


TCSS 422: OPERATING SYSTEMS

Condition Variables II, Concurrency Problems



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

May 9, 2024 TCSS422: Operating Systems [Spring 2024]
School of Engineering and Technology, University of Washington Tacoma

1

OFFICE HOURS – THURSDAY

- SPECIAL TIME THURSDAY MAY 9
- Thursday Office Hours
 - 6:00pm to 7:00 pm – CP 229 and Zoom

May 9, 2024 TCSS422: Operating Systems [Spring 2024]
School of Engineering and Technology, University of Washington - Tacoma L13.2

2

OBJECTIVES – 5/9

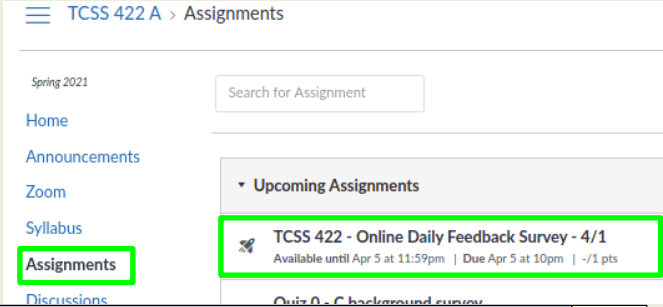
- **Questions from 5/9**
- Pthread Tutorial-May 24 / Assignment 2 posted next week
- Quiz 3 – Synchronized Array (class activity-next week)
- Chapter 30: Condition Variables
 - Covering Conditions
- Chapter 32: Concurrency Problems
 - Non-deadlock concurrency bugs
 - Deadlock causes
 - Deadlock prevention
- Chapter 13: Address Spaces
- Chapter 14: The Memory API

May 9, 2024	TCSS422: Operating Systems [Spring 2024] School of Engineering and Technology, University of Washington - Tacoma	L13.3
-------------	---	-------

3

ONLINE DAILY FEEDBACK SURVEY

- Daily Feedback Quiz in Canvas – Available After Each Class
- Extra credit available for completing surveys **ON TIME**
- Tuesday surveys: due by ~ Wed @ 11:59p
- Thursday surveys: due ~ Mon @ 11:59p



May 9, 2024	TCSS422: Computer Operating Systems [Spring 2024] School of Engineering and Technology, University of Washington - Tacoma	L13.4
-------------	--	-------

4

TCSS 422 - Online Daily Feedback Survey - 4/1

Quiz Instructions

Question 1 0.5 pts

On a scale of 1 to 10, please classify your perspective on material covered in today's class:

1 2 3 4 5 6 7 8 9 10

Mostly Review To Me Equal New and Review Mostly New To Me

Question 2 0.5 pts

Please rate the pace of today's class:

1 2 3 4 5 6 7 8 9 10

Slow Just Right Fast

May 9, 2024 TCSS422: Computer Operating Systems [Spring 2024] School of Engineering and Technology, University of Washington - Tacoma L13.5

5

MATERIAL / PACE

- Please classify your perspective on material covered in today's class (26 respondents):
- 1-mostly review, 5-equal new/review, 10-mostly new
- **Average - 6.58 (↑ - previous 6.56)**

- Please rate the pace of today's class:
- 1-slow, 5-just right, 10-fast
- **Average - 5.31 (no change - previous 5.31)**

May 9, 2024 TCSS422: Computer Operating Systems [Spring 2024] School of Engineering and Technology, University of Washington - Tacoma L13.6

6

FEEDBACK FROM 5/7

- **Why does the (bounded) buffer consist of only 1 element?**
 - In the textbook example, the initial bounded buffer example is a single integer
 - Think of this as an integer array of 1 element
 - This is initially done for simplicity in teaching the bounded buffer concepts
 - By not having an array, it is not necessary to track the index for where elements are added (**fill index**), and where elements are removed (**use index**) from the bounded buffer
- **Why isn't the buffer an array of multiple pointers that stores data in a FIFO order?**
 - This can be done. For example, we could have an array of matrices, where each matrix is a 2-D array of integers on the heap

