

High Performance Computing - Session 3

Shared Memory Programming with `pthread`s

Jayanti Prasad

<http://www.iucaa.ernet.in/jayanti/>

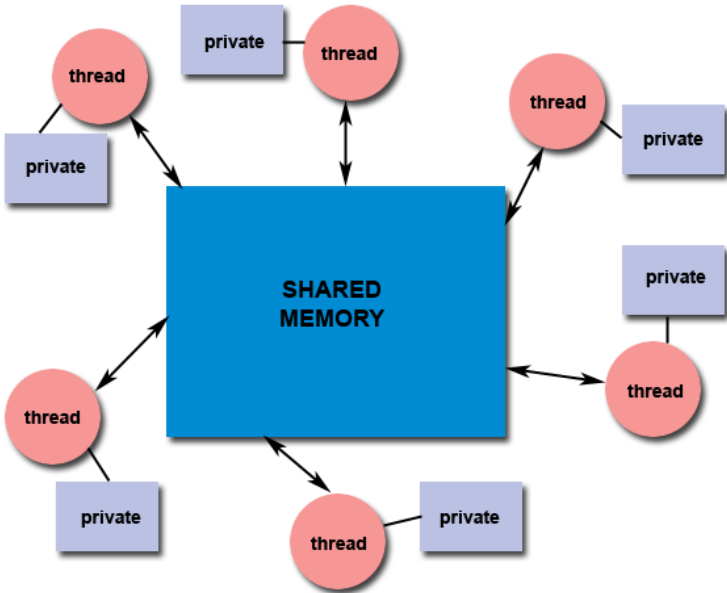
Inter-University Centre for Astronomy & Astrophysics
Pune, India (411007)

February 01, 2012

- Why we should bother about threads ?
 - ① At present almost all desktops/laptops come with multi-core processors which can be fully exploited only by multi-thread programming.
 - ② Multi-thread programming is much more efficient than multi-processors programming (MPI) due to faster communication (P-M : 30 GBPS, P-P: 5 GBPS).
 - ③ Multi-GPU systems and clusters with multi-core processors can be used more efficiently by incorporating multi-thread programming (hybrid programming).
- Why `pthread` ?
 - ① `pthread` provide better control on threads.
 - ② `pthread` are more natural to a Linux system with C programming.
 - ③ `pthread` are very light and versatile.

Plan of the Talk

- Introduction
 - ▶ Processes
 - ▶ Threads
- Posix threads
 - ▶ What is `pthread`
 - ▶ Creating and joining `pthread`
 - ▶ Race condition and Mutex locks
- Examples
- Problems



Linux Processes

- A process on a Linux system is an “object” through which the resources used by a program like memory, processor time and IO are managed and monitored.
- Processes are building blocks of any Linux system and can run in the “kernel” space or in the “user space” .
- A process consists of an address space (a set of memory pages) and a set of data structures.
- The address space of a process contains the code and libraries that the process is executing, the process variables, its stacks, and various extra information needed by the kernel while the process is running.
- Some of the common Linux command to monitor and manage processes are `ps -ef`, `top`, `strace`, `kill` etc.

Process

Program = Instruction + data

Process = Program in action

Process address contains

- The current status of the process (sleeping, stopped, runnable, etc.)
 - The execution priority of the process
 - Information about the resources the process has used
 - Information about the files and network ports the process has opened
 - The process signal mask (a record of which signals are blocked)
 - The owner of the process
-
- In a Linux/Unix system processes execute asynchronously (independent from each other) even when there is only one processor.
 - Processes are created using **fork** command on a Linux/Unix system.
 - Processes can be killed by **kill -9** command on a Linux/Unix system.

Threads

- Threads are *light weight* sub-processes within a process (examples). For example, when we open a new tab in Internet browser we launch a new thread.
- Threads inherit many attributes of a the parent process (such as the processs address space).
- Multiple threads can execute concurrently (may be not at the same time) within a single process under a model called **multi-threading**.
- In a multi-thread process the processor can switch execution resources between threads, resulting in concurrent execution.
- Concurrency on a single processor (core) system indicates that more than one thread is making progress, but the threads are not actually running simultaneously.
- On a multi-core system each thread in the process can run concurrently on a separate core i.e., true parallelism.

What is `pthread` ?

- `pthread` (Portable Operating System Interface Threads) is an Application Programming Interface (API) or library, which can be used for shared memory/address space programming.
- `pthread` library provides more than one hundred functions to manage threads, however, a very few (less than 10) are in general commonly used.
- When a thread is created, a new thread of control is added to the current process.
- Every process has at least one thread of control, in the program's `main()` routine.
- Each thread in the process runs simultaneously, and has access to the calling process's global data.
- In addition each thread has its own private attributes and call stack.

