



Introduction to Parallel Computing

George Karypis

Programming Shared Address Space
Platforms



Outline

- Shared Address-Space Programming Models
- Thread-based programming
 - POSIX API/Pthreads
- Directive-based programming
 - OpenMP API



Shared Memory Programming

- Communication is implicitly specified
- Focus on constructs for expressing concurrency and synchronization
 - Minimize data-sharing overheads



Commonly Used Models

- Process model
 - All memory is local unless explicitly specified/allocated as shared.
 - Unix processes.
- Light-weight process/thread model
 - All memory is global and can be accessed by all the threads.
 - Runtime stack is local but it can be shared.
 - POSIX thread API/Pthreads
 - Low-level & system-programming flavor to it.
- Directive model
 - Concurrency is specified in terms of high-level compiler directives.
 - High-level constructs that leave some of the error-prone details to the compiler.
 - OpenMP has emerged as a standard.



POSIX API/Pthreads

- Has emerged as the de-facto standard supported by most OS vendors.
 - Aids in the portability of threaded applications.
- Provides a good set of functions that allow for the creation, termination, and synchronization of threads.
 - However, these functions are low-level and the API is missing some high-level constructs for efficient data-sharing
 - There are no collective communication operation like those provided by MPI.



Pthreads Overview

- Thread creation & termination
- Synchronization primitives
 - Mutual exclusion locks
 - Conditional variables
- Object attributes

Thread Creation & Termination

```
1 #include <pthread.h>
2 int
3 pthread_create (
4     pthread_t *thread_handle,
5     const pthread_attr_t *attribute,
6     void * (*thread_function)(void *),
7     void *arg);
```

The `pthread_create` function creates a single thread that corresponds to the invocation of the function `thread_function` (and any other functions called by `thread_function`). On successful creation of a thread, a unique identifier is associated with the thread and assigned to the location pointed to by `thread_handle`. The thread has the attributes described by the `attribute` argument. When this argument is `NULL`, a thread with default attributes is created. We will discuss the `attribute` parameter in detail in Section 7.6. The `arg` field specifies a pointer to the argument to function `thread_function`. This argument is typically used to pass the workspace and other thread-specific data to a thread. In the `compute_pi` example, it is used to pass an integer id that is used as a seed for randomization. The `thread_handle` variable is written before the the function `pthread_create` returns; and the new thread is ready for execution as soon as it is created. If the thread is scheduled on the same processor, the new thread may, in fact, preempt its creator. This is important to note because all thread initialization procedures must be completed before creating the thread. Otherwise, errors may result based on thread scheduling. This is a very common class of errors caused by race conditions for data access that shows itself in some execution instances, but not in others. On successful creation of a thread, `pthread_create` returns 0; else it returns an error code. The reader is referred to the Pthreads specification for a detailed description of the error-codes.

```
1 int
2 pthread_join (
3     pthread_t thread,
4     void **ptr);
```

A call to this function waits for the termination of the thread whose id is given by `thread`. On a successful call to `pthread_join`, the value passed to `pthread_exit` is returned in the location pointed to by `ptr`. On successful completion, `pthread_join` returns 0, else it returns an error-code.

Computing the value of π

```
1 #include <pthread.h>
2 #include <stdlib.h>
3
4 #define MAX_THREADS    512
5 void *compute_pi (void *);
6
7 int total_hits, total_misses, hits[MAX_THREADS],
8     sample_points, sample_points_per_thread, 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, PTHREAD_SCOPE_SYSTEM);
21     printf("Enter number of sample points: ");
22     scanf("%d", &sample_points);
23     printf("Enter number of threads: ");
24     scanf("%d", &num_threads);
25
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
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 }
```


Synchronization Primitives

- Access to shared variable need to be controlled to remove race conditions and ensure serial semantics.

```
1 /* each thread tries to update variable best_cost as follows */
2 if (my_cost < best_cost)
3     best_cost = my_cost;
```

To understand the problem with shared data access, let us examine one execution instance of the above code fragment. Assume that there are two threads, the initial value of `best_cost` is 100, and the values of `my_cost` are 50 and 75 at threads `t1` and `t2`, respectively. If both threads execute the condition inside the `if` statement concurrently, then both threads enter the `then` part of the statement. Depending on which thread executes first, the value of `best_cost` at the end could be either 50 or 75. There are two problems here: the first is the non-deterministic nature of the result; second, and more importantly, the value 75 of `best_cost` is inconsistent in the sense that no serialization of the two threads can possibly yield this result. This is an undesirable situation, sometimes also referred to as a race condition (so called because the result of the computation depends on the race between competing threads).