May 9, 2024

TCSS422: Operating Systems [Spring 2024]
School of Engineering and Technology, University of Washington - Tacoma

L13.7

7

FEEDBACK - 2

- **Is the buffer in the stack or in the heap?**
 - In `signal.c` (chapter 30), the bounded buffer is a matrix pointer which is defined as a global variable in the program's data segment

```
int ** bigmatrix;
```
 - Globals are not on the stack or heap, but in the data segment
- **How are the memory addresses of each element in a matrix assigned (in `signal.c`)?**
 - In `GenMatrix()`, the 2-D matrix is represented as an array of integer arrays
 - `int *mm` is made to point at a row: `matrix[i]`

```
int * mm = matrix[i];
```
 - Then, we access the j^{th} element of the row to assign the column value

```
mm[j] = rand() % ELEMENT_SIZE;
```

May 9, 2024

TCSS422: Operating Systems [Spring 2024]
School of Engineering and Technology, University of Washington - Tacoma

L13.8

8

OBJECTIVES – 5/9

- Questions from 5/9
- **Pthread Tutorial-May 24 / Assignment 2 posted next week**
- Quiz 3 – Synchronized Array (class activity-next week)
- Chapter 30: Condition Variables
 - Covering Conditions
- Chapter 32: Concurrency Problems
 - Non-deadlock concurrency bugs
 - Deadlock causes
 - Deadlock prevention
- Chapter 13: Address Spaces
- Chapter 14: The Memory API

May 9, 2024	TCSS422: Operating Systems [Spring 2024] School of Engineering and Technology, University of Washington - Tacoma	L13.9
-------------	---	-------

9

OBJECTIVES – 5/9

- Questions from 5/9
- Pthread Tutorial-May 24 / **Assignment 2 posted next week**
- Quiz 3 – Synchronized Array (class activity-next week)
- Chapter 30: Condition Variables
 - Covering Conditions
- Chapter 32: Concurrency Problems
 - Non-deadlock concurrency bugs
 - Deadlock causes
 - Deadlock prevention
- Chapter 13: Address Spaces
- Chapter 14: The Memory API

May 9, 2024	TCSS422: Operating Systems [Spring 2024] School of Engineering and Technology, University of Washington - Tacoma	L13.10
-------------	---	--------


10

OBJECTIVES – 5/9

- Questions from 5/9
- Pthread Tutorial-May 24 / Assignment 2 posted next week
- **Quiz 3 – Synchronized Array (class activity-next week)**
- Chapter 30: Condition Variables
 - Covering Conditions
- Chapter 32: Concurrency Problems
 - Non-deadlock concurrency bugs
 - Deadlock causes
 - Deadlock prevention
- Chapter 13: Address Spaces
- Chapter 14: The Memory API

May 9, 2024	TCSS422: Operating Systems [Spring 2024] School of Engineering and Technology, University of Washington - Tacoma	L13.11
-------------	---	--------

11



CHAPTER 30 – CONDITION VARIABLES

May 9, 2024	TCSS422: Operating Systems [Spring 2024] School of Engineering and Technology, University of Washington - Tacoma	L13.1 2
-------------	---	------------

12

OBJECTIVES – 5/9

- Questions from 5/9
- Pthread Tutorial-May 24 / Assignment 2 posted next week
- Quiz 3 – Synchronized Array (class activity-next week)
- Chapter 30: Condition Variables
 - **Covering Conditions**
- Chapter 32: Concurrency Problems
 - Non-deadlock concurrency bugs
 - Deadlock causes
 - Deadlock prevention
- Chapter 13: Address Spaces
- Chapter 14: The Memory API

May 9, 2024	TCSS422: Operating Systems [Spring 2024] School of Engineering and Technology, University of Washington - Tacoma	L13.13
-------------	---	--------

13

COVERING CONDITIONS

- A condition that covers **all** cases (conditions):
- Excellent use case for **pthread_cond_broadcast**

- Consider memory allocation:
 - When a program deals with huge memory allocation/deallocation on the heap
 - Access to the heap must be managed when memory is scarce

- PREVENT: Out of memory:
 - queue requests until memory is free

- Which thread should be woken up?

