

15-213
 "The course that gives CMU its Zip!"

Synchronization
 November 19, 2008

Topics

- Synchronizing with semaphores
- Races and deadlocks
- Thread safety and reentrancy

lecture-24.ppt

badcnt.c: An Improperly Synchronized Threaded Program

```

/* shared */
volatile unsigned int cnt = 0;
#define NITERS 100000000

int main() {
    pthread_t tid1, tid2;
    Pthread_create(&tid1, NULL,
                  count, NULL);
    Pthread_create(&tid2, NULL,
                  count, NULL);

    Pthread_join(tid1, NULL);
    Pthread_join(tid2, NULL);

    if (cnt != (unsigned)NITERS*2)
        printf("BOOM! cnt=%d\n",
              cnt);
    else
        printf("OK cnt=%d\n",
              cnt);
}
    
```

```

/* thread routine */
void *count(void *arg) {
    int i;
    for (i=0; i<NITERS; i++)
        cnt++;
    return NULL;
}
    
```

```

linux> ./badcnt
BOOM! cnt=198841183

linux> ./badcnt
BOOM! cnt=198261801

linux> ./badcnt
BOOM! cnt=198269672
    
```

cnt should be
 equal to 200,000,000.
 What went wrong?!

- 2 -

15-213, F08

Assembly Code for Counter Loop

C code for counter loop

```

for (i=0; i<NITERS; i++)
    cnt++;
    
```

Corresponding asm code

```

Head (Hi) {
.L9:
    movl -4(%ebp),%eax
    cmpl $99999999,%eax
    jle .L12
    jmp .L10
}
Load cnt (Li)
Update cnt (Ui)
Store cnt (Si)
Tail (Ti) {
.L11:
    movl -4(%ebp),%eax
    leal 1(%eax),%edx # Update
    movl %edx,cnt # Store
}
.L10:
    
```

- 3 -

15-213, F08

Concurrent Execution

Key idea: In general, any sequentially consistent interleaving is possible, but some are incorrect!

- I_i denotes that thread i executes instruction I
- $\%eax_i$ is the contents of $\%eax$ in thread i 's context

i (thread)	$instr_i$	$\%eax_1$	$\%eax_2$	cnt
1	H_1	-	-	0
1	L_1	0	-	0
1	U_1	1	-	0
1	S_1	1	-	1
2	H_2	-	-	1
2	L_2	-	1	1
2	U_2	-	2	1
2	S_2	-	2	2
2	T_2	-	2	2
1	T_1	1	-	2

OK

- 4 -

15-213, F08

Concurrent Execution (cont)

Incorrect ordering: two threads increment the counter, but the result is 1 instead of 2.

i (thread)	instr _i	%eax ₁	%eax ₂	cnt
1	H ₁	-	-	0
1	L ₁	0	-	0
1	U ₁	1	-	0
2	H ₂	-	-	0
2	L ₂	-	0	0
1	S ₁	1	-	1
1	T ₁	1	-	1
2	U ₂	-	1	1
2	S ₂	-	1	1
2	T ₂	-	1	1

Oops!

- 5 -

15-213, F'08

Concurrent Execution (cont)

How about this ordering?

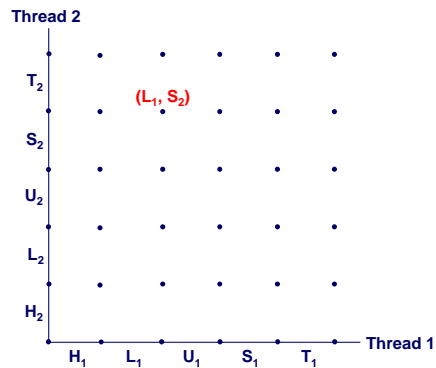
i (thread)	instr _i	%eax ₁	%eax ₂	cnt
1	H ₁			
1	L ₁			
2	H ₂			
2	L ₂			
2	U ₂			
2	S ₂			
1	U ₁			
1	S ₁			
1	T ₁			
2	T ₂			

We can clarify our understanding of concurrent execution with the help of the *progress graph*

- 6 -

15-213, F'08

Progress Graphs



A *progress graph* depicts the discrete *execution state space* of concurrent threads.

Each axis corresponds to the sequential order of instructions in a thread.


Each point corresponds to a possible *execution state* (Inst₁, Inst₂).

E.g., (L₁, S₂) denotes state where thread 1 has completed L₁ and thread 2 has completed S₂.

- 7 -

15-213, F'08

Trajectories in Progress Graphs



A *trajectory* is a sequence of legal state transitions that describes one possible concurrent execution of the threads.

