

Synchronization: Basics

15-213 / 18-213: Introduction to Computer Systems
24th Lecture, April. 17, 2012

Instructors:

Todd C. Mowry & Anthony Rowe

1

Today

- Threads review
- Sharing
- Mutual exclusion
- Semaphores

2

Process: Traditional View


- Process = process context + code, data, and stack

Process context

Program context:
Data registers
Condition codes
Stack pointer (SP)
Program counter (PC)

Kernel context:
VM structures
Descriptor table
brk pointer

Code, data, and stack



3

Process: Alternative View


- Process = thread + code, data, and kernel context

Thread

Program context:
Data registers
Condition codes
Stack pointer (SP)
Program counter (PC)



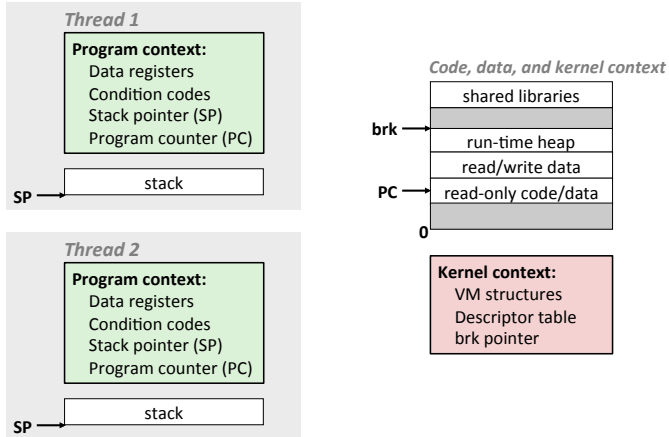
Code, data, and kernel context



Kernel context:
VM structures
Descriptor table
brk pointer

4

Process with Two Threads



5


Threads vs. Processes

- **Threads and processes: similarities**
 - Each has its own logical control flow
 - Each can run concurrently with others
 - Each is context switched (scheduled) by the kernel
- **Threads and processes: differences**
 - Threads share code and data, processes (typically) do not
 - Threads are less expensive than processes
 - Process control (creating and reaping) is more expensive as thread control
 - Context switches for processes more expensive than for threads

6

Threads vs. Processes (cont.)

- **Processes form a tree hierarchy**
- **Threads form a pool of peers**
 - Each thread can kill any other
 - Each thread can wait for any other thread to terminate
 - Main thread: first thread to run in a process



7

Posix Threads (Pthreads) Interface

- **Pthreads:** Standard interface for ~60 functions that manipulate threads from C programs
 - Threads run thread routines:
 - `void *threadroutine(void *vargp)`
 - Creating and reaping threads
 - `pthread_create(pthread_t *tid, ..., func *f, void *arg)`
 - `pthread_join(pthread_t tid, void **thread_return)`
 - Determining your thread ID
 - `pthread_self()`
 - Terminating threads
 - `pthread_cancel(pthread_t tid)`
 - `pthread_exit(void *thread_return)`
 - `return` (in primary thread routine terminates the thread)
 - `exit` (terminates all threads)

8

The Pthreads "Hello, world" Program

```

/*
 * hello.c - Pthreads "hello, world" program
 */
#include "csapp.h"

void *thread(void *vargp);

int main() {
    pthread_t tid;

    Pthread_create(&tid, NULL, thread, NULL);
    Pthread_join(tid, NULL);
    exit(0);
}

/* thread routine */
void *thread(void *vargp) {
    printf("Hello, world!\n");
    return NULL;
}

```

Thread attributes
(usually NULL)

Thread arguments
(void *p)

assigns return value
(void **p)

9

Pros and Cons of Thread-Based Designs

- + Easy to share data structures between threads
 - e.g., logging information, file cache
- + Threads are more efficient than processes
- – Unintentional sharing can introduce subtle and hard-to-reproduce errors!

10

Today

- Threads review
- Sharing
- Mutual exclusion
- Semaphores

11

Shared Variables in Threaded C Programs

- Question: Which variables in a threaded C program are shared?
 - The answer is not as simple as "global variables are shared" and "stack variables are private"
- Requires answers to the following questions:
 - What is the memory model for threads?
 - How are instances of variables mapped to memory?
 - How many threads might reference each of these instances?
- Def: A variable x is *shared* if and only if multiple threads reference some instance of x .

