

COMP4300 - Course Update

- COMP4300 is a fast-paced 4000-level course: **It assumes you are mature and independent programmers.**
- The course introduces the basics of the semantics of the programming models (Pthreads, OpenMP, CUDA).
- **You are left with the task and the responsibility of their further exploration and practice to master these programming models.**
- This is particularly important for the second half of the course.

SHARED MEMORY PARALLEL COMPUTING

USING MULTIPLE CORES



1



2

TEQSA PROVIDER ID: PRV1202 (AUSTRALIAN UNIVERSITY) CRICOS PROVIDER CODE: 00120C



References

- Chapter 12 from Computer Systems A Programmer's Perspective, Third Edition, Randal E. Bryant and David R. O'Hallaron, Pearson Education Heg USA, ISBN 9781292101767.
- Programming with POSIX Threads, David R. Butenhof, Addison-Wesley Professional, ISBN-13 : 978-0201633924.

How can we avoid the flaws of process-based concurrency?

3



4

TEQSA PROVIDER ID: PRV1202 (AUSTRALIAN UNIVERSITY) CRICOS PROVIDER CODE: 00120C



Parallel Programming with Threads

A **thread** is a logical flow that runs within the context of a process.

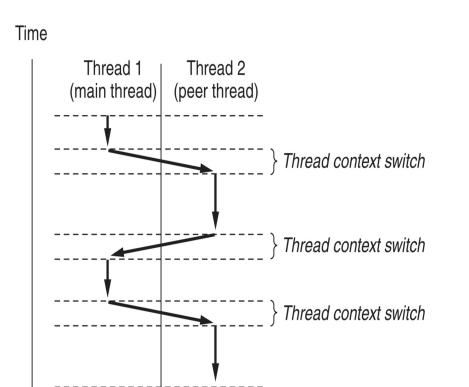
- So far we have discussed programs that consisted of a single thread per process.
- Multiple “independent” threads can be added to an existing process rather than starting a new process.
- Threads are scheduled by the OS and run as independent entities largely because they duplicate only the bare essential resources that enable them to exist as executable code.

Understanding Pthreads concepts helps developers grasp fundamental parallel programming principles

5

Thread Execution Model

- Each process is created with a single thread called the main thread.
- The main thread can create a peer thread, and the two run concurrently.
- If they are scheduled on the same CPU/core, the OS at some point will pass control from the main to the peer thread via a context switch.
- If they are scheduled by the OS on different cores they will run in parallel.
- Threads associated with a process form a pool of peers, where each one can kill or wait for any other to terminate.
- Each peer can read and write the same shared data.

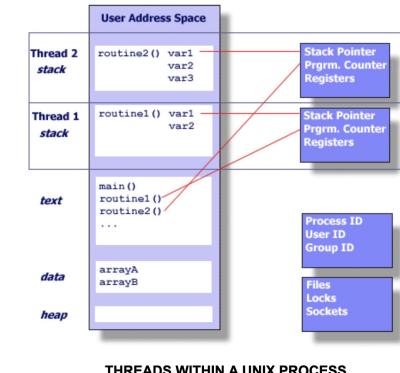


7

Parallel Programming with Threads

- A thread is a logical flow that runs within the context of a process.
- So far we discussed programs that consisted of a single thread per process.
- Multiple “independent” threads can be added to an existing process rather than starting a new process.
- Threads are able to be scheduled by the OS and run as independent entities largely because they duplicate only the bare essential resources that enable them to exist as executable code.

6



TEGDA PROVIDER ID: PRV12302 (AUSTRALIAN UNIVERSITY) CRICOS PROVIDER CODE: 00120C



POSIX Threads (Pthreads)

- Historically, hardware vendors have implemented their own proprietary versions of threads, making it difficult for programmers to develop portable threaded applications.
- POSIX Threads provide a standard interface, specified by the IEEE 1003.1 (1995) standard, for manipulating threads from C programs.
- Pthreads are defined as a set of C language programming types and procedure calls, implemented with an #include <pthread.h> header file and a thread library (this may be part of another library, such as libc, in some implementations).
- The specification contains about 60 functions that allow programs to create, kill, and reap threads, to share data safely with peer threads, and to notify peers about changes in the system state.

Compiler / Platform	Compiler Command	Description
INTEL Linux	icc -lpthread	C
	icpc -lpthread	C++
PGI Linux	pgcc -lpthread	C
	pgc++ -lpthread	C++
GNU Linux, Blue Gene	gcc -lpthread	GNU C
	g++ -lpthread	GNU C++

8

TEGDA PROVIDER ID: PRV12302 (AUSTRALIAN UNIVERSITY) CRICOS PROVIDER CODE: 00120C



POSIX Threads (Pthreads)

Detaching threads

- A thread is either *joinable* or *detachable*.
- *Joinable*: its memory resources (such as the stack) are not freed until it is reaped by another thread.
- A detached thread cannot be reaped by other threads. Its memory resources are freed automatically by the system when it terminates.
- To avoid memory leaks, each joinable thread should either be explicitly reaped by another thread, or detached by a call to `pthread_detach`

```
#include <pthread.h>
int pthread_detach(pthread_t tid);
Returns: 0 if OK, nonzero on error
```

Threads can detach themselves by calling `pthread_detach` with an argument of `pthread_self`.