Mutual Exclusion Locks

- Pthreads provide a special variable called a *mutex* lock that can be used to guard critical sections of the program.
 - The idea is for a thread to acquire the lock before entering the critical section and release on exit.
 - If the lock is already owned by another thread, the thread blocks until the lock is released.
- Lock represent serialization points, so too many locks can decrease the performance.

```
1 int
2 pthread_mutex_lock (
3     pthread_mutex_t *mutex_lock);
```

A call to this function attempts a lock on the mutex-lock `mutex_lock`. (The data type of a `mutex_lock` is predefined to be `pthread_mutex_t`.) If the mutex-lock is already locked, the calling thread blocks; otherwise the mutex-lock is locked and the calling thread returns. A successful return from the function returns a value 0. Other values indicate error conditions such as deadlocks.

```
1 int
2 pthread_mutex_unlock (
3     pthread_mutex_t *mutex_lock);
```

On calling this function, in the case of a normal mutex-lock, the lock is relinquished and one of the blocked threads is scheduled to enter the critical section. The specific thread is determined by the scheduling policy. There are other types of locks (other than normal locks), which are discussed in Section 7.6 along with the associated semantics of the function `pthread_mutex_unlock`. **If a programmer attempts a `pthread_mutex_unlock` on a previously unlocked mutex or one that is locked by another thread, the effect is undefined.**

```
1 int
2 pthread_mutex_init (
3     pthread_mutex_t *mutex_lock,
4     const pthread_mutexattr_t *lock_attr);
```

This function initializes the mutex-lock `mutex_lock` to an unlocked state. The attributes of the mutex-lock are specified by `lock_attr`. If this argument is set to NULL, the default mutex-lock attributes are used (normal mutex-lock). Attributes objects for threads are discussed in greater detail in Section 7.6.

It is often possible to reduce the idling overhead associated with locks using an alternate function, `pthread_mutex_trylock`. This function attempts a lock on `mutex_lock`. If the lock is successful, the function returns a zero. **If it is already locked by another thread, instead of blocking the thread execution, it returns a value `EBUSY`.** This allows the thread to do other work and to poll the mutex for a lock. Furthermore, `pthread_mutex_trylock` is typically much faster than `pthread_mutex_lock` on typical systems since it does not have to deal with queues associated with locks for multiple threads waiting on the lock. The prototype of `pthread_mutex_trylock` is:

```
1 int
2 pthread_mutex_trylock (
3     pthread_mutex_t *mutex_lock);
```

Computing the minimum element of an array.

```
1  #include <pthread.h>
2  void *find_min(void *list_ptr);
3  pthread_mutex_t minimum_value_lock;
4  int minimum_value, partial_list_size;
5
6  main() {
7      /* declare and initialize data structures and list */
8      minimum_value = MIN_INT;
9      pthread_init();
10     pthread_mutex_init(&minimum_value_lock, NULL);
11
12     /* initialize lists, list_ptr, and partial_list_size */
13     /* create and join threads here */
14 }
15
16 void *find_min(void *list_ptr) {
17     int *partial_list_pointer, my_min, i;
18     my_min = MIN_INT;
19     partial_list_pointer = (int *) list_ptr;
20     for (i = 0; i < partial_list_size; i++)
21         if (partial_list_pointer[i] < my_min)
22             my_min = partial_list_pointer[i];
23     /* lock the mutex associated with minimum_value and
24     update the variable as required */
25     pthread_mutex_lock(&minimum_value_lock);
26     if (my_min < minimum_value)
27         minimum_value = my_min;
28     /* and unlock the mutex */
29     pthread_mutex_unlock(&minimum_value_lock);
30     pthread_exit(0);
31 }
```

Producer Consumer Queues

```
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
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
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 }
```



Conditional Variables

- Waiting-queue like synchronization principles.
 - Based on the outcome of a certain condition a thread may attach itself to a waiting queue.
 - At a later point in time, another thread that change the outcome of the condition, will wake up one/all of the threads so that they can see if they can proceed.
- Conditional variables are always associated with a mutex lock.

Conditional Variables API

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

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 `task_available` and the predicate would never be satisfied. 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. When the condition is signaled (using `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).

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

The function unblocks at least one thread that is currently waiting on the condition variable `cond`. The producer then relinquishes its lock on `mutex` by explicitly calling `pthread_mutex_unlock`, allowing one of the blocked consumer threads to consume the task.