May 9, 2024	TCSS422: Operating Systems [Spring 2024] School of Engineering and Technology, University of Washington - Tacoma	L13.14
-------------	---	--------

14

COVERING CONDITIONS - 2

```
1 // how many bytes of the heap are free?
2 int bytesLeft = MAX_HEAP_SIZE;
3
4 // need lock and condition too
5 cond_t c;
6 mutex_t m;
7
8 void *
9 allocate(int size) {
10 pthread_mutex_lock(&m);
11 while (bytesLeft < size)
12 pthread_cond_wait(&c, &m);
13 void *ptr = ...; // get mem from heap
14 bytesLeft -= size;
15 pthread_mutex_unlock(&m);
16 return ptr;
17 }
18
19 void free(void *ptr, int size) {
20 pthread_mutex_lock(&m);
21 bytesLeft += size;
22 pthread_cond_signal(&c); // Broadcast
23 pthread_mutex_unlock(&m);
24 }
```

May 9, 2024 TCSS422: Operating Systems [Spring 2024] L13.15
School of Engineering and Technology, University of Washington - Tacoma

15

COVER CONDITIONS - 3


- Broadcast awakens all blocked threads requesting memory
- Each thread evaluates if there's enough memory: (bytesLeft < size)
 - Reject: requests that cannot be fulfilled- go back to sleep
 - *Insufficient memory*
 - Run: requests which can be fulfilled
 - with newly available memory!
- **Another use case:** coordinate a group of busy threads to gracefully end, to EXIT the program
- **Overhead**
 - Many threads may be awoken which can't execute

May 9, 2024 TCSS422: Operating Systems [Spring 2024] L13.16
School of Engineering and Technology, University of Washington - Tacoma

16

CHAPTER 31: SEMAPHORES

- Offers a combined C language construct that can assume the role of a lock or a condition variable depending on usage
 - Allows fewer concurrency related variables in your code
 - Potentially makes code more ambiguous
 - For this reason, with limited time in a 10-week quarter, we do not cover semaphores in TCSS 422
- **Ch. 31.6 – Dining Philosophers Problem**
 - Classic computer science problem about sharing eating utensils
 - Each philosopher tries to obtain two forks in order to eat
 - Mimics deadlock as there are not enough forks
 - Solution is to have one left-handed philosopher that grabs forks in opposite order



May 9, 2024	TCSS422: Operating Systems [Spring 2024] School of Engineering and Technology, University of Washington - Tacoma	L13.17
-------------	---	--------


17

OBJECTIVES – 5/9

- Questions from 5/9
- Pthread Tutorial-May 24 / Assignment 2 posted next week
- Quiz 3 – Synchronized Array (class activity-next week)
- Chapter 30: Condition Variables
 - Producer/Consumer
 - Covering Conditions
 - **Chapter 32: Concurrency Problems**
 - Non-deadlock concurrency bugs
 - Deadlock causes
 - Deadlock prevention
- Chapter 13: Address Spaces
- Chapter 14: The Memory API

May 9, 2024	TCSS422: Operating Systems [Spring 2024] School of Engineering and Technology, University of Washington - Tacoma	L13.18
-------------	---	--------

18



CHAPTER 32 – CONCURRENCY PROBLEMS

May 9, 2024

TCSS422: Operating Systems [Spring 2024]
School of Engineering and Technology, University of Washington - Tacoma

L13.19

19

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
Total		74	31

May 9, 2024

TCSS422: Operating Systems [Spring 2024]
School of Engineering and Technology, University of Washington - Tacoma

