

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

Synchronization December 4, 2007

Topics

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

lectures-28.ppt

badcnt.c: An Improperly Synchronized Threaded Program

```

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

int main() {
    pthread_t t1, t2;
    pthread_mutex_t mutex = PTHREAD_MUTEX_INITIALIZER;
    pthread_mutex_t mutex2 = PTHREAD_MUTEX_INITIALIZER;

    pthread_join(t1, NULL);
    pthread_join(t2, NULL);

    if (cnt != (unsigned)NTHREADS)
        printf("BAD! cnt=%d\n", cnt);
    else
        printf("OK! cnt=%d\n", cnt);
}
    
```

```

/* shared routine */
void Pthread(void *arg) {
    int i;
    for (i=0; i<NTHREADS; i++)
        cnt++;
}
    
```

```

$./badcnt
BAD! cnt=988813383
BAD! cnt=988813801
BAD! cnt=988888712
    
```

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

-3-

10.01.F07

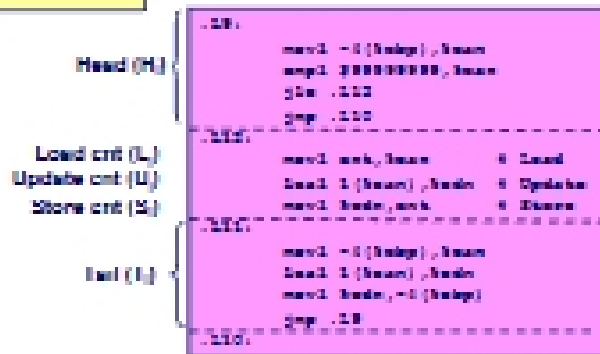
Assembly Code for Counter Loop

C code for counter loop

```

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

Corresponding asm code



-3-

10.01.F07

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)	I_{i1}	$\%eax_1$	$\%eax_2$	cnt
1	I_{11}	-	-	0
1	L_1	0	-	0
1	U_1	1	-	0
1	S_1	1	-	1
2	I_{21}	-	-	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-

10.01.F07

Concurrent Execution (cont)

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

i (thread)	I_{i1}	$\%eax_1$	$\%eax_2$	cnt
1	I_{11}	-	-	0
1	L_1	0	-	0
1	U_1	1	-	0
2	I_{21}	-	-	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-

10.01.F07

Concurrent Execution (cont)

How about this ordering?

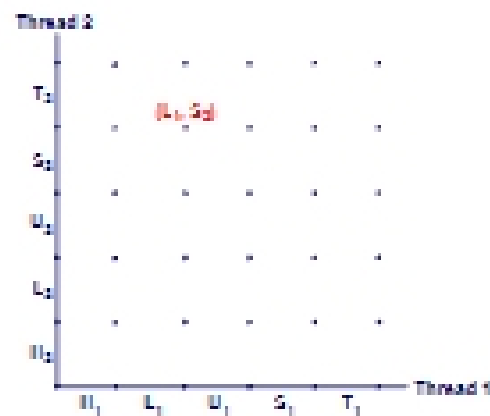
i (thread)	I_{i1}	$\%eax_1$	$\%eax_2$	cnt
1	I_{11}			
1	L_1			
2	I_{21}			
2	L_2			
2	U_2			
2	S_2			
1	U_1			
1	S_1			
1	T_1			
2	T_2			

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

-6-

10.01.F07

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.

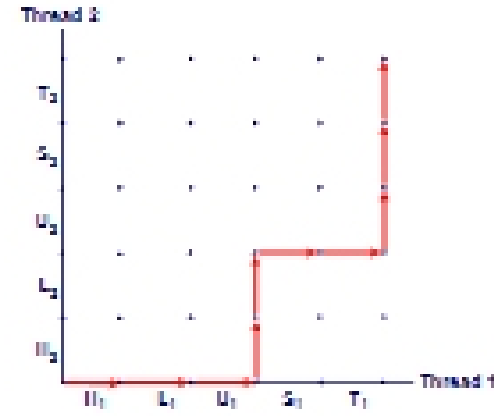
Each point corresponds to a possible **execution state** $(\text{inst}_1, \text{inst}_2)$.

E.g., (L_1, S_2) denotes state where thread 1 has completed L_1 and thread 2 has completed S_2 .

- 7 -

10.310.F07

Trajectories in Progress Graphs



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

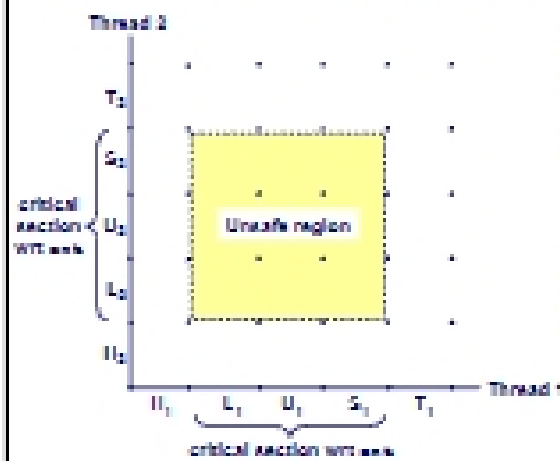
Example:

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

- 8 -

10.310.F07

Critical Sections and Unsafe Regions



L_i , U_i and S_i 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 -

10.310.F07

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 -

10.310.F07

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.

- **semaphores:** 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 -

10.310.F07

Safe Sharing with Semaphores

Here is how we would use P and V operations to synchronize the threads that update cnt .

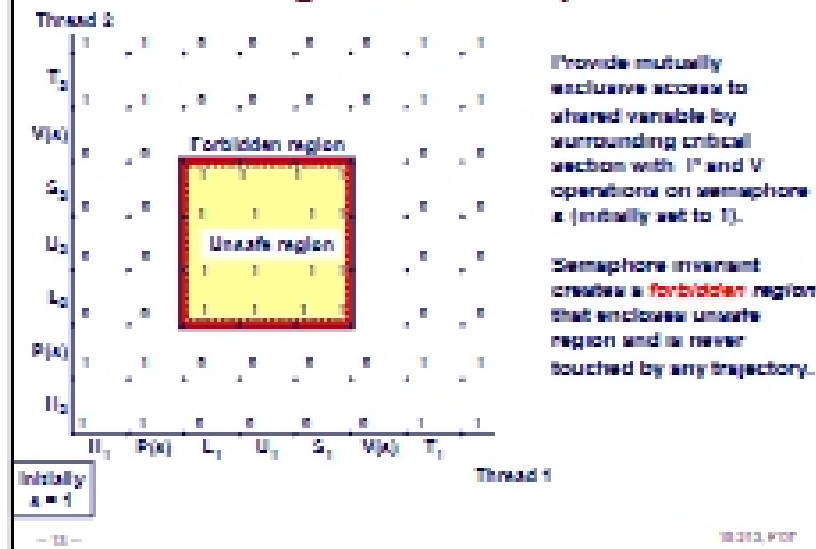
```
/* Semaphore s is initially 1 */
/* Thread routine */
void Thread(void *arg)
{
    int i;

