


TCSS 422: OPERATING SYSTEMS

Concurrency Problems, Introduction to memory virtualization

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

May 12, 2020

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



OBJECTIVES – 5/12

- **Questions from 5/7**
- Midterm review
- Assignment 2 (based on Ch. 30)
- **Chapter 32: Concurrency Problems**
 - Deadlock causes
 - Deadlock prevention
- **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

May 12, 2020

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

L12.2

MATERIAL / PACE

- Please classify your perspective on material covered in today’s class (44 respondents):
 - 1-mostly review, 5-equal new/review, 10-mostly new
 - **Average – 6.87 (↑ from 6.27)**
- Please rate the pace of today’s class:
 - 1-slow, 5-just right, 10-fast
 - **Average – 5.81 (↑ from 5.77)**

May 12, 2020

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

L12.3

FEEDBACK FROM 5/7

May 12, 2020

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

L12.4

OBJECTIVES – 5/12

- Questions from 5/7
- Midterm review
- Assignment 2 (based on Ch. 30)
- Chapter 32: Concurrency Problems
 - Deadlock causes
 - Deadlock prevention
- 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

May 12, 2020

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

L12.5

MIDTERM RESULTS

Average: 85.41
Mode: 95.5
Median: 88.3
1st quartile:
0 to 79.5
2nd quartile:
79.5 to 88.3
3rd quartile:
88.3 to 94.5
4th quartile:
94.5 to 99.5
Std. dev:
10.96

Distribution is log normal

Score Range	# of students
100-100	2
97-10-94	17
94-10-91	7
91-10-88	9
88-10-85	6
85-10-82	4
82-10-79	6
79-10-76	5
76-10-73	4
73-10-70	0
70-10-67	1
67-10-64	2
64-10-61	2
61-10-58	0
58-10-55	0
55-10-52	2

May 12, 2020

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

L12.6

NOTES ON SCORING

- **Questions 5, 6, 10, 11:**
- When I transferred this question to Canvas, I changed job D's arrival time from T=20 to T=10
- My solution accidentally did not capture this change
- Initial answers in Canvas incorrectly scored these questions
- Scores have been updated, please verify scoring:
- **Question 5:** 25 (turnaround time job D) +1
- **Question 6:** 28.8 (average turnaround time all jobs) +1
- **Question 10:** 20 (response time job D) +1
- **Question 11:** 16.3 (average response time all jobs) +1

May 12, 2020

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

L12.7

NOTES ON SCORING - 2

- **Question 34:**
- MLFQ Scheduler Question
- How many jobs execute in the medium priority queue?
- This question was simply asking of jobs A, B, C, how many of them ran at some point in the medium priority queue
(this question can be answered using only the starter graph)
- Answer: 3
- I have also accepted as an answer to the questions, the total number of timer units jobs ran in the medium priority queue
- Second answer: 24 (ok)

May 12, 2020

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

L12.8

NOTES ON SCORING - 3

- **Question 48:**
- Classify the API:
- **pthread_mutex_timelock()**
- This API is both **BLOCKING** (for a while) and **NON-BLOCKING**, so the best answer was “all of the above”
- Partial credit for **BLOCKING** or **NON-BLOCKING**

May 12, 2020

TCS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L12.9

NOTES ON SCORING - 4

- All other questions have been reviewed for partial credit, and Canvas grading interpretation
- Canvas autograder is sensitive and wants exact answers
- Where answers were not recognized, I have manually added points
- Please review scoring for correctness, and notify of any issues / questions
- Midterm questions ?
- Special session - midterm review:
recording LIVE on Wed May 13 @ 5:50p

May 12, 2020

TCS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L12.10


OBJECTIVES – 5/12

- Questions from 5/7
- Midterm review
- Assignment 2 (based on Ch. 30)
- Chapter 32: Concurrency Problems
 - Deadlock causes
 - Deadlock prevention
- 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

May 12, 2020

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

L12.11



CHAPTER 32 –
CONCURRENCY
PROBLEMS

May 12, 2020

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

L12.12

OBJECTIVES – 5/12

- Questions from 5/7
- Midterm review
- Assignment 2 (based on Ch. 30)
- Chapter 32: Concurrency Problems
 - Deadlock causes
 - Deadlock prevention
- 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

May 12, 2020

TCS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L12.13

NON-DEADLOCK BUGS

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

May 12, 2020

TCS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L12.14

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);
```

May 12, 2020

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

L12.15

ORDER VIOLATION - SOLUTION

- Use condition & signal 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 Thread 1::
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    ...
```

May 12, 2020

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

L12.16

ORDER VIOLATION – SOLUTION - 2

- Use condition & signal to enforce order

```
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 }
```

May 12, 2020

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

L12.17


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



May 12, 2020

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

L12.18

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

May 12, 2020

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

L12.19

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

May 12, 2020

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

L12.20

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  }
```

May 12, 2020

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

L12.21

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)

May 12, 2020

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

L12.22

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 }
```

