

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

Concurrency II: Synchronization

April 12, 2001

Topics

- Progress graphs
- Semaphores
- Mutex and condition variables
- Barrier synchronization
- Timeout waiting

class23.ppt

Aversion of badcnt.c with a simple counter loop

```

int ctr = 0; /* shared */

/* main routine creates */
/* two count threads */

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

    for (i=0; i<NITERS; i++)
        ctr++;
    return NULL;
}
            
```

note: counters should be equal to 200,000,000

```

linux> badcnt
BOOM! ctr=198841183

linux> badcnt
BOOM! ctr=198261801

linux> badcnt
BOOM! ctr=198269672
            
```

What went wrong?

class23.ppt - 2 - CS213S'01

Assembly code for counter loop

C code for counter loop

```

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

**Corresponding asm code
(gcc -O0 -fforce-mem)**

```

.L9:
    movl -4(%ebp),%eax
    cmpl $99999999,%eax
    jle .L12
    jmp .L10

.L12:
    movl ctr,%eax # Load
    leal 1(%eax),%edx # Update
    movl %edx,ctr # Store

.L11:
    movl -4(%ebp),%eax
    leal 1(%eax),%edx
    movl %edx,-4(%ebp)
    jmp .L9

.L10:
            
```

Head (H_i) { .L9: ... jmp .L10 }

Load ctr (L_i) { .L12: movl ctr,%eax # Load }

Update ctr (U_i) { .L12: leal 1(%eax),%edx # Update }

Store ctr (S_i) { .L12: movl %edx,ctr # Store }

Tail (T_i) { .L11: ... jmp .L9 }

class23.ppt - 3 - CS213S'01

Concurrent execution

Key thread 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 | ctr |
|-----------|--------------------|-------------------|-------------------|-----|
| 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

class23.ppt - 4 - CS213S'01

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 | ctr |
|-----------|--------------------|-------------------|-------------------|-----|
| 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!

class23.ppt

- 5 -

CS213S'01

Concurrent execution (cont)

How about this ordering?

| i(thread) | instr _i | %eax ₁ | %eax ₂ | ctr |
|-----------|--------------------|-------------------|-------------------|-----|
| 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*

class23.ppt

- 6 -

CS213S'01

Progress graphs



A *progress graph* depicts the discrete *execution statespace* 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.


class23.ppt

- 7 -


CS213S'01

Legal state transitions

Interleaved concurrent execution (one processor):



Parallel concurrent execution (multiple processors)



Keypoint: Always reason about concurrent threads as if each thread had its own CPU.

class23.ppt

- 8 -

CS213S'01

Trajectories

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

Example:
 $H_1, L_2, U_1, H_2, L_2, S_1, T_1, U_2, S_2, T_2$

class23.ppt - 9 - CS213S'01

Critical sections and unsafe regions

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

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

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

class23.ppt - 10 - CS213S'01

Safe trajectories

Def: A *safe trajectory* is a sequence of legal transitions that does not touch any states in an unsafe region.

Claim: Any safe trajectory results in a correct value for the shared variable ctr .

class23.ppt - 11 - CS213S'01

Unsafe trajectories

Touching a state of type x is always incorrect.

Touching a state of type y may or may not be OK:

- correct because store completes before load.
- incorrect because order of load and store are indeterminate.

Moral: be conservative and disallow all unsafe trajectories.

class23.ppt - 12 - CS213S'01

Semaphore operations

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.

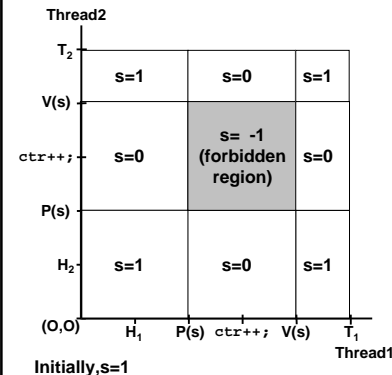
- **semaphore:** non-negative integers synchronization variable.
- **P(s):** [while (s == 0) wait(); s--;]
– Dutch for "Proberen" (test)
- **V(s):** [s++;]
– Dutch for "Verhogen" (increment)
- **OS guarantees** that operations between brackets [] are executed indivisibly.
– Only one P or V operation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

class23.ppt

- 13 -

CS213S'01

Sharing with semaphores



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.

Semaphore used in this way is often called a *mutex* (mutual exclusion).

class23.ppt

- 14 -

CS213S'01

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");
}
    
```

class23.ppt

- 15 -

CS213S'01

Sharing with Posix semaphores

```

/* goodcnt.c - properly synch'd */
/* version of badcnt.c */
#include <ics.h>
#define NITERS 10000000

void *count(void *arg);

struct {
    int ctr; /* shared ctr */
    sem_t mutex; /* semaphore */
} shared;

int main() {
    pthread_t tid1, tid2;

    /* init mutex semaphore to 1 */
    sem_init(&shared.mutex, 0, 1);

    /* create 2 ctr threads and wait */
    ...
}
    
```