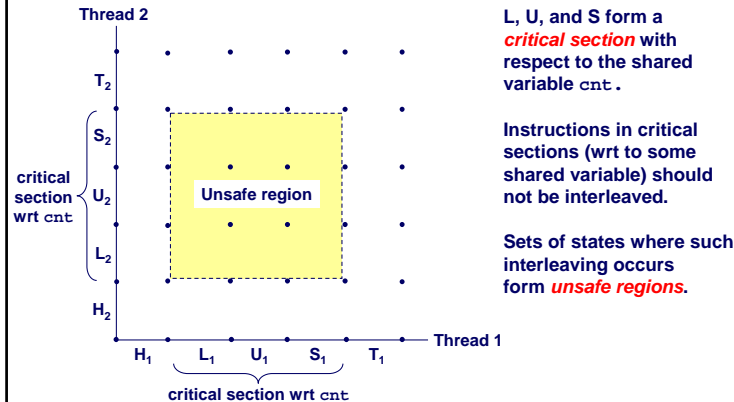
Example:

H1, L1, U1, H2, L2, S1, T1, U2, S2, T2

- 8 -

15-213, F'08

Critical Sections and Unsafe Regions



L, U, and S form a **critical section** with respect to the shared variable *cnt*.


Instructions in critical sections (wrt to some shared variable) should not be interleaved.

Sets of states where such interleaving occurs form **unsafe regions**.

- 9 -

15-213, F'08

Safe and Unsafe Trajectories



Def: A trajectory is **safe** iff it doesn't touch any part of an unsafe region.

Claim: A trajectory is correct (wrt *cnt*) iff it is safe.

- 10 -

15-213, F'08

Semaphores

Question: How can we guarantee a safe trajectory?

- We must **synchronize** the threads so that they never enter an unsafe state.

Classic solution: Dijkstra's P and V operations on **semaphores**.

- **semaphore:** non-negative integer synchronization variable
 - P(s): [while (s == 0) wait(); s--;]
 - » Dutch for "Proberen" (test)
 - V(s): [s++;]
 - » Dutch for "Verhogen" (increment)
- OS guarantees that operations between brackets [] are executed indivisibly
 - Only one P or V operation at a time can modify s
 - When while loop in P terminates, only that P can decrement s

Semaphore invariant: ($s \geq 0$)

- 11 -

15-213, F'08

Locking with Semaphores

Here is one way we could use P and V operations to **synchronize the threads that update *cnt***

- Semaphore used like this referred to as a "lock"

```

/* Semaphore s is initially 1 */


/* Thread routine */
void *count(void *arg)
{
    int i;

    for (i=0; i<NITERS; i++) {
        P(s);
        cnt++;
        V(s);
    }
    return NULL;
}
    
```

- 12 -

15-213, F'08

Safe Sharing With Locks



Provide mutually exclusive access to shared variable by surrounding critical section with P and V operations on semaphore s (initially set to 1).

Semaphore invariant creates a **forbidden region** that encloses unsafe region and is never touched by any trajectory.

- 13 -

15-213, F08

Wrappers on POSIX Semaphores

```

/* Initialize semaphore sem to value */
/* pshared=0 if thread, pshared=1 if process */
void Sem_init(sem_t *sem, int pshared, unsigned int value) {
    if (sem_init(sem, pshared, value) < 0)
        unix_error("Sem_init");
}

/* P operation on semaphore sem */
void P(sem_t *sem) {
    if (sem_wait(sem))
        unix_error("P");
}

/* V operation on semaphore sem */
void V(sem_t *sem) {
    if (sem_post(sem))
        unix_error("V");
}
    
```

- 14 -

15-213, F08

Sharing With POSIX Semaphores

```

/* properly sync'd counter program */
#include "csapp.h"
#define NITERS 10000000

volatile unsigned int cnt;
sem_t sem; /* semaphore */

int main() {
    pthread_t tid1, tid2;

    Sem_init(&sem, 0, 1); /* sem=1 */

    /* create 2 threads and wait */
    ...

    if (cnt != (unsigned)NITERS*2)
        printf("BOOM! cnt=%d\n", cnt);
    else
        printf("OK cnt=%d\n", cnt);
    exit(0);
}
    
```

```

/* thread routine */
void *count(void *arg)
{
    int i;

    for (i=0; i<NITERS; i++) {
        P(&sem);
        cnt++;
        V(&sem);
    }
    return NULL;
}
    
```

Warning:
It's really slow!