L13.20

20

OBJECTIVES – 5/9

- Questions from 5/9
- Pthread Tutorial-May 24 / Assignment 2 posted next week
- Quiz 3 – Synchronized Array (class activity-next week)
- Chapter 30: Condition Variables
 - Covering Conditions
- Chapter 32: Concurrency Problems
 - **Non-deadlock concurrency bugs**
 - Deadlock causes
 - Deadlock prevention
- Chapter 13: Address Spaces
- Chapter 14: The Memory API

May 9, 2024	TCSS422: Operating Systems [Spring 2024] School of Engineering and Technology, University of Washington - Tacoma	L13.21
-------------	---	--------

21

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 9, 2024	TCSS422: Operating Systems [Spring 2024] School of Engineering and Technology, University of Washington - Tacoma	L13.22
-------------	---	--------

22

ATOMICITY VIOLATION - MYSQL

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

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;
```

May 9, 2024	TCSS422: Operating Systems [Spring 2024] School of Engineering and Technology, University of Washington - Tacoma	L13.23
-------------	---	--------

23

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 9, 2024	TCSS422: Operating Systems [Spring 2024] School of Engineering and Technology, University of Washington - Tacoma	L13.24
-------------	---	--------

24

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?

May 9, 2024

TCSS422: Operating Systems [Spring 2024]
School of Engineering and Technology, University of Washington - Tacoma

L13.25

25

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 9, 2024

TCSS422: Operating Systems [Spring 2024]
School of Engineering and Technology, University of Washington - Tacoma

L13.26

26

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 9, 2024

TCSS422: Operating Systems [Spring 2024]
School of Engineering and Technology, University of Washington - Tacoma

L13.27

27

NON-DEADLOCK BUGS - 1

- 97% of Non-Deadlock Bugs were
 - Atomicity
 - Order violations
- Consider what is involved in “spotting” these bugs in code
 - >> *no use of locking constructs to search for*
- Desire for automated tool support (IDE)

May 9, 2024

TCSS422: Operating Systems [Spring 2024]
School of Engineering and Technology, University of Washington - Tacoma

L13.28

28

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?
 - Legal uses: before threads are created, after threads exit
 - Must verify the scope
- **Order violation**
 - Must consider all variable accesses
 - Must know desired order

May 9, 2024	TCSS422: Operating Systems [Spring 2024] School of Engineering and Technology, University of Washington - Tacoma	L13.29
-------------	---	--------

29


WE WILL RETURN AT 4:50PM



May 9, 2024	TCSS422: Operating Systems [Spring 2024] School of Engineering and Technology, University of Washington - Tacoma	L13.30
-------------	---	--------

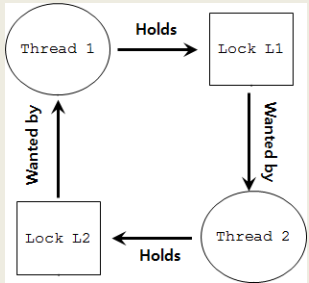
30

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 9, 2024	TCSS422: Operating Systems [Spring 2024] School of Engineering and Technology, University of Washington - Tacoma	L13.31
-------------	---	--------

31

OBJECTIVES – 5/9

- Questions from 5/9
- Pthread Tutorial-May 24 / Assignment 2 posted next week
- Quiz 3 – Synchronized Array (class activity-next week)
- Chapter 30: Condition Variables
 - Covering Conditions
- Chapter 32: Concurrency Problems
 - Non-deadlock concurrency bugs
 - **Deadlock causes**
 - Deadlock prevention
- Chapter 13: Address Spaces
- Chapter 14: The Memory API

May 9, 2024	TCSS422: Operating Systems [Spring 2024] School of Engineering and Technology, University of Washington - Tacoma	L13.32
-------------	---	--------

32

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 9, 2024
TCSS422: Operating Systems [Spring 2024]
 School of Engineering and Technology, University of Washington - Tacoma
L13.33

33

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 9, 2024
TCSS422: Operating Systems [Spring 2024]
 School of Engineering and Technology, University of Washington - Tacoma
L13.34

34

OBJECTIVES – 5/9

- Questions from 5/9
- Pthread Tutorial-May 24 / Assignment 2 posted next week
- Quiz 3 – Synchronized Array (class activity-next week)
- Chapter 30: Condition Variables
 - Covering Conditions
- Chapter 32: Concurrency Problems
 - Non-deadlock concurrency bugs
 - Deadlock causes