    for (i=0; i<100000; i++) {
        P(s);
        cnt++;
        V(s);
    }
    return NULL;
}
```

- 12 -

10.310.F07

Safe Sharing With Semaphores



Wrappers on POSIX Semaphores

```

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

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

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

- 14 - 1023.PDF

Sharing With POSIX Semaphores

```

/* properly synchronized program */
#include "csapp.h"
#define N 1000000

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

int main() {
    pshared_t pshd;
    Sem_Init(&sem, 0, 1); /* sem=1 */

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

    if (x != (unsigned)N)
        printf("Error: x=%d\n", x);
    else
        printf("OK: x=%d\n", x);
    wait(0);
}

/* shared routine */
void *work(void *vpp) {
    int i;

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

- 15 - 1023.PDF

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() {
    pshared_t pshd(0);
    int i;

    for (i=0; i<N; i++)
        Pshared_mutex(&id[i], NULL, shared, 0);
    for (i=0; i<N; i++)
        Pshared_join(&id[i], NULL);
    wait(0);
}

/* shared routine */
void *work(void *vpp) {
    int i;
    printf("Hello from shared %d\n", i);
    return NULL;
}
    
```

- 16 - 1023.PDF

Deadlock

- Processes wait for condition that will never be true

Typical Scenario

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

- 17 - 1023.PDF

Deadlocking With POSIX Semaphores

```

int main() {
    pshared_t pshd(0);
    Sem_Init(&sem0, 0, 1); /* mutex[0] = 1 */
    Sem_Init(&sem1, 0, 1); /* mutex[1] = 1 */
    Pshared_mutex(&id[0], NULL, shared, (read?) 0);
    Pshared_mutex(&id[1], NULL, shared, (read?) 0);
    Pshared_join(&id[0], NULL);
    Pshared_join(&id[1], NULL);
    printf("x=%d\n", x);
    wait(0);
}
    
```

```

void *work(void *vpp) {
    int i;
    int id = (int) vpp;
    for (i=0; i<N; i++) {
        P(&sem[id]); P(&sem[1-id]);
        x++;
        V(&sem[id]); V(&sem[1-id]);
    }
    return NULL;
}
    
```

Tid(0):
P(x);
P(y);
crit;
V(x);

Tid(1):
P(y);
P(x);
crit;
V(y);

- 18 - 1023.PDF