

Lecture 9

Programming Shared Memory III

Synchronization Primitives; Condition Variables

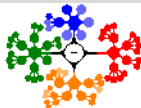
Ceng505 *Parallel Computing* at December 06, 2010

Thread Examples

Computing the value of π
Producer-consumer work
queues

Condition Variables for
Synchronization

Dr. Cem Özdoğan
Computer Engineering Department
Çankaya University



1 Thread Examples

Computing the value of π

Producer-consumer work queues

Thread Examples

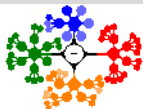
Computing the value of π
Producer-consumer work
queues

Condition Variables for Synchronization

2 Condition Variables for Synchronization

Computing the value of π I

- **Computing the value of π .**
- Based on generating random numbers in a *unit length square* and counting the number of points that fall within the largest circle inscribed in the square.
- Since the area of the circle (πr^2) is equal to $\pi/4$, and the area of the square is 1×1 , the fraction of random points that fall in the circle should approach $\pi/4$.
- **Threaded strategy:**
 - assigns a fixed number of points to each thread.
 - Each thread generates these random points and keeps track of the number of points in the circle locally.
 - After all threads finish execution, their counts are combined to compute the value of π (by calculating the fraction over all threads and multiplying by 4).



Thread Examples

Computing the value of π

Producer-consumer work
queues

Condition Variables for Synchronization

Computing the value of π II

```
51 void *compute_pi (void *s) {
52     int seed, i, *hit_pointer;
53     double rand_no_x, rand_no_y;
54     int local_hits;
55
56     hit_pointer = (int *) s;
57     seed = *hit_pointer;
58     local_hits = 0;
59     for (i = 0; i < sample_points_per_thread; i++) {
60         rand_no_x = (double) (rand_r(&seed)) / (double) ((2<<14) - 1);
61         rand_no_y = (double) (rand_r(&seed)) / (double) ((2<<14) - 1);
62         if (((rand_no_x - 0.5) * (rand_no_x - 0.5) +
63             (rand_no_y - 0.5) * (rand_no_y - 0.5)) < 0.25)
64             local_hits ++;
65         seed += i;
66     }
67     *hit_pointer = local_hits;
68     pthread_exit(0);
69 }
```

The *arg* field is used to pass an integer id that is used as a seed for randomization.



Thread Examples

Computing the value of π

Producer-consumer work
queues

Condition Variables for Synchronization

Computing the value of π III

```
1  #include <pthread.h>
2  #include <stdlib.h>
3  #define MAX_THREADS 512
4  //cat /proc/sys/kernel/threads-max
5  void*compute_pi (void*);
6  int total_hits, total_misses, hits[MAX_THREADS],
7     sample_points, sample_points_per_thread,
8     num_threads;
9
10 main() {
11     int i;
12     pthread_t p_threads[MAX_THREADS];
13     pthread_attr_t attr;
14     double computed_pi;
15     double time_start, time_end;
16     struct timeval tv;
17     struct timezone tz;
18
19     pthread_attr_init (&attr);
20     pthread_attr_setscope (&attr,
21                            PTHREAD_SCOPE_SYSTEM);
22     printf("Enter number of sample points: ");
23     scanf("%d", &sample_points);
24     printf("Enter number of threads: ");
25     scanf("%d", &num_threads);
```



Thread Examples

Computing the value of π

Producer-consumer work
queues

Condition Variables for
Synchronization

Computing the value of π IV

```
26     gettimeofday(&tv, &tz);
27     time_start = (double)tv.tv_sec +
28                 (double)tv.tv_usec / 1000000.0;
29
30     total_hits = 0;
31     sample_points_per_thread=sample_points/num_threads;
32     for (i=0; i< num_threads; i++) {
33         hits[i] = i;
34         pthread_create(&p_threads[i], &attr, compute_pi,
35                      (void *) &hits[i]);
36     }
37     for (i=0; i< num_threads; i++) {
38         pthread_join(p_threads[i], NULL);
39         total_hits += hits[i];
40     }
41     computed_pi = 4.0*(double) total_hits /
42                 ((double)(sample_points));
43     gettimeofday(&tv, &tz);
44     time_end = (double)tv.tv_sec +
45              (double)tv.tv_usec / 1000000.0;
46
47     printf("Computed PI = %lf\n", computed_pi);
48     printf(" %lf\n", time_end - time_start);
49 }
50
```



Thread Examples

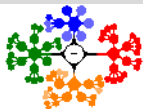
Computing the value of π

Producer-consumer work
queues

Condition Variables for
Synchronization

Computing the value of π V

- For computing the value of π ,
- First read in the desired number of threads, *num_threads*, and the desired number of sample points, *sample_points*.
- These points are divided equally among the threads.
- The program uses an array, **hits**, for assigning an integer id to each thread (this id is used as a seed for randomizing the random number generator).
- The same array is used to keep track of the number of hits (points inside the circle) encountered by each thread upon return.