 - **Deadlock prevention**
- Chapter 13: Address Spaces
- Chapter 14: The Memory API

May 9, 2024	TCSS422: Operating Systems [Spring 2024] School of Engineering and Technology, University of Washington - Tacoma	L13.35
-------------	---	--------

35

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 9, 2024	TCSS422: Operating Systems [Spring 2024] School of Engineering and Technology, University of Washington - Tacoma	L13.36
-------------	---	--------

36

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 9, 2024

TCSS422: Operating Systems [Spring 2024]
School of Engineering and Technology, University of Washington - Tacoma

L13.37

37

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 9, 2024

TCSS422: Operating Systems [Spring 2024]
School of Engineering and Technology, University of Washington - Tacoma

L13.38

38

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 9, 2024

TCSS422: Operating Systems [Spring 2024]
School of Engineering and Technology, University of Washington - Tacoma

L13.39

39

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 9, 2024

TCSS422: Operating Systems [Spring 2024]
School of Engineering and Technology, University of Washington - Tacoma

L13.40

40

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 9, 2024
TCSS422: Operating Systems [Spring 2024]
School of Engineering and Technology, University of Washington - Tacoma
L13.41

41

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 9, 2024
TCSS422: Operating Systems [Spring 2024]
School of Engineering and Technology, University of Washington - Tacoma
L13.42

42

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 9, 2024
TCSS422: Operating Systems [Spring 2024]
School of Engineering and Technology, University of Washington - Tacoma
L13.43

43


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 9, 2024
TCSS422: Operating Systems [Spring 2024]
School of Engineering and Technology, University of Washington - Tacoma
L13.44


44

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 9, 2024	TCSS422: Operating Systems [Spring 2024] School of Engineering and Technology, University of Washington - Tacoma	L13.45
-------------	---	--------

45

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 9, 2024	TCSS422: Operating Systems [Spring 2024] School of Engineering and Technology, University of Washington - Tacoma	L13.46
-------------	---	--------

46

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 9, 2024	TCSS422: Operating Systems [Spring 2024] School of Engineering and Technology, University of Washington - Tacoma	L13.47
-------------	---	--------

47

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 9, 2024	TCSS422: Operating Systems [Spring 2024] School of Engineering and Technology, University of Washington - Tacoma	L13.48
-------------	---	--------

48

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

49

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 9, 2024 TCSS422: Operating Systems [Spring 2024]
School of Engineering and Technology, University of Washington - Tacoma L13.50

50

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 9, 2024
TCSS422: Operating Systems [Spring 2024]
 School of Engineering and Technology, University of Washington - Tacoma
L13.51

51

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

May 9, 2024
TCSS422: Operating Systems [Spring 2024]
 School of Engineering and Technology, University of Washington - Tacoma
L13.52

52

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 9, 2024	TCSS422: Operating Systems [Spring 2024] School of Engineering and Technology, University of Washington - Tacoma	L13.53
-------------	---	--------


53

OBJECTIVES – 5/9

- Questions from 5/9
- Pthread Tutorial-May 24 / Assignment 2 posted next week
- Quiz 3 – Synchronized Array (class activity-next week)
- Chapter 30: Condition Variables
 - Covering Conditions
- Chapter 32: Concurrency Problems
 - Non-deadlock concurrency bugs
 - Deadlock causes
 - Deadlock prevention
- **Chapter 13: Address Spaces**
- Chapter 14: The Memory API

May 9, 2024	TCSS422: Operating Systems [Spring 2024] School of Engineering and Technology, University of Washington - Tacoma	L13.54
-------------	---	--------

54



CHAPTER 13: ADDRESS SPACES

May 9, 2024

TCSS422: Operating Systems [Spring 2024]
