked before it is set
- Example:

```
1 Thread1::
2 void init(){
3     mThread = PR_CreateThread(mMain, ...);
4 }
5
6 Thread2::
7 void mMain(...){
8     mState = mThread->State
9 }
```

- What if `mThread` is not initialized?

February 27, 2019

TCSS422: Operating Systems [Winter 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L12.12

## ORDER VIOLATION - SOLUTION

### ■ Use condition variable to enforce order

```
1 pthread_mutex_t mtLock = PTHREAD_MUTEX_INITIALIZER;
2 pthread_cond_t mtCond = PTHREAD_COND_INITIALIZER;
3 int mtInit = 0;
4
5 Thread1::
6 void init(){
7     ...
8     mThread = PR_CreateThread(mMain,...);
9
10    // signal that the thread has been created.
11    pthread_mutex_lock(&mtLock);
12    mtInit = 1;
13    pthread_cond_signal(&mtCond);
14    pthread_mutex_unlock(&mtLock);
15    ...
16 }
17
18 Thread2::
19 void mMain(...) {
20     ...
```

February 27, 2019

TCSS422: Operating Systems [Winter 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L12.13

## ORDER VIOLATION – SOLUTION 2

```
21 // wait for the thread to be initialized ...
22 pthread_mutex_lock(&mtLock);
23 while(mtInit == 0)
24     pthread_cond_wait(&mtCond, &mtLock);
25 pthread_mutex_unlock(&mtLock);
26
27 mState = mThread->State;
28 ...
29 }
```

February 27, 2019

TCSS422: Operating Systems [Winter 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L12.14

## NON-DEADLOCK BUGS - 1

- 97% of Non-Deadlock Bugs were
  - Atomicity
  - Order violations
- Consider what is involved in “spotting” these bugs in code
- Desire for automated tool support (IDE)

February 27, 2019

TCSS422: Operating Systems [Winter 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L12.15

## NON-DEADLOCK BUGS - 2

- Atomicity
  - How can we tell if a given variable is shared?
    - Can search the code for uses
  - How do we know if all instances of its use are shared?
    - Can some non-synchronized (non-atomic) uses be legal?
    - Before threads are created, after threads exit
    - Must verify the scope
- Order violation
  - Must consider all variable accesses
  - Must known desired order

February 27, 2019

TCSS422: Operating Systems [Winter 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L12.16




## DEADLOCK BUGS



- Presence of a cycle in code
- Thread 1 acquires lock L1, waits for lock L2
- Thread 2 acquires lock L2, waits for lock L1

```
Thread 1:      Thread 2:  
lock(L1);      lock(L2);  
lock(L2);      lock(L1);
```

- Both threads can block, unless one manages to acquire both locks



February 27, 2019

TCSS422: Operating Systems [Winter 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L12.17

## REASONS FOR DEADLOCKS

- Complex code
  - Must avoid circular dependencies – can be hard to find...
- Encapsulation hides potential locking conflicts
  - Easy-to-use APIs embed locks inside
  - Programmer doesn't know they are there
  - Consider the Java Vector class:

```
1  Vector v1,v2;  
2  v1.AddAll(v2);
```

- Vector is thread safe (synchronized) by design
- If there is a `v2.AddAll(v1);` call at nearly the same time deadlock could result

February 27, 2019

TCSS422: Operating Systems [Winter 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L12.18

## CONDITIONS FOR DEADLOCK

- Four conditions are required for dead lock to occur

Condition	Description
Mutual Exclusion	Threads claim exclusive control of resources that they require.
Hold-and-wait	Threads hold resources allocated to them while waiting for additional resources
No preemption	Resources cannot be forcibly removed from threads that are holding them.
Circular wait	There exists a circular chain of threads such that each thread holds one more resources that are being requested by the next thread in the chain

February 27, 2019

TCSS422: Operating Systems [Winter 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L12.19

## PREVENTION – MUTUAL EXCLUSION

- Build wait-free data structures
  - Eliminate locks altogether
  - Build structures using CompareAndSwap atomic CPU (HW) instruction
- C pseudo code for CompareAndSwap
- Hardware executes this code atomically

```
1  int CompareAndSwap(int *address, int expected, int new) {
2      if(*address == expected){
3          *address = new;
4          return 1; // success
5      }
6      return 0;
7  }
```

February 27, 2019

TCSS422: Operating Systems [Winter 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L12.20

## PREVENTION – MUTUAL EXCLUSION - 2

- Recall atomic increment

```
1 void AtomicIncrement(int *value, int amount){
2     do{
3         int old = *value;
4     }while( CompareAndSwap(value, old, old+amount)!=0);
5 }
```

- Compare and Swap tries over and over until successful
- CompareAndSwap is guaranteed to be atomic
- When it runs it is **ALWAYS** atomic (at HW level)