- 15 -

15-213, F08

One worry: races

A **race** occurs when the correctness of the program depends on one thread reaching point x before another thread reaches point y

```

/* a threaded program with a race */
int main() {
    pthread_t tid[N];
    int i;
    for (i = 0; i < N; i++)
        Pthread_create(&tid[i], NULL, thread, &i);
    for (i = 0; i < N; i++)
        Pthread_join(tid[i], NULL);
    exit(0);
}

/* thread routine */
void *thread(void *vargp) {
    int myid = *((int *)vargp);
    printf("Hello from thread %d\n", myid);
    return NULL;
}
    
```

- 16 -

15-213, F08

Race Elimination

Make sure don't have unintended sharing of state

```

/* a threaded program with a race */
int main() {
    pthread_t tid[N];
    int i;
    for (i = 0; i < N; i++) {
        int *valp = malloc(sizeof(int));
        *valp = i;
        Pthread_create(&tid[i], NULL, thread, valp);
    }
    for (i = 0; i < N; i++) {
        Pthread_join(tid[i], NULL);
    }
    exit(0);

    /* thread routine */
    void *thread(void *vargp) {
        int myid = *((int *)vargp);
        free(vargp);
        printf("Hello from thread %d\n", myid);
        return NULL;
    }
}

```

- 17 -

15-213, F'08

Another worry: Deadlock

- Processes wait for condition that will never be true

Typical Scenario

- Processes 1 and 2 needs two resources (A and B) to proceed
- Process 1 acquires A, waits for B
- Process 2 acquires B, waits for A
- Both will wait forever!

- 18 -

15-213, F'08

Deadlocking With POSIX Semaphores

```

int main()
{
    pthread_t tid[2];
    Sem_init(&mutex[0], 0, 1); /* mutex[0] = 1 */
    Sem_init(&mutex[1], 0, 1); /* mutex[1] = 1 */
    Pthread_create(&tid[0], NULL, count, (void*) 0);
    Pthread_create(&tid[1], NULL, count, (void*) 1);
    Pthread_join(tid[0], NULL);
    Pthread_join(tid[1], NULL);
    printf("cnt=%d\n", cnt);
    exit(0);
}

```

```

void *count(void *vargp)
{
    int i;
    int id = (int) vargp;
    for (i = 0; i < NITERS; i++) {
        P(&mutex[id]); P(&mutex[1-id]);
        cnt++;
        V(&mutex[id]); V(&mutex[1-id]);
    }
    return NULL;
}

```

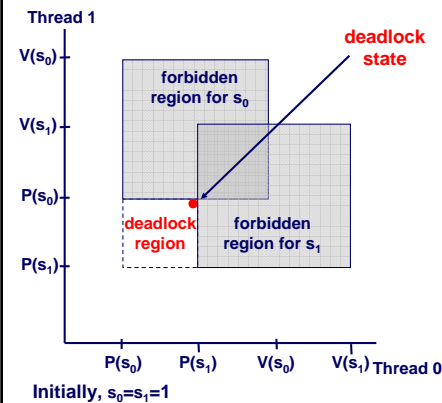
Tid[0]:
P(s₀);
P(s₁);
cnt++;
V(s₀);
V(s₁);

Tid[1]:
P(s₁);
P(s₀);
cnt++;
V(s₁);
V(s₀);

- 19 -

15-213, F'08

Deadlock



Locking introduces the potential for **deadlock**: waiting for a condition that will never be true.

Any trajectory that enters the **deadlock region** will eventually reach the **deadlock state**, waiting for either s_0 or s_1 to become nonzero.

Other trajectories luck out and skirt the deadlock region.

Unfortunate fact: deadlock is often non-deterministic.

- 20 -

15-213, F'08

Avoiding Deadlock

Acquire shared resources in same order

```
int main()
{
    pthread_t tid[2];
    Sem_init(&mutex[0], 0, 1); /* mutex[0] = 1 */
    Sem_init(&mutex[1], 0, 1); /* mutex[1] = 1 */
    Pthread_create(&tid[0], NULL, count, (void*) 0);
    Pthread_create(&tid[1], NULL, count, (void*) 1);
    Pthread_join(tid[0], NULL);
    Pthread_join(tid[1], NULL);
    printf("cnt=%d\n", cnt);
    exit(0);
}
```

```
void *count(void *vargp)
{
    int i;
    int id = (int) vargp;
    for (i = 0; i < NITERS; i++) {
        P(&mutex[0]); P(&mutex[1]);
        cnt++;
        V(&mutex[id]); V(&mutex[1-id]);
    }
    return NULL;
}
```