School of Engineering and Technology, University of Washington - Tacoma

L13.55

55

OBJECTIVES – 5/9

- **Chapter 13: Introduction to memory virtualization**
 - The address space
 - Goals of OS memory virtualization
- **Chapter 14: Memory API**
 - Common memory errors

May 9, 2024

TCSS422: Operating Systems [Spring 2024]
School of Engineering and Technology, University of Washington - Tacoma

L13.56

56

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, 2024	TCSS422: Operating Systems [Spring 2024] School of Engineering and Technology, University of Washington - Tacoma	L13.57
-------------	---	--------

57


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, 2024	TCSS422: Operating Systems [Spring 2024] School of Engineering and Technology, University of Washington - Tacoma	L13.58
-------------	---	--------

58

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

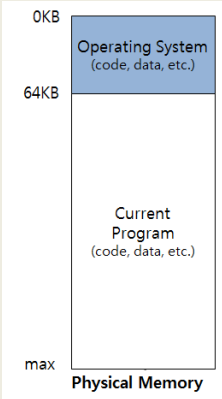
- **Protection**
 - From other processes
 - From programmer error (segmentation fault)

May 9, 2024	TCSS422: Operating Systems [Spring 2024] School of Engineering and Technology, University of Washington - Tacoma	L13.59
-------------	---	--------

59

EARLY MEMORY MANAGEMENT

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



The diagram illustrates physical memory management. It shows a vertical bar representing memory. The top portion is labeled 'Operating System (code, data, etc.)' and extends to the '0KB' mark. Below that is a larger section labeled 'Current Program (code, data, etc.)' extending to the '64KB' mark. The bottom of the bar is labeled 'max' and 'Physical Memory'. This indicates that only the OS and one program are loaded at a time, leaving the rest of the memory unused.

May 9, 2024	TCSS422: Operating Systems [Spring 2024] School of Engineering and Technology, University of Washington - Tacoma	L13.60
-------------	---	--------

60

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

The diagram shows a vertical stack of memory segments. From top to bottom: Operating System (code, data, etc.) from 0KB to 64KB; Free space from 64KB to 128KB; Process C (code, data, etc.) from 128KB to 192KB; Process B (code, data, etc.) from 192KB to 256KB; Free space from 256KB to 320KB; Process A (code, data, etc.) from 320KB to 384KB; Free space from 384KB to 448KB; and Free space from 448KB to 512KB. The total is labeled as Physical Memory.

May 9, 2024

TCSS422: Operating Systems [Spring 2024]
School of Engineering and Technology, University of Washington - Tacoma

L13.61

61

ADDRESS SPACE

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

The diagram shows a vertical stack of address space components. From top to bottom: Program Code from 0KB to 1KB; Heap from 1KB to 2KB; a large (free) region from 2KB to 15KB; and Stack from 15KB to 16KB. Arrows indicate the direction of growth: the heap grows downwards and the stack grows upwards.

May 9, 2024

TCSS422: Operating Systems [Spring 2024]
School of Engineering and Technology, University of Washington - Tacoma

L13.62

62

ADDRESS SPACE - 2

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

The diagram shows a vertical axis for address space from 0KB to 16KB. At the top (0KB) is 'Program Code'. Below it (1KB) is 'Heap'. Below that (2KB) is a large shaded area labeled '(free)'. At the bottom (15KB) is 'Stack'. An arrow points down from the Heap, and another arrow points up to the Stack, indicating they grow towards each other.

May 9, 2024	TCSS422: Operating Systems [Spring 2024] School of Engineering and Technology, University of Washington - Tacoma	L13.63
-------------	---	--------

63

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

The diagram shows a vertical axis for address space from 0KB to 16KB. At the top (0KB) is 'Program Code'. Below it (1KB) is 'Heap'. Below that (2KB) is a large shaded area labeled '(free)'. At the bottom (15KB) is 'Stack'. An arrow points down from the Heap, and another arrow points up to the Stack, indicating they grow towards each other.

May 9, 2024	TCSS422: Operating Systems [Spring 2024] School of Engineering and Technology, University of Washington - Tacoma	L13.64
-------------	---	--------

64

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, 2024	TCSS422: Operating Systems [Spring 2024] School of Engineering and Technology, University of Washington - Tacoma	L13.65