12

Threads Memory Model

■ Conceptual model:

- Multiple threads run within the context of a single process
- Each thread has its own separate thread context
 - Thread ID, stack, stack pointer, PC, condition codes, and GP registers
- All threads share the remaining process context
 - Code, data, heap, and shared library segments of the process virtual address space
 - Open files and installed handlers

■ Operationally, this model is not strictly enforced:

- Register values are truly separate and protected, but...
- Any thread can read and write the stack of any other thread

The mismatch between the conceptual and operation model is a source of confusion and errors

13

Example Program to Illustrate Sharing

```
char **ptr; /* global */

int main()
{
    int i;
    pthread_t tid;
    char *msgs[2] = {
        "Hello from foo",
        "Hello from bar"
    };
    ptr = msgs;

    for (i = 0; i < 2; i++)
        Pthread_create(&tid,
                       NULL,
                       thread,
                       (void *)i);
    Pthread_exit(NULL);
}
```

```
/* thread routine */
void *thread(void *vargp)
{
    int myid = (int) vargp;
    static int cnt = 0;

    printf("[%d]: %s (svar=%d)\n",
           myid, ptr[myid], ++cnt);
}
```

Peer threads reference main thread's stack indirectly through global ptr variable

14

Mapping Variable Instances to Memory

■ Global variables

- Def: Variable declared outside of a function
- Virtual memory contains exactly one instance of any global variable

■ Local variables

- Def: Variable declared inside function without `static` attribute
- Each thread stack contains one instance of each local variable

■ Local static variables

- Def: Variable declared inside function with the `static` attribute
- Virtual memory contains exactly one instance of any local static variable.

15

Mapping Variable Instances to Memory

Global var: 1 instance (ptr [data])

Local vars: 1 instance (i, m, msgs, m)

```
char **ptr; /* global */

int main()
{
    int i;
    pthread_t tid;
    char *msgs[2] = {
        "Hello from foo",
        "Hello from bar"
    };
    ptr = msgs;

    for (i = 0; i < 2; i++)
        Pthread_create(&tid,
                       NULL,
                       thread,
                       (void *)i);
    Pthread_exit(NULL);
}
```

Local var: 2 instances (myid.p0 [peer thread 0's stack], myid.p1 [peer thread 1's stack])

```
/* thread routine */
void *thread(void *vargp)
{
    int myid = (int)vargp;
    static int cnt = 0;

    printf("[%d]: %s (svar=%d)\n",
           myid, ptr[myid], ++cnt);
}
```

Local static var: 1 instance (cnt [data])

16

Shared Variable Analysis

■ Which variables are shared?

Variable instance	Referenced by main thread?	Referenced by peer thread 0?	Referenced by peer thread 1?
ptr	yes	yes	yes
cnt	no	yes	yes
i.m	yes	no	no
msgs.m	yes	yes	yes
myid.p0	no	yes	no
myid.p1	no	no	yes

Answer: A variable x is shared iff multiple threads reference at least one instance of x . Thus:

- ptr, cnt, and msgs are shared
- i and myid are *not* shared

17

Today

- Threads review
- Sharing
- Mutual exclusion
- Semaphores

18

badcnt.c: Improper Synchronization

```
volatile int cnt = 0; /* global */

int main(int argc, char **argv)
{
    int niters = atoi(argv[1]);
    pthread_t tid1, tid2;

    pthread_create(&tid1, NULL,
                  thread, &niters);
    pthread_create(&tid2, NULL,
                  thread, &niters);
    pthread_join(tid1, NULL);
    pthread_join(tid2, NULL);

    /* Check result */
    if (cnt != (2 * niters))
        printf("BOOM! cnt=%d\n", cnt);
    else
        printf("OK cnt=%d\n", cnt);
    exit(0);
}
```

```
/* Thread routine */
void *thread(void *vargp)
{
    int i, niters = *((int *)vargp);

    for (i = 0; i < niters; i++)
        cnt++;

    return NULL;
}
```

```
linux> ./badcnt 10000
OK cnt=20000
linux> ./badcnt 10000
BOOM! cnt=13051
linux>
```

cnt should equal 20,000.

What went wrong?