Thread Examples

Computing the value of π

Producer-consumer work
queues

Condition Variables for Synchronization

Computing the value of π VI

- The program creates *num_threads* threads, each invoking the same function *compute_pi*, using the **pthread_create** function.
- Once the respective *compute_pi* threads have generated assigned number of random points and computed their hit ratios, the results must be combined to determine π .
- Once all threads have joined, the value of π is computed by multiplying the combined hit ratio by 4.0.
- The use of the function *rand_r* (instead of superior random number generators such as *drand48*).
- The reason for this is that many functions (including *rand* and *drand48*) are not reentrant.

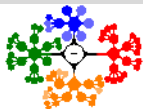


Thread Examples

Computing the value of π

Producer-consumer work
queues

Condition Variables for Synchronization



Thread Examples

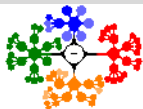
Computing the value of π

Producer-consumer work
queues

Condition Variables for
Synchronization

- **Producer-consumer work queues**
- A common use of mutex-locks is in establishing a producer-consumer relationship between threads.
- The **producer creates tasks and inserts** them into a work-queue.
- The **consumer threads pick up tasks** from the task queue and execute them.
- Consider that the task queue can hold only one task.
- In a general case, the task queue may be longer but is typically of bounded size.
- A simple (and incorrect) threaded program would associate a producer thread with creating a task
- and placing it in a **shared data structure**
- and the consumer threads with picking up tasks from this shared data structure and executing them.

Producer-consumer work queues II



Thread Examples

Computing the value of π

Producer-consumer work
queues

Condition Variables for
Synchronization

- However, this simple version does not account for the following possibilities:
 - 1 The producer thread must not overwrite the shared buffer when the previous task has not been picked up by a consumer thread.
 - 2 The consumer threads must not pick up tasks until there is something present in the shared data structure.
 - 3 Individual consumer threads should pick up tasks one at a time.
- To implement this, we can use a variable called *task_available*.
 - If this variable is 0, consumer threads must wait, but the producer thread can insert tasks into the shared data structure *task_queue*.
 - If *task_available* is equal to 1, the producer thread must wait to insert the task into the shared data structure but one of the consumer threads can pick up the task available.

Producer-consumer work queues III



All of these operations on the variable *task_available* should be protected by **mutex-locks** to ensure that only one thread is executing test-update on it.

```
1  pthread_mutex_t task_queue_lock;
2  int task_available;
3
4  /* other shared data structures here */
5
6  main() {
7      /* declarations and initializations */
8      task_available = 0;
9      pthread_init();
10     pthread_mutex_init(&task_queue_lock, NULL);
11 /* create and join producer and consumer threads */
12 }
13
```

Thread Examples

Computing the value of π

Producer-consumer work
queues

Condition Variables for Synchronization

Producer-consumer work queues IV



Thread Examples

Computing the value of π

Producer-consumer work
queues

Condition Variables for Synchronization

```
14 void *producer(void *producer_thread_data) {
15     int inserted;
16     struct task my_task;
17     while (!done()) {
18         inserted = 0;
19         create_task(&my_task);
20         while (inserted == 0) {
21             pthread_mutex_lock(&task_queue_lock);
22             if (task_available == 0) {
23                 insert_into_queue(my_task);
24                 task_available = 1;
25                 inserted = 1;
26             }
27             pthread_mutex_unlock(&task_queue_lock);
28         }
29     }
30 }
31
```

Producer-consumer work queues V

```
32 void *consumer(void *consumer_thread_data) {
33     int extracted;
34     struct task my_task;
35     /* local data structure declarations */
36     while (!done()) {
37         extracted = 0;
38         while (extracted == 0) {
39             pthread_mutex_lock(&task_queue_lock);
40             if (task_available == 1) {
41                 extract_from_queue(&my_task);
42                 task_available = 0;
43                 extracted = 1;
44             }
45             pthread_mutex_unlock(&task_queue_lock);
46         }
47         process_task(my_task);
48     }
49 }
```



