

Engenharia de Sistemas da Computação

João Barbosa and Vitor Oliveira

Universidade do Minho

2021

Introduction - Syllabus - Grades

Syllabus

- C++11/C++17 POSIX Threads Programming Model && Mutual Exclusion
- Basics of Concurrency in C++ 11/17/20
- SIMD/SPMD/STMD Programming Models
- Distributed Memory Model Parallel Programming
- HPC Clusters Architecture
- HPC Program performance analysis and debugging

Grading

- Two projects
 - Basic raytracer with acceleration structure - April 6th (35%)
 - Performance Analysis - May 18th (35%)
 - Maximum of two students per group
 - No deadline extensions
 - Report in a "Paper" Format
- Exam at the end of the semester (30%)

C++11/17 Thread Programming Model

POSIX Execution Threads

- OS Dependent Execution Model (UNIX only)
- Allows multiple time overlapping control flows in shared memory (thread)
- Standard POSIX.1c, Threads extensions (IEEE Std 1003.1c-1995)
- Defines API for:
 - Thread creation, join, etc...
 - Mutual exclusion (MUTEX)
 - Conditional Variables
 - Synchronization

C++ 11/17 Threads

- OS Independent Execution Model (Resource Acquisition Is Initialization)
 - UNIX - Pthreads
 - Windows - Windows Threads
- Allows multiple time overlapping control flows in shared memory (thread)
- Defines API for:
 - Thread creation, join, etc...
 - Mutual exclusion (MUTEX)
 - Conditional Variables
 - Synchronization
- Introduces parallelism primitives for C++ and allows compiler to “reason” about optimizations

C++ 11/17 Threads

Why not OpenMP?

■ Pros:

- Easy to implement on pragmas
- Thread Pool matches the number of cores
- Dynamic Scheduling
- Implicit shared variables and explicit privates
- Good Support in Unix systems

■ Cons:

- Not everything fits in a for loop fork and join
- Good graph algorithms are "hard" to implement
- Irregular algorithms are hard to map and to choose a good scheduling strategy
- Ultimately maps to POSIX Threads
- *"Easy to induce a false sense of performance"*

C++ 11/17 Threads - Example

```
A = INT[N]
SUM = 0
(...)
function LocalSum(START_BLOCK, END_BLOCK,
                  A, SUM)

    for I in START_BLOCK .. END_BLOCK
        LOCALSUM = LOCALSUM + A[I]

SUM += LOCALSUM
```

C++ 11/17 Threads - OpenMP Example

```
A = INT[N]
SUM = 0
(...)
function LocalSum(START_BLOCK, END_BLOCK,
                  A, SUM)
#pragma omp parallel for private(I)
    reduce(localsum:+)
    schedule(dynamic, CHUNK)
for I in START_BLOCK .. END_BLOCK
    LOCALSUM += A[I]
SUM = LOCALSUM
schedule(dynamic, CHUNK)
```

C++ 11/17 Threads - Example

```
A = INT[N]
SUM = 0
(...)
function LocalSum(START_BLOCK, END_BLOCK,
                  A, SUM)

    for I in START_BLOCK .. END_BLOCK
        LOCALSUM = LOCALSUM + A[I]
        Acquire (Mutex)
        SUM += LOCALSUM
        Release (Mutex)
```

C++ 11/17 Threads - Example

```
1  #include <thread>
2  #include <mutex>
3  #include <vector>
4  #define N (1u << 16)
5
6  int main(int argc, char* argv) {
7      float A(N);
8      float sum = 0;
9      std::vector<std::thread> threads;
10     (...)
11 }
```

C++ 11/17 Threads - Example

```
1   for(auto start=0u; start<N; start+=512u){
2       uint end = std::min(start+512,N);
3       threads.push_back(std::thread(
4           (&)(const float A(),
5               const uint start,
6               const uint end) {
7
8           for(auto i = start; i < end; i++) {
9               A(i) = float(i) / float(N);
10          }
11
12          },A,start ,end));
13    }
14
15    for(auto &thr : threads) {
16        thr.join();
17    }
```

C++ 11/17 Threads - Example

```
1  std::mutex protect_sum;
2  for(auto start=0u; start<N; start+=512u){
3      uint end = std::min(start+512,N);
4      threads.push_back(std::thread(
5          (&)(const float A(),
6              const uint start,
7              const uint end) {
8          float localsum;
9          for(auto i = start; i < end; i++) {
10             localsum+=A(i);
11         }
12         std::lock_guard lk(protect_sum);
13         sum += localsum;
14         },A,start,end));
15     )
16 }
17 for(auto &thr : threads) {
18     thr.join();
19 }
20 (...)
```

C++ 11/17 Threads - Example

- Access search6
- module use `/home/jbarbosa/software/modulefiles`
- module load cmake gcc

Try to implement the fastest parallel sum using C++11/17
Threads (30min)

C++ 11/17 Threads - Example

Problems:

- Too many threads $2^{16}/2^7 = 2^9$
 - Creating and Destroying threads is expensive
 - Typically we want one thread per core
- Too many locks $2^{16}/2^7 = 2^9$
 - Locking is expensive

Possible solutions:

- Atomics
- Task based programming

C++ 11/17 Threads - Exercise

- Implement a prefix sum (reduction)
- The problem that we discussed using the locks

C++ 11/17 ASYNC - Example

```
1  std::vector<std::future<float>> tasks;
2  for(auto start=0u; start<N; start+=512u){
3      uint end = std::min(start+512,N);
4      threads.push_back(std::async(
5          (&)(const float A(),
6              const uint start,
7              const uint end) -> float {
8          float localsum;
9          for(auto i = start; i < end; i++) {
10             localsum+=A(i);
11         }
12         return localsum;
13     },A.start ,end));
14     )
15 }
16 for(auto &thr : threads) {
17     sum += thr.get();
18 }
```

C++ 11/17 ASYNC

- Creates a thread pool that equals the number of cores
- Issues a *task* using a FIFO scheduling approach
- Can get a result from a *task* through a ***std::future***

C++ 11/17 ASYNC - Exercise

- Implement a prefix sum (reduction)
- The problem that we discussed using the futures

C++ 11/17 ASYNC - Advantages

- Allows task based parellism
- e.g. Traverse a graph to compute a single shortest path
- Not everything can be implemented using for loops (OpenMP)

Single Source Shortest Path Problem