19

Assembly Code for Counter Loop

C code for counter loop in thread i

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

Corresponding assembly code

```
movl (%rdi), %ecx
movl $0, %edx
cmpl %ecx, %edx
jge .L13
.L11:
movl cnt(%rip), %eax
incl %eax
movl %eax, cnt(%rip)
incl %edx
cmpl %ecx, %edx
jl .L11
.L13:
```

Head (H_i)

Load cnt (L_i)
Update cnt (U_i)
Store cnt (S_i)

Tail (T_i)

20

Concurrent Execution

- **Key idea:** In general, any sequentially consistent interleaving is possible, but some give an unexpected result!
 - I_i denotes that thread i executes instruction I
 - $\%eax_i$ is the content of $\%eax$ in thread i 's context

i (thread)	instr _{<i>i</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

Thread 1 critical section
 Thread 2 critical section

OK

21

Concurrent Execution (cont)

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

i (thread)	instr _{<i>i</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!

22

Concurrent Execution (cont)

- How about this ordering?


i (thread)	instr _{<i>i</i>}	$\%eax_1$	$\%eax_2$	cnt
1	H ₁	-	-	0
1	L ₁	0	-	0
2	H ₂	-	-	0
2	L ₂	-	0	0
2	U ₂	-	1	0
2	S ₂	-	1	1
1	U ₁	1	-	1
1	S ₁	1	-	1
1	T ₁	-	-	1
2	T ₂	-	-	1

Oops!

- We can analyze the behavior using a **progress graph**

23

Progress Graphs



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

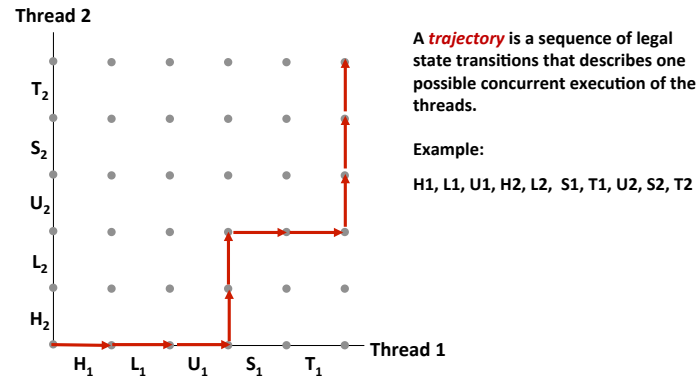
Each axis corresponds to the sequential order of instructions in a thread.

Each point corresponds to a possible **execution state** $(Inst_1, Inst_2)$.

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

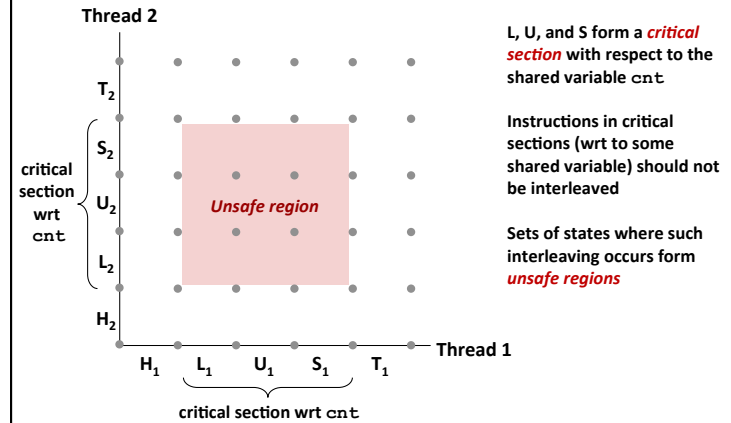
24

Trajectories in Progress Graphs



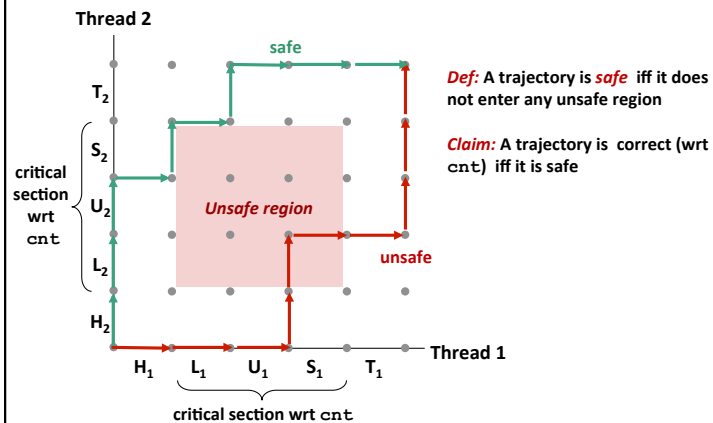
25

Critical Sections and Unsafe Regions



26

Critical Sections and Unsafe Regions



27

Enforcing Mutual Exclusion

- **Question:** How can we guarantee a safe trajectory?
- **Answer:** We must **synchronize** the execution of the threads so that they never have an unsafe trajectory.
 - i.e., need to guarantee **mutually exclusive access** to critical regions
- **Classic solution:**
 - Semaphores (Edsger Dijkstra)
- **Other approaches (out of our scope)**
 - Mutex and condition variables (Pthreads)
 - Locks and rwlocks (Pthreads)
 - Monitors (Java)