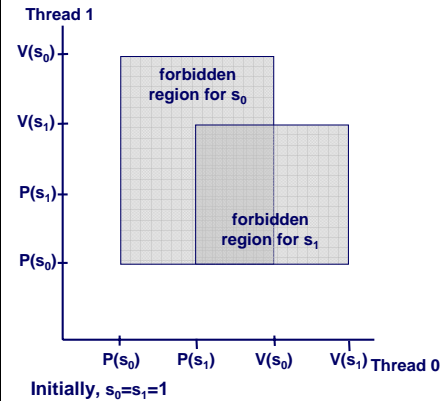
Tid[0]:
P(s₀);
P(s₁);
cnt++;
V(s₀);
V(s₁);

Tid[1]:
P(s₀);
P(s₁);
cnt++;
V(s₁);
V(s₀);

- 21 -

15-213, F'08

Removed Deadlock



No way for trajectory to get stuck

Processes acquire locks in same order

Order in which locks released immaterial

- 22 -

15-213, F'08

Crucial concept: Thread Safety

Functions called from a thread (without external synchronization) must be **thread-safe**

- i.e., it must be safe for multiple threads to be calling it concurrently

Some examples of thread-unsafe functions:

- Failing to protect shared variables
- Relying on persistent state across invocations
- Returning a pointer to a static variable
- Calling thread-unsafe functions

- 23 -

15-213, F'08

Thread-Unsafe Functions (1)

Failing to protect shared variables

- Fix: Use P and V semaphore operations
- Example: goodcnt.c
- Issue: Synchronization operations will slow down code
 - e.g., badcnt requires 0.5s, goodcnt requires 7.9s

- 24 -

15-213, F'08

Thread-Unsafe Functions (2)

Relying on persistent state across multiple function invocations

- Example: Random number generator that relies on static state

```
/* rand - return pseudo-random integer on 0..32767 */
static unsigned int next = 1;
int rand(void)
{
    next = next*1103515245 + 12345;
    return (unsigned int)(next/65536) % 32768;
}

/* srand - set seed for rand() */
void srand(unsigned int seed)
{
    next = seed;
}
```

- 25 -

15-213, F08

Making Thread-Safe RNG

Pass state as part of argument

- and, thereby, eliminate static state

```
/* rand - return pseudo-random integer on 0..32767 */
int rand_r(int *nextp)
{
    *nextp = *nextp*1103515245 + 12345;
    return (unsigned int)(*nextp/65536) % 32768;
}
```

- Consequence: programmer using rand must maintain seed

- 26 -

15-213, F08

Thread-Unsafe Functions (cont)

Returning a ptr to a static variable

Fixes:

- 1. Rewrite code so caller passes pointer to struct
 - » Issue: Requires changes in caller and callee
- 2. **Lock-and-copy**
 - » Issue: Requires only simple changes in caller (and none in callee)
 - » However, caller must free memory

```
struct hostent
*gethostbyname(char name)
{
    static struct hostent h;
    <contact DNS and fill in h>
    return &h;
}
```

```
hostp = Malloc(...);
gethostbyname_r(name, hostp);
```

```
struct hostent
*gethostbyname_ts(char *name)
{
    struct hostent *q = Malloc(...);
    struct hostent *p;
    P(&mutex); /* lock */
    p = gethostbyname(name);
    *q = *p; /* copy */
    V(&mutex);
    return q;
}
```

- 27 -

15-213, F08

Thread-Unsafe Functions

Calling thread-unsafe functions

- Calling one thread-unsafe function makes the entire function that calls it thread-unsafe
- Fix: Modify the function so it calls only thread-safe functions ©

- 28 -

15-213, F08

Thread-Safe Library Functions

All functions in the Standard C Library (at the back of your K&R text) are thread-safe

- Examples: malloc, free, printf, scanf


Most Unix system calls are thread-safe, with a few exceptions:

Thread-unsafe function	Class	Reentrant version
asctime	3	asctime_r
ctime	3	ctime_r
gethostbyaddr	3	gethostbyaddr_r
gethostbyname	3	gethostbyname_r
inet_ntoa	3	(none)
localtime	3	localtime_r
rand	2	rand_r

- 29 -

15-213, F08

Notifying With Semaphores



Common synchronization pattern:

- Producer waits for slot, inserts item in buffer, and notifies consumer
- Consumer waits for item, removes it from buffer, and notifies producer