February 27, 2019

TCSS422: Operating Systems [Winter 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L12.21

## MUTUAL EXCLUSION: LIST INSERTION

- Consider list insertion

```
1 void insert(int value){
2     node_t * n = malloc(sizeof(node_t));
3     assert( n != NULL );
4     n->value = value ;
5     n->next = head;
6     head    = n;
7 }
```

February 27, 2019

TCSS422: Operating Systems [Winter 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L12.22

## MUTUAL EXCLUSION – LIST INSERTION - 2

### ■ Lock based implementation

```
1 void insert(int value){
2     node_t * n = malloc(sizeof(node_t));
3     assert( n != NULL );
4     n->value = value ;
5     lock(listlock); // begin critical section
6     n->next = head;
7     head = n;
8     unlock(listlock) ; //end critical section
9 }
```

February 27, 2019

TCSS422: Operating Systems [Winter 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L12.23

## MUTUAL EXCLUSION – LIST INSERTION - 3

### ■ Wait free (no lock) implementation

```
1 void insert(int value) {
2     node_t *n = malloc(sizeof(node_t));
3     assert(n != NULL);
4     n->value = value;
5     do {
6         n->next = head;
7     } while (!CompareAndSwap(&head, n->next, n));
8 }
```

- Assign &head to n (new node ptr)
- Only when head = n->next

February 27, 2019

TCSS422: Operating Systems [Winter 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L12.24

## CONDITIONS FOR DEADLOCK

- Four conditions are required for dead lock to occur

Condition	Description
Mutual Exclusion	Threads claim exclusive control of resources that they require.
Hold-and-wait	Threads hold resources allocated to them while waiting for additional resources
No preemption	Resources cannot be forcibly removed from threads that are holding them.
Circular wait	There exists a circular chain of threads such that each thread holds one more resources that are being requested by the next thread in the chain

February 27, 2019

TCSS422: Operating Systems [Winter 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L12.25

## PREVENTION LOCK – HOLD AND WAIT

- Problem: acquire all locks atomically
- Solution: use a “lock” “lock”... (*like a guard lock*)

```
1 lock(prevention);  
2 lock(L1);  
3 lock(L2);  
4 ...  
5 unlock(prevention);
```

- Effective solution – guarantees no race conditions while acquiring L1, L2, etc.
- Order doesn't matter for L1, L2
- Prevention (GLOBAL) lock decreases concurrency of code
  - Acts Lowers lock granularity
- Encapsulation: consider the Java Vector class...

February 27, 2019

TCSS422: Operating Systems [Winter 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L12.26

## CONDITIONS FOR DEADLOCK

- Four conditions are required for dead lock to occur

Condition	Description
Mutual Exclusion	Threads claim exclusive control of resources that they require.
Hold-and-wait	Threads hold resources allocated to them while waiting for additional resources
➡ No preemption	Resources cannot be forcibly removed from threads that are holding them.
Circular wait	There exists a circular chain of threads such that each thread holds one more resources that are being requested by the next thread in the chain

February 27, 2019

TCSS422: Operating Systems [Winter 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L12.27

## PREVENTION – NO PREEMPTION

- When acquiring locks, don't BLOCK forever if unavailable...
- `pthread_mutex_trylock()` - try once
- `pthread_mutex_timedlock()` - try and wait awhile

```
1  top:
2      lock(L1);
3      if( tryLock(L2) == -1 ){
4          unlock(L1);
5          goto top;
6      }
```



- Eliminates deadlocks

February 27, 2019

TCSS422: Operating Systems [Winter 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L12.28

## NO PREEMPTION – LIVENLOCKS PROBLEM

- Can lead to livelock

```
1  top:
2      lock(L1);
3      if( tryLock(L2) == -1 ){
4          unlock(L1);
5          goto top;
6      }
```

- Two threads execute code in parallel →  
always fail to obtain both locks

- Fix: add random delay

- Allows one thread to win the livelock race!



February 27, 2019

TCSS422: Operating Systems [Winter 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L12.29

## CONDITIONS FOR DEADLOCK

- Four conditions are required for dead lock to occur

Condition	Description
Mutual Exclusion	Threads claim exclusive control of resources that they require.
Hold-and-wait	Threads hold resources allocated to them while waiting for additional resources
No preemption	Resources cannot be forcibly removed from threads that are holding them.
→ Circular wait	There exists a circular chain of threads such that each thread holds one more resources that are being requested by the next thread in the chain

February 27, 2019

TCSS422: Operating Systems [Winter 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L12.30

## PREVENTION – CIRCULAR WAIT

- Provide total ordering of lock acquisition throughout code
  - Always acquire locks in same order
  - L1, L2, L3, ...
  - Never mix: L2, L1, L3; L2, L3, L1; L3, L1, L2....
- Must carry out same ordering through entire program

February 27, 2019

TCSS422: Operating Systems [Winter 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L12.31