Managing `pthread`s

- To create a new thread, a running thread calls the `pthread_create()` function, and passes a **pointer to a function** for the new thread to run.
- **One argument** for the new thread's function can also be passed, along with thread attributes.
- The execution of a thread begins with the successful return from the `pthread_create()` function.
- The thread ends when the function that was called with the thread completes normally.
- A thread can also be terminated if the thread calls a `pthread_exit()` routine, or if any other thread calls `pthread_cancel()` to explicitly terminate that thread.
- When two or more concurrently running threads access a shared data item and the final result depends on the order of execution, we have a **race condition**.

Creating pthreads

```
1  #include <pthread.h>
2
3  pthread_t  thread[num_threads];
4
5  void *foo(void *);
6  void *arg;
7
8  int i;
9
10 for(i=0; i < num_threads; i++){
11     pthread_create(&thread[i], NULL, foo, arg);
12 }
```

Compilation

```
gcc program.c -lpthread -lm
```

Joining threads

```
1  #include <pthread.h>
2
3  pthread_t thread[num_threads];
4
5  void *status;
6  /* waiting to join thread "tid" with status */
7
8  for(i=0; i < num_threads; i++){
9      pthread_join(thread[i], &status);
10 }
11
12 /* waiting to join thread "tid" without status */
13
14 for(i=0; i < num_threads; i++){
15     pthread_join(thread[i], NULL);
16 }
```

Race condition

Simultaneous updates lead to race conditions. Suppose two threads both execute `i++`. In machine code, this single statement becomes several operations:

Thread 1	Thread 2
<code>r1 = i</code>	<code>r2 = i</code>
<code>r1 += 1</code>	<code>r2 += 1</code>
<code>i = r1</code>	<code>i = r2</code>

If `i` starts at 0, what is the value of `i` after execution?

Race condition

Simultaneous updates lead to race conditions.

Thread 1

```
r1 = i
```

```
r1 += 1
```

```
i = r1
```

Thread 2

```
r2 = i
```

```
r2 += 1
```

```
i = r2
```

Final value of `i` is 2.

Race condition

Simultaneous updates lead to race conditions.

Thread 1

`r1 = i`

`r1 += 1`

`i = r1`

Thread 2

`r2 = i`

`r2 += 1`

`i = r2`

Final value of `i` is 1. The first update to `i` is lost.

Mutex Lock

```
1  #include <pthread.h>
2
3  pthread_mutex_t  mutex = PTHREAD_MUTEX_INITIALIZER;
4
5  pthread_mutex_lock(&mutex);
6
7  counter++;
8
9  pthread_mutex_unlock(&mutex);
```

Program 1 : hello_world1_pthreads.c

```
1  #include<stdio.h>
2  #include<stdlib.h>
3  #include<pthread.h> // must be included
4
5  //function which will be executed by every thread in parallel
6  void *tfunc(void *arg){ // one argument can be passed
7      long    my_id = (long) arg;
8      fprintf(stdout, " Hello World from %ld\n",my_id);
9  }
10
11 int main(int argc, char *argv[]){
12     int num_threads;
13     long i;
14     pthread_t *t; // thread object
15     if(argc < 2){
16         fprintf(stderr, "./hello_world <# threads>\n");
17         return(-1);
18     }
19     num_threads=atoi(argv[1]); // # of threads
20     t = (pthread_t *)malloc(num_threads*sizeof(pthread_t));
21     // create threads
22     for(i=0; i < num_threads; i++)
23         pthread_create(&t[i],NULL,tfunc,(void *)i);
24
25     //join threads
26     for(i=0; i < num_threads; i++)
27         pthread_join(t[i],NULL) ;
28
29     return(0);
30 }
```


Program 2 : hello_world2_pthreads.c

```
1  /* This structure can be used to send data to the function
2     executed by every thread in parallel */
3
4  typedef struct{
5     int thread_id;
6     char msg[100];
7 } thread_data;
8
9  // this function still done not return anything
10
11 void *tfunc(void *arg){
12     thread_data *p = (thread_data *)arg; // p is a pointer to a structure
13     int my_id = p->thread_id;
14     fprintf(stdout, " %s from %d\n",p->msg,my_id);
15 }
16
17 int main(int argc, char *argv[]){
18     pthread_t *t;
19     thread_data *q; // this is an array
20
21     q = (thread_data *)malloc(num_threads*sizeof(thread_data));
22
23     for(i=0; i < num_threads; i++){
24         q[i].thread_id = i;
25         sprintf(q[i].msg, "Hello World");
26         pthread_create(&t[i], NULL, tfunc, (void *) (q+i));
27     }
28
29     // join thread
30
31 }
```