28

Today

- Threads review
- Sharing
- Mutual exclusion
- Semaphores

29

Semaphores

- **Semaphore:** non-negative global integer synchronization variable
- **Manipulated by *P* and *V* operations:**
 - *P*(*s*): [while (*s* == 0) wait(); *s*--;]
 - Dutch for "Proberen" (test)
 - *V*(*s*): [*s*++;]
 - Dutch for "Verhogen" (increment)
- **OS kernel 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)**

30

C Semaphore Operations

Pthreads functions:

```
#include <semaphore.h>

int sem_init(sem_t *sem, 0, unsigned int val); /* s = val */

int sem_wait(sem_t *s); /* P(s) */
int sem_post(sem_t *s); /* V(s) */
```

CS:APP wrapper functions:

```
#include "csapp.h"

void P(sem_t *s); /* Wrapper function for sem_wait */
void V(sem_t *s); /* Wrapper function for sem_post */
```

31

badcnt.c: Improper Synchronization

```
volatile int cnt = 0; /* global */

/* Thread routine */
void *thread(void *vargp)
{
    int i, niters = *((int *)vargp);
    for (i = 0; i < niters; i++)
        cnt++;
    return NULL;
}

int main(int argc, char **argv)
{
    int niters = atoi(argv[1]);
    pthread_t tid1, tid2;

    Pthread_create(&tid1, NULL,
                  thread, &niters);
    Pthread_create(&tid2, NULL,
                  thread, &niters);
    Pthread_join(tid1, NULL);
    Pthread_join(tid2, NULL);

    /* Check result */
    if (cnt != (2 * niters))
        printf("BOOM! cnt=%d\n", cnt);
    else
        printf("OK cnt=%d\n", cnt);
    exit(0);
}
```

How can we fix this using semaphores?

32

Using Semaphores for Mutual Exclusion

Basic idea:

- Associate a unique semaphore *mutex*, initially 1, with each shared variable (or related set of shared variables).
- Surround corresponding critical sections with $P(mutex)$ and $V(mutex)$ operations.

Terminology:

- Binary semaphore:** semaphore whose value is always 0 or 1
- Mutex:** binary semaphore used for mutual exclusion
 - P operation: “locking” the mutex
 - V operation: “unlocking” or “releasing” the mutex
 - “Holding” a mutex: locked and not yet unlocked.
- Counting semaphore:** used as a counter for set of available resources.

33

goodcnt.c: Proper Synchronization

- Define and initialize a mutex for the shared variable *cnt*:

```
volatile int cnt = 0; /* Counter */
sem_t mutex; /* Semaphore that protects cnt */

Sem_init(&mutex, 0, 1); /* mutex = 1 */
```

- Surround critical section with P and V :

```
for (i = 0; i < niters; i++) {
    P(&mutex);
    cnt++;
    V(&mutex);
}
```

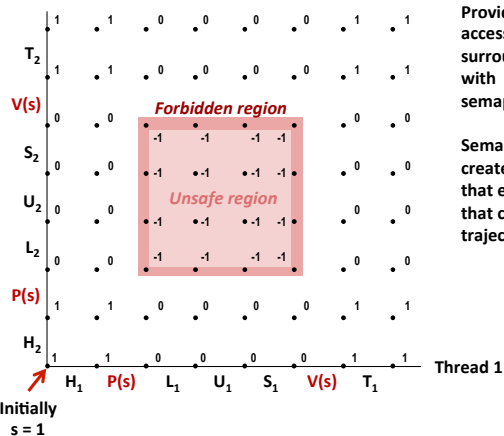
```
linux> ./goodcnt 10000
OK cnt=20000
linux> ./goodcnt 10000
OK cnt=20000
linux>
```

Warning: It's much slower than *badcnt.c*.

34

Why Mutexes Work

Thread 2



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 that cannot be entered by any trajectory.

35

Summary

- Programmers need a clear model of how variables are shared by threads.
- Variables shared by multiple threads must be protected to ensure mutually exclusive access.
- Semaphores are a fundamental mechanism for enforcing mutual exclusion.

36