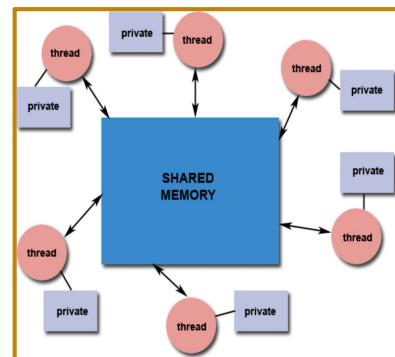
13

TEGDA PROVIDER ID: PRV12302 (AUSTRALIAN UNIVERSITY) CRICOS PROVIDER CODE: 00120C



Pthreads Memory Model

- **Thread context:** TID, stack, stack pointer (SP), program counter (PC), condition codes (CC), registers values.
- **Shared:** process virtual address space → code, read/write data, heap, shared libraries, open files.

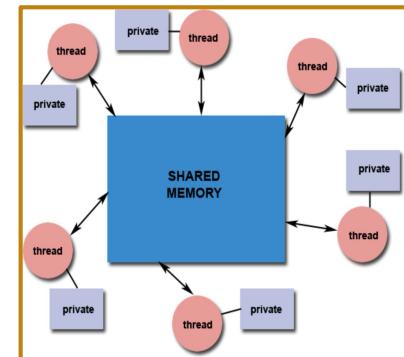


15

TEGDA PROVIDER ID: PRV12302 (AUSTRALIAN UNIVERSITY) CRICOS PROVIDER CODE: 00120C

Pthreads Memory Model

- The key advantage of threads is the efficiency and ease with which they can share data
- The key disadvantage is the potential for complicated bugs that do not appear in serial code.
- To write correct threaded programs a clear understanding of the sharing mechanisms and risks is required



14

TEGDA PROVIDER ID: PRV12302 (AUSTRALIAN UNIVERSITY) CRICOS PROVIDER CODE: 00120C



Pthreads Memory Model

- Registers, Condition Codes (CC), Program Counter (PC), the Stack Pointer (SP) are private. Thread stack is usually accessed independently by each thread. However, thread stacks are not protected!
- If a thread accesses a stack pointer to another thread's stack, it can read and write to it!

```
1 #include "csapp.h"
2 #define N 2
3 void *thread(void *vargp);
4
5 char **ptr; /* Global variable */
6
7 int main()
8 {
9     int i;
10    pthread_t tid;
11    char *msg[N] = {
12        "Hello from foo",
13        "Hello from bar"
14    };
15
16    ptr = msg;
17    for (i = 0; i < N; i++)
18        Pthread_create(&tid, NULL, thread, (void *)i);
19        Pthread_exit(NULL);
20    }
21
22 void *thread(void *vargp)
23 {
24     int myid = (int)vargp;
25     static int cnt = 0;
26     printf("%d: %s (cnt=%d)\n", myid, ptr[myid], ++cnt);
27     return NULL;
28 }
```

16

TEGDA PROVIDER ID: PRV12302 (AUSTRALIAN UNIVERSITY) CRICOS PROVIDER CODE: 00120C



Mapping Variables to Memory

- **global variables**: declared outside functions, only one copy in virtual memory area, readable/writable by any thread.
- **local automatic variables**: declared inside functions, each thread's stack contains its own instance of it. E.g. tid and myid.
- **local static variables**: declared inside functions with the static attribute. There is only one copy in virtual memory area, readable/writable by any thread.
- **shared variables**: a variable is shared only, and only if, one of its instances is referenced by more than one thread. E.g. cnt.

```
1 #include "csapp.h"
2 #define N 2
3 void *thread(void *vargp);
4
5 char **ptr; /* Global variable */
6
7 int main()
8 {
9     int i;
10    pthread_t tid;
11    char *msgs[N] = {
12        "Hello from foo",
13        "Hello from bar"
14    };
15
16    ptr = msgs;
17    for (i = 0; i < N; i++)
18        Pthread_create(&tid, NULL, thread, (void *)i);
19    Pthread_exit(NULL);
20 }
21
22 void *thread(void *vargp)
23 {
24     int myid = (int)vargp;
25     static int cnt = 0;
26     printf("[%d]: %s (cnt=%d)\n", myid, ptr[myid], ++cnt);
27     return NULL;
28 }
```