## DEADLOCK AVOIDANCE VIA INTELLIGENT SCHEDULING

- Consider a smart scheduler
  - Scheduler knows which locks threads use
- Consider this scenario:
  - 4 Threads (T1, T2, T3, T4)
  - 2 Locks (L1, L2)
- Lock requirements of threads:

	T1	T2	T3	T4
L1	yes	yes	no	no
L2	yes	yes	yes	no

February 27, 2019

TCSS422: Operating Systems [Winter 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L12.32



INTELLIGENT SCHEDULING - 2

■ Scheduler produces schedule:

CPU 1	T3	T4
CPU 2	T1	T2

■ No deadlock can occur

■ Consider:

	T1	T2	T3	T4
L1	yes	yes	yes	no
L2	yes	yes	yes	no

February 27, 2019

TCSS422: Operating Systems [Winter 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L12.33

INTELLIGENT SCHEDULING - 3

■ Scheduler produces schedule

CPU 1	T4		
CPU 2	T1	T2	T3

■ Scheduler must be conservative and not take risks

■ Slows down execution – many threads

■ There has been limited use of these approaches given the difficulty having intimate lock knowledge about every thread

February 27, 2019

TCSS422: Operating Systems [Winter 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L12.34

## DETECT AND RECOVER

- Allow deadlock to occasionally occur and then take some action.
  - Example: When OS freezes, reboot...
- How often is this acceptable?
  - Once per year
  - Once per month
  - Once per day
  - *Consider the effort tradeoff of finding every deadlock bug*
- Many database systems employ deadlock detection and recovery techniques.

February 27, 2019

TCSS422: Operating Systems [Winter 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L12.35

## CHAPTER 13: ADDRESS SPACES



February 27, 2019

TCSS422: Operating Systems [Winter 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L12.36

## OBJECTIVES – MEMORY VIRTUALIATION

- 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 27, 2019

TCSS422: Operating Systems [Winter 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L12.37

## 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 27, 2019

TCSS422: Operating Systems [Winter 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L12.38

## 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 27, 2019

TCSS422: Operating Systems [Winter 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L12.39

## 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 27, 2019

TCSS422: Operating Systems [Winter 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L12.40

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.)

February 27, 2019

TCSS422: Operating Systems [Winter 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L12.41

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

February 27, 2019

TCSS422: Operating Systems [Winter 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L12.42


## ADDRESS SPACE

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

- Main elements:

- Program code
- Stack
- Heap

- Example: 16KB address space



February 27, 2019

TCSS422: Operating Systems [Winter 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L12.43

## ADDRESS SPACE - 2

- Code


- Program code

- Stack

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

- Heap

- Dynamic storage
- Malloc() new()




February 27, 2019

TCSS422: Operating Systems [Winter 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L12.44

## 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



February 27, 2019

TCSS422: Operating Systems [Winter 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L12.45

## 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 27, 2019

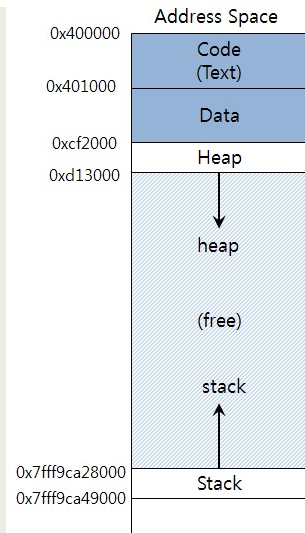
TCSS422: Operating Systems [Winter 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L12.46

## VIRTUAL ADDRESSING - 2

### ■ Output from 64-bit Linux:

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



February 27, 2019

TCSS422: Operating Systems [Winter 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L12.47

## 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 27, 2019

TCSS422: Operating Systems [Winter 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L12.48



## 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 evaluating memory virtualization schemes***

February 27, 2019

TCSS422: Operating Systems [Winter 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L12.49

## CHAPTER 14: THE MEMORY API



February 27, 2019

TCSS422: Operating Systems [Winter 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L12.50

## 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 27, 2019

TCSS422: Operating Systems [Winter 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L12.51

## 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 27, 2019

TCSS422: Operating Systems [Winter 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L12.52

## FREE()

```
#include <stdlib.h>

void free(void* ptr)
```

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

February 27, 2019

TCSS422: Operating Systems [Winter 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L12.53

```
#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;
}
```

54

```
#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?

55

## 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”

February 27, 2019

TCSS422: Operating Systems [Winter 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L12.57

## 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 27, 2019

TCSS422: Operating Systems [Winter 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L12.58

## 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 27, 2019


TCCS422: Operating Systems [Winter 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L12.59

## 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 27, 2019

TCCS422: Operating Systems [Winter 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L12.60

## 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 27, 2019

TCSS422: Operating Systems [Winter 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L12.61

## QUESTIONS