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

The function `pthread_cond_init` initializes a condition variable (pointed to by `cond`) whose attributes are defined in the attribute object `attr`. Setting this pointer to `NULL` assigns default attributes for condition variables. If at some point in a program a condition variable is no longer required, it can be discarded using the function `pthread_cond_destroy`. These functions for manipulating condition variables enable us to rewrite our producer-consumer segment as follows:

In the above example, each task could be consumed by only one consumer thread. Therefore, we choose to signal one blocked thread at a time. In some other computations, it may be beneficial to wake all threads that are waiting on the condition variable as opposed to a single thread. This can be done using the function `pthread_cond_broadcast`.

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

It is often useful to build time-outs into condition waits. Using the function `pthread_cond_timedwait`, 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. The prototype for this function is:

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

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.

Producer Consumer Example with Conditional Variables

```
1 pthread_cond_t cond_queue_empty, cond_queue_full;
2 pthread_mutex_t task_queue_cond_lock;
3 int task_available;
4
5 /* other data structures here */
6
7 main() {
8     /* declarations and initializations */
9     task_available = 0;
10    pthread_init();
11    pthread_cond_init(&cond_queue_empty, NULL);
12    pthread_cond_init(&cond_queue_full, NULL);
13    pthread_mutex_init(&task_queue_cond_lock, NULL);
14    /* create and join producer and consumer threads */
15 }
16
17 void *producer(void *producer_thread_data) {
18     int inserted;
19     while (!done()) {
20         create_task();
21         pthread_mutex_lock(&task_queue_cond_lock);
22         while (task_available == 1)
23             pthread_cond_wait(&cond_queue_empty,
24                               &task_queue_cond_lock);
25         insert_into_queue();
26         task_available = 1;
27         pthread_cond_signal(&cond_queue_full);
28         pthread_mutex_unlock(&task_queue_cond_lock);
29     }
30 }
31
32 void *consumer(void *consumer_thread_data) {
33     while (!done()) {
34         pthread_mutex_lock(&task_queue_cond_lock);
35         while (task_available == 0)
36             pthread_cond_wait(&cond_queue_full,
37                               &task_queue_cond_lock);
38         my_task = extract_from_queue();
39         task_available = 0;
40         pthread_cond_signal(&cond_queue_empty);
41         pthread_mutex_unlock(&task_queue_cond_lock);
42         process_task(my_task);
43     }
44 }
```



Attribute Objects

- Various attributes can be associated with threads, locks, and conditional variables.
 - Thread attributes:
 - scheduling parameters
 - stack size
 - detached state
 - Mutex attributes:
 - normal
 - only a single thread is allowed to lock it.
 - if a threads tries to lock it twice a deadlock occurs.
 - recursive
 - a thread can lock the mutex multiple time.
 - each successive lock increments a counter and each successive release decrements the counter.
 - a thread can lock a mutex only if its counter is zero.
 - errorcheck
 - like normal but an attempt to lock it again by the same thread leads to an error.
 - The book and the Posix thread API provide additional details.



OpenMP

- A standard directive-based shared memory programming API
 - C/C++/Fortran versions of the API exist
- API consists of a set of compiler directive along with a set of API functions.

```
1  #pragma omp directive [clause list]
```

Parallel Region

- Parallel regions are specified by the `parallel` directive:

```
1  #pragma omp parallel [clause list]
2  /* structured block */
3
```

- The *clause list* contains information about:

- conditional parallelization

- `if (scalar expression)`

- degree of concurrency

- `num_threads (integer expression)`

- data handling

- `private (var list), firstprivate (var list), shared (var list)`

- `default(shared|private|none)`

```
1  #pragma omp parallel if (is_parallel == 1) num_threads(8) \
2  private (a) shared (b) firstprivate(c)
3  {
4  /* structured block */
5  }
```

Reduction clause

Just as `firstprivate` specifies how multiple local copies of a variable are initialized inside a thread, the `reduction` clause specifies how multiple local copies of a variable at different threads are combined into a single copy at the master when threads exit. The usage of the `reduction` clause is `reduction (operator: variable list)`. This clause performs a reduction on the scalar variables specified in the list using the operator. The variables in the list are implicitly specified as being private to threads. The operator can be one of `+`, `*`, `-`, `&`, `|`, `^`, `&&`, and `||`.