Program 3 : return_pthreads.c

```
1  void* child_thread( void * param ){
2  long id, jd;
3  id = (long )param;
4  jd = id *id;
5  return (void *)jd;
6  }
7
8  int main(int argc, char *argv[]){
9  pthread_t *thread;
10 long i, *return_value;
11
12 thread = (pthread_t *)malloc(num_threads*sizeof(pthread_t));
13 return_value=(long *)malloc(num_threads*sizeof(long ));
14
15 for (i=0; i<num_threads; i++ )
16     pthread_create(&thread[i],NULL,&child_thread,(void*)i );
17
18 for (i=0; i<num_threads ; i++ ) {
19     pthread_join(thread[i], (void**)&return_value[i] );
20     printf( "input = %ld output=%ld \n",i, return_value[i] );
21 }
22 }
```

Program 4 : summation1_threads.c

```
1  void* child_thread(void * param ){
2      long id,i,p,y;
3      id = (long)param;
4      p = n/num_threads;
5      y = 0;
6      for(i=id*p; i <(id+1)*p; i++)
7          y+=x[i];
8      return (void *)y;
9  }
10
11 int main(int argc, char *argv[]){
12     pthread_t *thread;
13     long i,sum,y;
14
15     thread = (pthread_t *)malloc(num_threads*sizeof(pthread_t));
16     x=(long *)malloc(n*sizeof(long));
17
18     for(i=0; i < n; i++)
19         x[i] = (long) i;
20
21     for (i=0; i<num_threads; i++ )
22         pthread_create(&thread[i],0,&child_thread,(void*)i);
23
24     sum = 0;
25     for (i=0; i<num_threads ; i++ ) {
26         pthread_join(thread[i], (void**)&y);
27         sum+=y;
28     }
29     printf("sum=%ld\n",sum);
30 }
```

Program 5 : summation2_pthreads.c

```
1  float *x;
2  /*-----*/
3  typedef struct{
4      int thread_id;  int chunk;  float data;
5  } thread_data;
6  /*-----*/
7  void *tfunc(void *arg){
8      thread_data *p = (thread_data *)arg;
9      long id = p->thread_id;
10     int i, np = p->chunk;  float y = 0.0;
11     for(i=id*np; i <(id+1)*np; i++)
12         y+=x[i];
13
14     p->data = y;
15     // this is how the pointer to a structure is feed.
16     return (void *)p;
17 }
18 /*-----*/
19 int main(int argc, char *argv[]){
20     pthread_t *t;  thread_data *q;
21     np = (int) (num_points/num_threads);
22     /*---- Create threads ----*/
23     for(i=0; i < num_threads; i++){
24         q[i].thread_id = i; q[i].chunk = np;
25         pthread_create(&t[i], NULL, tfunc, (void *) (q+i));
26     }
27     /*-----Join threads -----*/
28     for(sum=0.0, i=0; i < num_threads; i++){
29         pthread_join(t[i], (void **)&q[i] );
30         sum+=q[i].data;
31     }
32 }
```

Program 6 : mutex1_pthreads.c

```
1
2 pthread_mutex_t mutex1 = PTHREAD_MUTEX_INITIALIZER;
3 int counter = 0
4
5 void *functionC(){
6     pthread_mutex_lock(&mutex1 );
7     counter++;
8     printf("Counter value: %d\n",counter);
9     pthread_mutex_unlock( &mutex1 );
10 }
11
12 int main(int argc, char *argv[]){
13
14
15     for(l=0; l < num_threads; l++)
16         pthread_create(&thread[l],NULL,&functionC,NULL);
17
18     for(l=0; l < num_threads; l++)
19         pthread_join(thread[l],NULL);
20
21 }
```

All programs

- 1 hello_world1_threads.c
- 2 hello_world2_threads.c
- 3 return_threads.c
- 4 summation1_threads.c
- 5 summation2_threads.c
- 6 mutex1_threads.c
- 7 mutex2_threads.c
- 8 create_fork.c
- 9 create_threads.c
- 10 stack_threads.c
- 11 scalar_prod_threads.c
- 12 compute_pi_threads.c
- 13 nbody_kernel_threads.c

Exercise

- 1 Compute the value of π by numerically integrating $1/(1 + x^2)$ between limits $[0 - 1]$ in parallel using `pthread` . You need to split the limit of integration and send the sub-limits to threads which return the answer for that limit.
- 2 Write a parallel program using `pthread` for finding the number of prime number up to some number n and compare the performance with the `OpenMP` program for the same (you can modify the `OpenMP` program which was provided in `OpenMP` session).
- 3 Write a parallel program for matrix multiplication using `pthread` .

Thank You !