class23.ppt

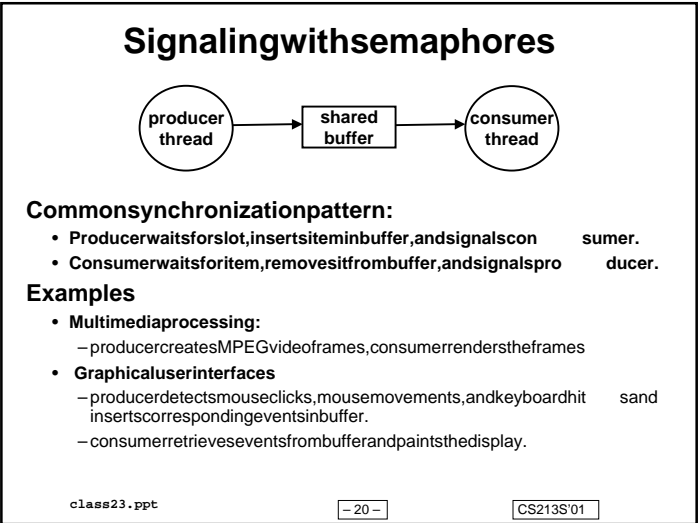
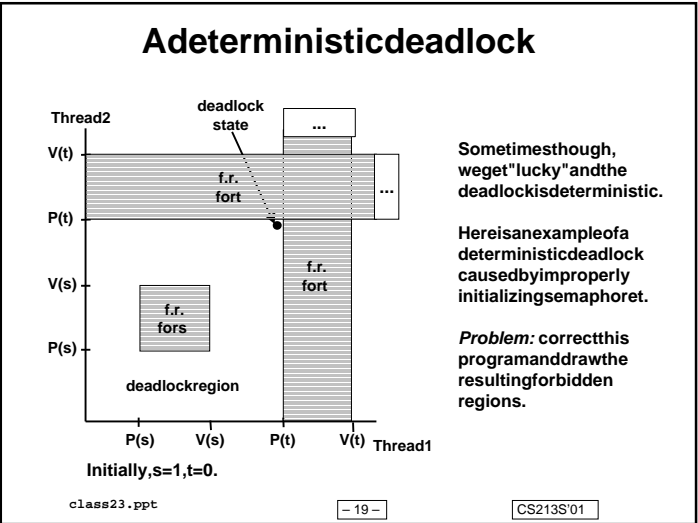
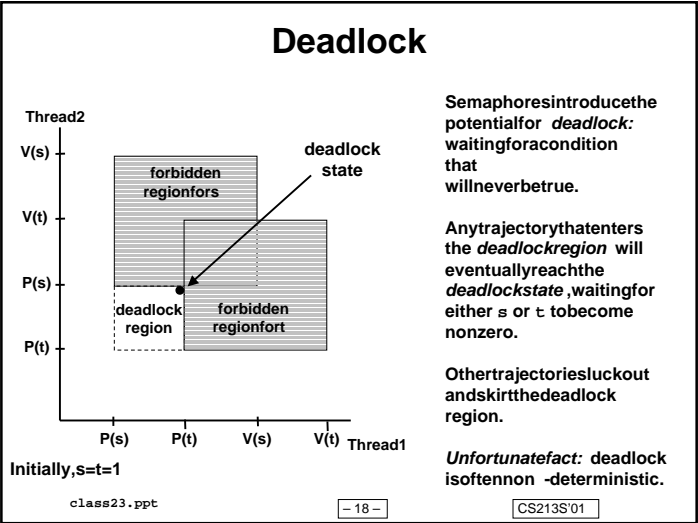

- 16 -

CS213S'01

```

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

    for (i=0; i<NITERS; i++) {
        P(&shared.mutex);
        shared.ctr++;
        V(&shared.mutex);
    }
    return NULL;
}
    
```



Producer-consumer(1 -buffer)

```

/* buf1.c - producer-consumer
on 1-element buffer */
#include <ics.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);
}

```

class23.ppt

- 21 -

CS213S'01

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;
}


```

class23.ppt

- 22 -

CS213S'01

Producer-consumer progress graph



class23.ppt

- 23 -

CS213S'01

Limitations of semaphores

Semaphores are sound and fundamental, but they have limitations.

- Difficult to broadcast a signal to a group of threads.
 - e.g., *barriersynchronization*: no thread returns from the barrier function until every other thread has called the barrier function.
- Impossible to do timeout waiting.
 - e.g., wait for at most 1 second for a condition to become true.

For these we must use Pthreads mutex and condition variables.

class23.ppt

- 24 -

CS213S'01

Basic operations on mutex variables

```
int pthread_mutex_init(pthread_mutex_t *mutex,
                      pthread_mutexattr_t *attr)
```

Initializes a mutex variable (`mutex`) with some attributes (`attr`).

- attributes are usually NULL.
- like initializing a mutex semaphore to 1.

```
int pthread_mutex_lock(pthread_mutex_t *mutex)
```

Indivisibly waits for `mutex` to be unlocked and then locks it.

- like P(`mutex`)

```
int pthread_mutex_unlock(pthread_mutex_t *mutex)
```

Unlocks `mutex`.

- like V(`mutex`)

class23.ppt

- 25 -

CS213S'01

Basic operations on condition variables

```
int pthread_cond_init(pthread_cond_t *cond,
                    pthread_condattr_t *attr)
```

Initializes a condition variable (`cond`) with some attributes (`attr`).

- attributes are usually NULL.