Example 7.10 Using the `reduction` clause

```
1      #pragma omp parallel reduction(+: sum) num_threads(8)
2      {
3          /* compute local sums here */
4      }
5      /* sum here contains sum of all local instances of sums */
```

In this example, each of the eight threads gets a copy of the variable `sum`. When the threads exit, the sum of all of these local copies is stored in the single copy of the variable (at the master thread). ■

Computing the value of π

```
1  /* *****  
2  An OpenMP version of a threaded program to compute PI.  
3  ***** */  
4  
5  #pragma omp parallel default(private) shared (npoints) \  
6      reduction(+: sum) num_threads(8)  
7  {  
8      num_threads = omp_get_num_threads();  
9      sample_points_per_thread = npoints / num_threads;  
10     sum = 0;  
11     for (i = 0; i < sample_points_per_thread; i++) {  
12         rand_no_x = (double) (rand_r(&seed)) / (double) ((2<<14) - 1);  
13         rand_no_y = (double) (rand_r(&seed)) / (double) ((2<<14) - 1);  
14         if (((rand_no_x - 0.5) * (rand_no_x - 0.5) +  
15             (rand_no_y - 0.5) * (rand_no_y - 0.5)) < 0.25)  
16             sum ++;  
17     }  
18 }
```



Specifying concurrency

- Concurrent tasks are specified using the `for` and `sections` directives.
 - The `for` directive splits the iterations of a loop across the different threads.
 - The `sections` directive assigns each thread to explicitly identified tasks.

The for directive

The `for` directive is used to split parallel iteration spaces across threads. The general form of a `for` directive is as follows:

```
1      #pragma omp for [clause list]
2      /* for loop */
3
```

The clauses that can be used in this context are: `private`, `firstprivate`, `lastprivate`, `reduction`, `schedule`, `nowait`, and `ordered`. The first four clauses deal with data handling and have identical semantics as in the case of the `parallel` directive. The `lastprivate` clause deals with how multiple local copies of a variable are written back into a single copy at the end of the parallel `for` loop. When using a `for` loop (or `sections` directive as we shall see) for farming work to threads, it is sometimes desired that the last iteration (as defined by serial execution) of the `for` loop update the value of a variable. This is accomplished using the `lastprivate` directive.

The `schedule` clause of the `for` directive deals with the assignment of iterations to threads. The general form of the `schedule` directive is `schedule(scheduling_class[, parameter])`. OpenMP supports four scheduling classes: `static`, `dynamic`, `guided`, and `runtime`.

Often, it is desirable to have a sequence of `for`-directives within a parallel construct that do not execute an implicit barrier at the end of each `for` directive. OpenMP provides a clause – `nowait`, which can be used with a `for` directive to indicate that the threads can proceed to the next statement without waiting for all other threads to complete the `for` loop execution. This is illustrated in the following example:

An example

```
1  #pragma omp parallel default(private) shared (npoints) \  
2      reduction(+: sum) num_threads(8)  
3  {  
4      sum = 0;  
5      #pragma omp for  
6      for (i = 0; i < npoints; i++) {  
7          rand_no_x =(double) (rand_r(&seed))/(double) ((2<<14)-1);  
8          rand_no_y =(double) (rand_r(&seed))/(double) ((2<<14)-1);  
9          if (((rand_no_x - 0.5) * (rand_no_x - 0.5) +  
10             (rand_no_y - 0.5) * (rand_no_y - 0.5)) < 0.25)  
11              sum ++;  
12      }  
13 }
```

The loop index for the `for` directive is assumed to be private.



More one `for` directive

■ Loop scheduling schemes

- `schedule(static[, chunk-size])`
 - splits the iterations into consecutive chunks of size `chunk-size` and assigns them in round-robin fashion.
- `schedule(dynamic[, chunk-size])`
 - splits the iterations into consecutive chunks of size `chunk-size` and gives to each thread a chunk as soon as it finishes processing its previous chunk.
- `schedule(guided[, chunk-size])`
 - like `dynamic` but the `chunk-size` is reduced exponentially as each chunk is dispatched to a thread.
- `schedule(runtime)`
 - is determined by reading an environmental variable.



Restrictions on the `for` directive

- For loops must not have break statements.
- Loop control variables must be integers.
- The initialization expression of the control variable must be an integer.
- The logical expression must be one of `<`, `<=`, `>`, `>=`.
- The increment expression must have integer increments and decrements.

The sections directive

```
1      #pragma omp parallel
2      {
3          #pragma omp sections
4          {
5              #pragma omp section
6              {
7                  taskA();
8              }
9              #pragma omp section
10             {
11                 taskB();
12             }
13             #pragma omp section
14             {
15                 taskC();
16             }
17         }
18     }
```

If there are three threads, each section (in this case, the associated task) is assigned to one thread. At the end of execution of the assigned section, the threads synchronize (unless the `nowait` clause is used). **Note that it is illegal to branch in and out of section blocks.**

Synchronization Directives

■ barrier directive

```
1      #pragma omp barrier
```

■ single/master directives

```
1      #pragma omp single [clause list]
2      structured block
```

```
1      #pragma omp master
2      structured block
```

■ critical/atomic directives

```
1      #pragma omp critical [(name)]
2      structured block
```

■ ordered directive

```
1      #pragma omp ordered
2      structured block
```