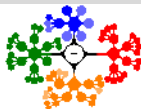
Thread Examples

Computing the value of π

Producer-consumer work
queues

Condition Variables for Synchronization

Producer-consumer work queues VI



- The *create_task* and *process_task* functions are left **outside the critical region**, making the critical section as small as possible.
- but *insert_into_queue* and *extract_from_queue* functions are left **inside the critical region**.
- Inside because if the lock is **relinquished** after updating *task_available* but not inserting or extracting the task,
- other threads may gain access to the shared data structure while the insertion or extraction is in progress, resulting in errors.

Thread Examples

Computing the value of π

Producer-consumer work
queues

Condition Variables for Synchronization

Producer-consumer work queues VII

- For producer-consumer work queues
- The producer thread creates a task and waits for space on the queue.
- This is indicated by the variable *task_available* being 0.
- The test and update of this variable as well as insertion and extraction from the shared queue are protected by a mutex called *task_queue_lock*.
- Once space is available on the task queue, the recently created task is inserted into the task queue and the availability of the task is signaled by setting *task_available* to 1.
- Within the producer thread, the fact that the recently created task has been inserted into the queue is signaled by the variable *inserted* being set to 1, which allows the producer to produce the next task.
- Irrespective of whether a recently created task is successfully inserted into the queue or not, the lock is relinquished.



Thread Examples

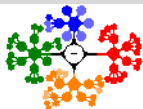
Computing the value of π

Producer-consumer work
queues

Condition Variables for
Synchronization

Producer-consumer work queues VIII

- This allows consumer threads to pick up work from the queue in case there is work on the queue to begin with.
- If the lock is not relinquished, threads would deadlock since a consumer would not be able to get the lock to pick up the task and the producer would not be able to insert its task into the task queue.
- The consumer thread waits for a task to become available and executes it when available.
- As was the case with the producer thread, the consumer relinquishes the lock in each iteration of the while loop to allow the producer to insert work into the queue if there was none.



Thread Examples

Computing the value of π

Producer-consumer work
queues

Condition Variables for Synchronization

Synchronization Primitives; Condition Variables I

- Indiscriminate use of locks can **result in idling overhead** from blocked threads.
- While the function **`pthread_mutex_trylock`** removes this overhead, it introduces the overhead of polling for availability of locks.
- For example, if the producer-consumer example is rewritten using *`pthread_mutex_trylock`* instead of *`pthread_mutex_lock`*,
- the producer and consumer threads would have to periodically poll for availability of lock (and subsequently availability of buffer space or tasks on queue).
- A natural solution to this problem is to **suspend the execution** of the polling thread until space becomes available.
- An **interrupt driven mechanism** as opposed to a **polled mechanism**.

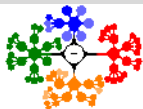


Thread Examples

Computing the value of π
Producer-consumer work
queues

Condition Variables for Synchronization

Synchronization Primitives; Condition Variables II



- The availability of space is signaled by the thread that holding the space.
- The functionality to accomplish this is provided by a **condition variable**.
- A condition variable is a data object used for synchronizing threads and always used in conjunction with a mutex lock.
- While mutexes implement synchronization by **controlling thread access to data**,
- condition variables allow threads to synchronize based upon **the actual value of data**.
- This variable allows a thread to block itself until specified data reaches a predefined state.

Thread Examples

Computing the value of π
Producer-consumer work
queues

Condition Variables for Synchronization



Thread Examples

Computing the value of π
Producer-consumer work
queues

Condition Variables for Synchronization

- `pthread_cond_wait`

```
1  int pthread_cond_wait(pthread_cond_t *cond,  
2      pthread_mutex_t *mutex);
```

- A thread locks this mutex and tests the predicate defined on the shared variable;
- if the predicate is not true, the thread waits on the condition variable associated with the predicate using this function.
- A call to this function blocks the execution of the thread until it receives a signal from another thread or is interrupted by an OS signal.
- In addition to blocking the thread, the `pthread_cond_wait` function **releases the lock on mutex**.
- This is important because otherwise no other thread will be able to work on the shared variable and the predicate would never be satisfied.