May 12, 2020

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

L12.23

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 }
```

May 12, 2020

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

L12.24

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

May 12, 2020

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

L12.25

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

May 12, 2020

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

L12.26

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...

May 12, 2020

TCS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L12.27

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

May 12, 2020

TCS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L12.28

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

May 12, 2020

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

L12.29

NO PREEMPTION – LIVELOCKS 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!



May 12, 2020

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

L12.30

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

May 12, 2020

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

L12.31

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

May 12, 2020

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

L12.32

CONDITIONS FOR DEADLOCK

- If any of the following conditions DOES NOT EXSIST, describe why deadlock can not 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

May 12, 2020

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

L12.33

The dining philosophers problem where 5 philosophers compete for 5 forks, and where a philosopher must hold two forks to eat involves which deadlock condition(s)?

Mutual Exclusion

Hold-and-wait

No preemption

Circular wait

All of the above

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

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

May 12, 2020

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

L12.35

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


May 12, 2020

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

L12.36

INTELLIGENT SCHEDULING - 3

- Scheduler produces schedule



- 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

May 12, 2020

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

L12.37

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.

May 12, 2020


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

L12.38

TCSS 422 WILL RETURN AT ~2:48PM

May 12, 2020

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




L12.3
9

CHAPTER 13: ADDRESS SPACES

May 12, 2020

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



L12.40

OBJECTIVES – 5/12

- Questions from 5/7
- Midterm review
- Assignment 2 (based on Ch. 30)
- Chapter 32: Concurrency Problems
 - Deadlock causes
 - Deadlock prevention
- 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

May 12, 2020

TCS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L12.41

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 12, 2020

TCS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L12.42

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 12, 2020

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

L12.43

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 12, 2020

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

L12.44

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 12, 2020

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

L12.45

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 12, 2020

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

L12.46


ADDRESS SPACE

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

- Main elements:

- Program code
- Stack
- Heap

- Example: 16KB address space



May 12, 2020

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

L12.47

ADDRESS SPACE - 2

- Code


- Program code

- Stack

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

- Heap

- Dynamic storage
- Malloc() new()




May 12, 2020

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

L12.48

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 12, 2020

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

L12.49

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 12, 2020


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

L12.50

VIRTUAL ADDRESSING - 2

■ Output from 64-bit Linux:

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



May 12, 2020

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

L12.51

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 12, 2020

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

L12.52

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***

May 12, 2020

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

L12.53

CHAPTER 14: THE MEMORY API



May 12, 2020

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

L12.54

OBJECTIVES – 5/12

- Questions from 5/7
- Midterm review
- Assignment 2 (based on Ch. 30)
- Chapter 32: Concurrency Problems
 - Deadlock causes
 - Deadlock prevention
- 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

May 12, 2020

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

L12.55

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 12, 2020

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

L12.56

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 12, 2020

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

L12.57

FREE()

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

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

May 12, 2020

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

L12.58

```
#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;
}
```

59

```
#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;
}
```

Output:

```
$ ./pointer_error
The magic number is=53247
The magic number is=11111
```

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

Why?

60

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 12, 2020

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

L12.61

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 12, 2020

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

L12.62

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 12, 2020

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

L12.63

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 12, 2020


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

L12.64

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 12, 2020

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

L12.65

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 12, 2020

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

L12.66

CHAPTER 15: ADDRESS TRANSLATION



May 12, 2020

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

L12.67

OBJECTIVES – 5/12

- Questions from 5/7
- Midterm review
- Assignment 2 (based on Ch. 30)
- Chapter 32: Concurrency Problems
 - Deadlock causes
 - Deadlock prevention
- 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

May 12, 2020

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

L12.68

CH. 15: OBJECTIVES

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

May 12, 2020

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

L12.69

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 'Virtual mapping' section shows the 'Address Space' with segments: Program Code (0KB-16KB), Heap (16KB-32KB), heap (free) (32KB-48KB), stack (48KB-64KB), and Stack (64KB-16KB). On the right, the 'Physical Memory' section shows the layout: Operating System (0KB-16KB), (not in use) (16KB-32KB), Code (32KB-36KB), Heap (36KB-40KB), (allocated but not in use) (40KB-48KB), Stack (48KB-64KB), and (not in use) (64KB-16KB). A bracket on the right indicates the 'Relocated Process' spanning from 32KB to 64KB. Dashed lines show the mapping from the virtual address space to the physical memory.

May 12, 2020

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

L12.70

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 12, 2020

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

L12.71

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

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

0KB128
1KB132
135

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

Program Code

2KB

3KB

4KB

Heap

↓

heap

(free)

↑

stack

14KB

15KB3000

16KB

Int x
Stack

May 12, 2020

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

L12.72

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 12, 2020

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

L12.73

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 12, 2020

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

L12.74

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 12, 2020

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

L12.75

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 12, 2020

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

L12.76

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 12, 2020

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

L12.77

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 12, 2020

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

L12.78

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 12, 2020


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

L12.79

QUESTIONS



TCSS 422 OFFICE HOURS



WILL RETURN IN A FEW MINUTES