-------------	---	--------

65

VIRTUAL ADDRESSING - 2

- Output from 64-bit Linux:

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

The diagram illustrates the virtual address space layout. It is divided into several segments: Code (Text) at 0x400000, Data at 0x401000, Heap at 0xc2000, a (free) region at 0xd13000, stack at 0x7ff9ca28000, and Stack at 0x7ff9ca49000. Arrows indicate the direction of memory growth: the heap grows downwards from 0xc2000, and the stack grows upwards from 0x7ff9ca49000.

May 9, 2024	TCSS422: Operating Systems [Spring 2024] School of Engineering and Technology, University of Washington - Tacoma	L13.66
-------------	---	--------

66

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, 2024	TCSS422: Operating Systems [Spring 2024] School of Engineering and Technology, University of Washington - Tacoma	L13.67
-------------	---	--------

67

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 9, 2024	TCSS422: Operating Systems [Spring 2024] School of Engineering and Technology, University of Washington - Tacoma	L13.68
-------------	---	--------

68


OBJECTIVES – 5/9

- Questions from 5/9
- Pthread Tutorial-May 24 / Assignment 2 posted next week
- Quiz 3 – Synchronized Array (class activity-next week)
- Chapter 30: Condition Variables
 - Covering Conditions
- Chapter 32: Concurrency Problems
 - Non-deadlock concurrency bugs
 - Deadlock causes
 - Deadlock prevention
- Chapter 13: Address Spaces
- **Chapter 14: The Memory API**

May 9, 2024	TCSS422: Operating Systems [Spring 2024] School of Engineering and Technology, University of Washington - Tacoma	L13.69
-------------	---	--------

69

CHAPTER 14: THE MEMORY API



May 9, 2024	TCSS422: Operating Systems [Spring 2024] School of Engineering and Technology, University of Washington - Tacoma	L13.70
-------------	---	--------

70

OBJECTIVES – 5/9

- **Chapter 13: Introduction to memory virtualization**
 - The address space
 - Goals of OS memory virtualization
- **Chapter 14: Memory API**
 - Common memory errors

May 9, 2024	TCSS422: Operating Systems [Spring 2024] School of Engineering and Technology, University of Washington - Tacoma	L13.71
-------------	---	--------

71

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 9, 2024	TCSS422: Operating Systems [Spring 2024] School of Engineering and Technology, University of Washington - Tacoma	L13.72
-------------	---	--------

72

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 9, 2024	TCSS422: Operating Systems [Spring 2024] School of Engineering and Technology, University of Washington - Tacoma	L13.73
-------------	---	--------

73

FREE()

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

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

May 9, 2024	TCSS422: Operating Systems [Spring 2024] School of Engineering and Technology, University of Washington - Tacoma	L13.74
-------------	---	--------

74

```
#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?

75

75

```
#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?

76

76

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, 2024

TCSS422: Operating Systems [Spring 2024]
School of Engineering and Technology, University of Washington - Tacoma

L13.77

77

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, 2024

TCSS422: Operating Systems [Spring 2024]
School of Engineering and Technology, University of Washington - Tacoma

L13.78

78

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, 2024	TCSS422: Operating Systems [Spring 2024] School of Engineering and Technology, University of Washington - Tacoma	L13.79
-------------	---	--------

79

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, 2024	TCSS422: Operating Systems [Spring 2024] School of Engineering and Technology, University of Washington - Tacoma	L13.80
-------------	---	--------

80

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

Undefined Error

May 9, 2024	TCSS422: Operating Systems [Spring 2024] School of Engineering and Technology, University of Washington - Tacoma	L13.81
-------------	---	--------

81

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, 2024	TCSS422: Operating Systems [Spring 2024] School of Engineering and Technology, University of Washington - Tacoma	L13.82
-------------	---	--------

82



83