17

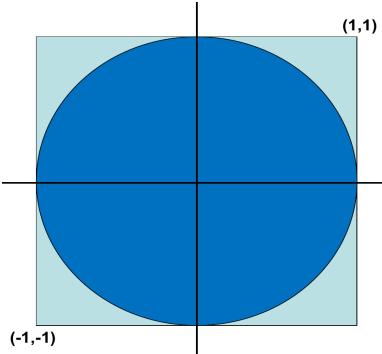
TEGDA PROVIDER ID: PRV1202 (AUSTRALIAN UNIVERSITY) CRICOS PROVIDER CODE: 00120C



Pthreads

Example: Computing Pi

- The ratio of area of circle $radius=1$ to the square is $\pi/4$
- Generate random numbers for x and y within the domain of square ie range [-1,1] for each axis
- identify those that are distance less than 1 from origin
- The ratio of points in circle to the total points is $\pi/4$



18

TEGDA PROVIDER ID: PRV1202 (AUSTRALIAN UNIVERSITY) CRICOS PROVIDER CODE: 00120C



Pthreads

Example: Computing Pi

- The ratio of area of circle $radius=1$ to the square is $\pi/4$
- Generate random numbers for x and y within the domain of square ie range [-1,1] for each axis
- identify those that are distance less than 1 from origin
- The ratio of points in circle to the total points is $\pi/4$

```
#include <stdio.h>
#include <stdlib.h>
#include <math.h>
#include <pthread.h>

#define NUM_THREADS 4
#define NUM_THROWNS 1000000

double total_points_inside_circle = 0;

int main() {
    pthread_t threads[NUM_THREADS];
    int thread_ids[NUM_THREADS];

    srand(time(NULL));

    for (int i = 0; i < NUM_THREADS; i++) {
        thread_ids[i] = i;
        pthread_create(&threads[i], NULL, computePi, (void *)&thread_ids[i]);
    }

    for (int i = 0; i < NUM_THREADS; i++) {
        pthread_join(threads[i], NULL);
    }

    double pi_estimate = 4.0 * total_points_inside_circle / NUM_THROWNS;
    printf("Estimated value of pi: %f\n", pi_estimate);

    return 0;
}
```

19

TEGDA PROVIDER ID: PRV1202 (AUSTRALIAN UNIVERSITY) CRICOS PROVIDER CODE: 00120C



Pthreads

Example: Computing Pi

- The ratio of area of circle $radius=1$ to the square is $\pi/4$
- Generate random numbers for x and y within the domain of square ie range [-1,1] for each axis
- identify those that are distance less than 1 from origin
- The ratio of points in circle to the total points is $\pi/4$

```
double total_points_inside_circle = 0;

void* computePi(void* arg) {
    int myNum = *(int*)arg;
    double x, y;

    for (int i = myNum; i < NUM_THROWNS; i += NUM_THREADS) {
        x = (double)rand() / RAND_MAX * 2.0 - 1.0;
        y = (double)rand() / RAND_MAX * 2.0 - 1.0;

        if (x * x + y * y <= 1.0) {
            total_points_inside_circle++;
        }
    }

    pthread_exit(NULL);
}
```

20

TEGDA PROVIDER ID: PRV1202 (AUSTRALIAN UNIVERSITY) CRICOS PROVIDER CODE: 00120C



DEMO

Pthreads based Concurrency

- **Problem Definition**
 - Computing π
 - The exact result is 3.1415..., which makes it easy to verify our numerical approximation
 - We use Mont Carlo integration
- **Sequential Algorithm**
 - Computes the integration over N samples
 - Evaluates the function at each point
 - Check if the point is inside the unit circle
- **Parallelization Strategy**
 - Each process handles its own contiguous portion of the N samples
 - Added 64 bytes of padding to the `thread_data` structure to avoid false sharing (where threads invalidate each other's cache lines).
 - Each thread can maintain its own counter in an array element – no Mutex
 - The parent process sums the results from each thread
- **Performance Measurement**
 - Uses `clock()` to measure execution time for both sequential and parallel algorithms – need to allow for measurement of time on all threads
 - Calculates and reports speedup

21

TEQSA PROVIDER ID: PRV1202 (AUSTRALIAN UNIVERSITY) CRICOS PROVIDER CODE: 00323C



DEMO

Pthreads based Concurrency

- **Problem Definition**
 - Computing π
 - The exact result is 3.1415..., which makes it easy to verify our numerical approximation
 - We use Mont Carlo integration
- **Sequential Algorithm**
 - Computes the integration over N samples
 - Evaluates the function at each point
 - Check if the point is inside the unit circle
- **Parallelization Strategy**
 - Each process handles its own portion of the N samples
 - Local results are combined using pipes
 - The parent process collects the final result from the child processes
- **Performance Measurement**
 - Uses `clock()` to measure execution time for both sequential and parallel algorithms – need to allow for measurement of time on all threads
 - Calculates and reports speedup

22

TEQSA PROVIDER ID: PRV1202 (AUSTRALIAN UNIVERSITY) CRICOS PROVIDER CODE: 00323C