Examples

- Multimedia processing:
 - Producer creates MPEG video frames, consumer renders them
- Event-driven graphical user interfaces
 - Producer detects mouse clicks, mouse movements, and keyboard hits and inserts corresponding events in buffer
 - Consumer retrieves events from buffer and paints the display

- 30 -

15-213, F08

Producer-Consumer on a Buffer That Holds One Item

```

/* buf1.c - producer-consumer
on 1-element buffer */
#include "csapp.h"

#define NITERS 5

void *producer(void *arg);
void *consumer(void *arg);

struct {
    int buf; /* shared var */
    sem_t full; /* sems */
    sem_t empty;
} shared;
  
```

```

int main() {
    pthread_t tid_producer;
    pthread_t tid_consumer;

    /* initialize the semaphores */
    Sem_init(&shared.empty, 0, 1);
    Sem_init(&shared.full, 0, 0);

    /* create threads and wait */
    Pthread_create(&tid_producer, NULL,
                  producer, NULL);
    Pthread_create(&tid_consumer, NULL,
                  consumer, NULL);
    Pthread_join(tid_producer, NULL);
    Pthread_join(tid_consumer, NULL);

    exit(0);
}
  
```

- 31 -

15-213, F08

Producer-Consumer (cont)

Initially: empty = 1, full = 0.

```

/* producer thread */
void *producer(void *arg) {
    int i, item;

    for (i=0; i<NITERS; i++) {
        /* produce item */
        item = i;
        printf("produced %d\n",
              item);

        /* write item to buf */
        P(&shared.empty);
        shared.buf = item;
        V(&shared.full);
    }
    return NULL;
}
  
```

```

/* consumer thread */
void *consumer(void *arg) {
    int i, item;

    for (i=0; i<NITERS; i++) {
        /* read item from buf */
        P(&shared.full);
        item = shared.buf;
        V(&shared.empty);

        /* consume item */
        printf("consumed %d\n",
              item);
    }
    return NULL;
}
  
```

- 32 -

15-213, F08

Counting with Semaphores

Remember, it's a non-negative integer

- So, values greater than 1 are legal

Lets repeat thing_5() 5 times for every 3 of thing_3()

```
/* thing_5 and thing_3 */
#include "csapp.h"

sem_t five;
sem_t three;

void *five_times(void *arg);
void *three_times(void *arg);
```

```
int main() {
    pthread_t tid_five, tid_three;

    /* initialize the semaphores */
    Sem_init(&five, 0, 5);
    Sem_init(&three, 0, 3);

    /* create threads and wait */
    Pthread_create(&tid_five, NULL,
                  five_times, NULL);
    Pthread_create(&tid_three, NULL,
                  three_times, NULL);

    .
    .
    .
}
```

- 33 -

15-213, F08

Counting with semaphores (cont)

Initially: five = 5, three = 3

```
/* thing_5() thread */
void *five_times(void *arg) {
    int i;

    while (1) {
        for (i=0; i<5; i++) {
            /* wait & thing_5() */
            P(&five);
            thing_5();
        }
        V(&three);
        V(&three);
        V(&three);
    }
    return NULL;
}
```

```
/* thing_3() thread */
void *three_times(void *arg) {
    int i;

    while (1) {
        for (i=0; i<3; i++) {
            /* wait & thing_3() */
            P(&three);
            thing_3();
        }
        V(&five);
        V(&five);
        V(&five);
        V(&five);
        V(&five);
    }
    return NULL;
}
```

- 34 -

15-213, F08

Threads Summary

Threads provide another mechanism for writing concurrent programs

Threads are growing in popularity

- Somewhat cheaper than processes
- Easy to share data between threads

However, the ease of sharing has a cost:

- Easy to introduce subtle synchronization errors
- Tread carefully with threads!

For more info:

- D. Butenhof, "Programming with Posix Threads", Addison-Wesley, 1997.

- 35 -

Beware of Optimizing Compilers!

Code From Book

```
#define NITERS 100000000

/* shared counter variable */
unsigned int cnt = 0;

/* thread routine */
void *count(void *arg)
{
    int i;
    for (i = 0; i < NITERS; i++)
        cnt++;
    return NULL;
}
```

Generated Code

```
movl    cnt, %ecx
movl    $99999999, %eax
.L6:
    leal 1(%ecx), %edx
    decl %eax
    movl %edx, %ecx
    jns .L6
    movl %edx, cnt
```

- Global variable cnt shared between threads
- Multiple threads could be trying to update within their iterations

- Compiler moved access to cnt out of loop
- Only shared accesses to cnt occur before loop (read) or after (write)
- What are possible program outcomes?

- 36 -

15-213, F08