```
int pthread_cond_signal(pthread_cond_t *cond)
```

Awakens one thread waiting on condition `cond`.

- if no threads waiting on condition, then it does nothing.
- keypoint: signals are not queued!

```
int pthread_cond_wait(pthread_cond_t *cond, pthread_mutex_t *mutex)
```

Indivisibly unlocks `mutex` and waits for signal on condition `cond`

- When awakened, indivisibly locks `mutex`.

class23.ppt

- 26 -

CS213S'01

Advanced operations on condition variables

```
int pthread_cond_broadcast(pthread_cond_t *cond)
```

Awakens *all* threads waiting on condition `cond`.

- if no threads waiting on condition, then it does nothing.

```
int pthread_cond_timedwait(pthread_cond_t *cond,
                          pthread_mutex_t *mutex,
                          struct timespec *abstime)
```

Waits for condition `cond` until absolute wall clock time is `abstime`

- Unlocks `mutex` on entry, locks `mutex` on awakening.
- Use of absolute timer rather than relative time is strange.

class23.ppt

- 27 -

CS213S'01

Signaling and waiting on conditions

Basic pattern for signaling

```
Pthread_mutex_lock(&mutex);
Pthread_cond_signal(&cond);
Pthread_mutex_unlock(&mutex);
```

A mutex is always associated with a condition variable.

Guarantees that the condition cannot be signaled (and thus ignored) in the interval when the waiter locks the mutex and waits on the condition.

Basic pattern for waiting

```
Pthread_mutex_lock(&mutex);
Pthread_cond_wait(&cond, &mutex);
Pthread_mutex_unlock(&mutex);
```

class23.ppt

- 28 -

CS213S'01

```

#include <ics.h>

static pthread_mutex_t mutex;
static pthread_cond_t cond;
static int nthreads;
static int barriercnt = 0;

void barrier_init(int n) {
    nthreads = n;
    Pthread_mutex_init(&mutex, NULL);
    Pthread_cond_init(&cond, NULL);
}

void barrier() {
    Pthread_mutex_lock(&mutex);
    if (++barriercnt == nthreads) {
        barriercnt = 0;
        Pthread_cond_broadcast(&cond);
    }
    else
        Pthread_cond_wait(&cond, &mutex);
    Pthread_mutex_unlock(&mutex);
}

```

Barrier synchronization

Call to barrier will not return until every other thread has also called barrier.

Needed for tightly - coupled parallel applications that proceed in phases. E.g., physical simulations.

class23.ppt

- 29 -

CS213S'01

timebomb.c: timeout waiting example

A program that explodes unless the user hits a key within 5 seconds.

```

#include <ics.h>
#define TIMEOUT 5

/* function prototypes */
void *thread(void *vargp);
struct timespec *maketimeout(int secs);

/* condition variable and its associated mutex */
pthread_cond_t cond;
pthread_mutex_t mutex;

/* thread id */
pthread_t tid;

```

class23.ppt

- 30 -

CS213S'01

timebomb.c(cont)

A routine for building a timeout structure for pthread_cond_timewait.

```

/*
 * maketimeout - builds a timeout object that times out
 *               in secs seconds
 */
struct timespec *maketimeout(int secs) {
    struct timeval now;
    struct timespec *tp =
        (struct timespec *)malloc(sizeof(struct timespec));

    gettimeofday(&now, NULL);
    tp->tv_sec = now.tv_sec + secs;
    tp->tv_nsec = now.tv_usec * 1000;
    return tp;
}

```

class23.ppt

- 31 -

CS213S'01

Main routine for timebomb.c

```

int main() {
    int i, rc;

    /* initialize the mutex and condition variable */
    Pthread_cond_init(&cond, NULL);
    Pthread_mutex_init(&mutex, NULL);

    /* start getchar thread and wait for it to timeout */
    Pthread_mutex_lock(&mutex);
    Pthread_create(&tid, NULL, thread, NULL);
    for (i=0; i<TIMEOUT; i++) {
        printf("BEEP\n");
        rc = pthread_cond_timedwait(&cond, &mutex, maketimeout(1));
        if (rc != ETIMEDOUT) {
            printf("WHEW!\n");
            exit(0);
        }
    }
    printf("BOOM!\n");
    exit(0);
}

```

class23.ppt

- 32 -

CS213S'01

Thread routine for timebomb.c

```
/*  
 * thread - executes getchar in a separate thread  
 */  
void *thread(void *vargp) {  
  
    (void) getchar();  
  
    pthread_mutex_lock(&mutex);  
    pthread_cond_signal(&cond);  
    pthread_mutex_unlock(&mutex);  
    return NULL;  
}
```

class23.ppt

- 33 -

CS213S'01

Thread 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.

For more info:

- manpages(`man -k pthreads`)
- D. Butenhof, "Programming with Posix Threads", Addison-Wesley, 1997.

class23.ppt

- 34 -

CS213S'01