- **pthread_cond_signal**

```
1      int pthread_cond_signal(pthread_cond_t  
2                                     *cond);
```

- When the condition is signaled, **pthread_cond_signal**, one of these threads in the queue is unblocked,
- and when the mutex becomes available, it is handed to this thread (and the thread becomes runnable).
- When the thread is released on a signal, it waits to reacquire the lock on mutex before resuming execution.
- It is convenient to think of each condition variable as being associated with a queue.
- Threads performing a condition wait on the variable relinquish their lock and enter the queue.

Thread Examples

Computing the value of π
Producer-consumer work
queues

Condition Variables for Synchronization



- `pthread_cond_init` & `pthread_cond_destroy`


```
1 int pthread_cond_init(pthread_cond_t *cond,  
2     const pthread_condattr_t *attr);  
3 int pthread_cond_destroy(pthread_cond_t *cond);
```

- Function calls for initializing and destroying condition variables.
- Condition variables must be declared with type `pthread_cond_t`, and must be initialized before they can be used.
- There are two ways to initialize a condition variable:
 - 1 Statically, when it is declared. For example:
`pthread_cond_t myconvar = PTHREAD_COND_INITIALIZER;`
 - 2 Dynamically, with the `pthread_cond_init()` routine.

Thread Examples

Computing the value of π
Producer-consumer work
queues

Condition Variables for Synchronization




- The function **pthread_cond_init** initializes a condition variable (pointed to by *cond*).
- The ID of the created condition variable is returned to the calling thread through the condition parameter.
- This method permits setting condition variable object attributes, *attr*. (*NULL* assigns default attributes)
- If at some point in a program a condition variable is no longer required, it can be discarded using the function **pthread_cond_destroy**.

Thread Examples

Computing the value of π
Producer-consumer work
queues

Condition Variables for Synchronization

Synchronization Primitives; Condition Variables VII



Thread Examples

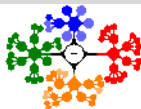
Computing the value of π
Producer-consumer work
queues

Condition Variables for Synchronization

Main Thread <ul style="list-style-type: none">○ Declare and initialize global data/variables which require synchronization (such as "count")○ Declare and initialize a condition variable object○ Declare and initialize an associated mutex○ Create threads A and B to do work	
Thread A <ul style="list-style-type: none">○ Do work up to the point where a certain condition must occur (such as "count" must reach a specified value)○ Lock associated mutex and check value of a global variable○ Call <code>pthread_cond_wait()</code> to perform a blocking wait for signal from Thread-B. Note that a call to <code>pthread_cond_wait()</code> automatically and atomically unlocks the associated mutex variable so that it can be used by Thread-B.○ When signalled, wake up. Mutex is automatically and atomically locked.○ Explicitly unlock mutex○ Continue	Thread B <ul style="list-style-type: none">○ Do work○ Lock associated mutex○ Change the value of the global variable that Thread-A is waiting upon.○ Check value of the global Thread-A wait variable. If it fulfills the desired condition, signal Thread-A.○ Unlock mutex.○ Continue
Main Thread Join / Continue	

Figure: A representative sequence for using condition variables.

Synchronization Primitives; Condition Variables VIII



Thread Examples

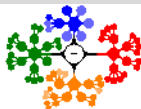
Computing the value of π
Producer-consumer work
queues

Condition Variables for Synchronization

- When a thread performs a condition wait, it takes itself off the runnable list consequently, it does **not use any CPU cycles** until it is woken up.
- This is in contrast to a mutex lock which **consumes** CPU cycles as it polls for the lock.
- **pthread_cond_broadcast**.

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

- In some cases, it may be beneficial to wake all threads that are waiting on the condition variable as opposed to a single thread.
- An example of this is in the producer-consumer scenario with large work queues and multiple tasks being inserted into the work queue on each insertion cycle.
- Another example is in the implementation of barriers.



- `pthread_cond_timedwait`,

```
1 int pthread_cond_timedwait(pthread_cond_t *cond,  
2     pthread_mutex_t *mutex,  
3     const struct timespec *abstime);
```

- It is often useful to build time-outs into condition waits.
- Using the function a thread can perform a wait on a condition variable until a specified time expires.
- At this point, the thread wakes up by itself if it does not receive a signal or a broadcast.
- If the absolute time *abstime* specified expires before a signal or broadcast is received, the function returns an error message.
- It also reacquires the lock on mutex when it becomes available.

Thread Examples

Computing the value of π
Producer-consumer work
queues

Condition Variables for Synchronization