

Parallel Programming course. C++ threading

Obolenskiy Arseniy, Nesterov Alexander

Nizhny Novgorod State University

April 29, 2025

- 1 Introduction to C++ threading API
- 2 C++ STL threading API (`std::thread`, `std::jthread`)
- 3 `std::future`, `std::promise` and `std::async`
- 4 Synchronization primitives (mutexes, condition variables, ...)
- 5 Best practices and recommendations

Introduction to C++ threading API

What is `std::thread`?

- Part of the C++11 thread support library (`<thread>`, `<mutex>`, etc.)
- Low-level, manual thread creation and management
- Relies on the OS native threads under the hood
- Provides:
 - `std::thread` for launching threads
 - Synchronization primitives (`std::mutex`, `std::condition_variable`, `std::atomic`)
 - Utilities for futures and promises (`std::future`, `std::promise`)

Threading before C++11

- POSIX threads (`pthread`): C-based API, widely used on Unix-like systems
- Windows threads: Win32 API with `CreateThread`, `CRITICAL_SECTION`, and events
- External libraries (e.g. `Boost.Thread`)

Challenges and inconveniences:

- non-standard APIs (towards C++)
- platform dependent APIs
- verbose initialization boilerplate
- manual resource management

- C++11 (2011): Introduced `std::thread`, `std::mutex`, `std::future`, etc.
- C++14/17/20: Incremental improvements (e.g., `std::hardware_constructive_interference_size`)
- Widely adopted as the base for custom thread pools and concurrency utilities
- Now the foundation of many higher-level C++ concurrency libraries

OpenMP vs TBB vs `std::thread`

- **OpenMP**

- Directive-based, compiler-driven shared-memory parallelism
- Simple loop and region parallelism

- **TBB**

- Template library, task-based parallelism with work-stealing
- High-level parallel patterns (`tbb::parallel_for`, `tbb::parallel_reduce`, etc.)

- **`std::thread`**

- Low-level manual thread API
- Full control over thread lifetime, but more boilerplate
- Foundation for building custom task systems or thread pools

Pros and cons of `std::thread`

Pros

- Complete control over thread creation and destruction
- No hidden scheduler—behavior is predictable
- Part of standard C++, portable across platforms
- Good for learning fundamentals

Cons

- Manual management—risk of leaks, detach/join errors
- Verbose boilerplate for synchronization
- No built-in task scheduling or work-stealing
- Harder to scale and tune compared to higher-level APIs

C++ STL threading API

(`std::thread`, `std::jthread`)

- Constructors:

- `std::thread(callable, args...)`: starts a new thread executing `callable` with `args`
- Default constructor: creates a thread object without an associated thread

- Member functions:

- `join()`: blocks until the thread finishes execution
- `detach()`: detaches the thread to run independently
- `joinable()`: checks whether the thread can be joined
- `get_id()`: returns the `std::thread::id` of the thread

- Properties:

- Move-only: supports move construction and assignment; copy operations are deleted
- Static `hardware_concurrency()`: returns the number of concurrent threads supported

Creating and launching threads

```
1  #include <thread>
2  #include <iostream>
3
4  void worker(int id) {
5      std::cout << "Worker " << id << " running\n";
6  }
7
8  int main() {
9      // Launch two threads with argument
10     std::thread t1(worker, 1);
11     std::thread t2(worker, 2);
12
13     // Wait for them to finish
14     if (t1.joinable()) {
15         t1.join();
16     }
17     if (t2.joinable()) {
18         t2.join();
19     }
20
21     return 0;
22 }
23
```

- Construct `std::thread` with a callable + optional args
- Must call `join()` or `detach()` before destruction
- Threads are move-only: no copy construction/assignment

Managing thread lifetime

- `join()`: waits for thread completion
- `detach()`: allows thread to run independently (daemon-like)

```
1 std::thread t(worker);  
2 // ...  
3 if (t.joinable()) {  
4     t.join();  
5 }  
6
```

- `joinable()`: check if thread can be joined
- Detached threads continue after `main`—dangerous if resources go out of scope
- Always ensure each thread is either joined or detached

Passing arguments to threads

```
1 #include <thread>
2 #include <string>
3 #include <iostream>
4
5 void greet(const std::string &name) {
6     std::cout << "Hello, " << name << "!\n";
7 }
8
9 int main() {
10     std::string user = "Alice";
11     // Pass by reference: use std::ref
12     std::thread t(greet, std::ref(user));
13     t.join();
14     return 0;
15 }
16
```

- By default, args are copied into the new thread
- Use `std::ref()` to pass references

Querying hardware concurrency capabilities

- `std::thread::hardware_concurrency()` returns the number of supported threads
- May potentially return 0

`unsigned n = std::thread::hardware_concurrency();` Usage example: to size thread pools or partition work

std::jthread

- Introduced in C++20 as a safer and more convenient alternative to std::thread.
- Automatically joins upon destruction, reducing the risk of detached threads.
- Supports cooperative cancellation via std::stop_token.

Example usage:

```
1  #include <iostream>
2  #include <thread>
3  #include <chrono>
4
5  void task(std::stop_token token) {
6      while (!token.stop_requested()) {
7          std::cout << "Working...\n";
8          std::this_thread::sleep_for(std::chrono::milliseconds(100));
9      }
10     std::cout << "Task stopping.\n";
11 }
12
13 int main() {
14     std::jthread jt(task);
15     std::this_thread::sleep_for(std::chrono::seconds(1));
16     // Request stop automatically during destruction or manually via jt.request_stop
17     // ();
18     return 0;
19 }
```

`std::future`, `std::promise` and
`std::async`

Future/Promise and Async Execution

- Future and promise provide a mechanism for asynchronous communication between threads
- A promise allows one thread to set a value or report an error
- The associated future retrieves the value, waiting if necessary
- `std::async` simplifies launching asynchronous tasks without explicit thread management
- These features promote decoupling of computation and enhance concurrency control

std::future and std::promise

```
1  #include <future>
2  #include <iostream>
3
4  int compute() {
5      int result;
6      ... // heavy computation
7      return result;
8  }
9
10 int main() {
11     std::promise<int> prom;
12     std::future<int> fut = prom.get_future();
13
14     std::thread t([&prom]{
15         prom.set_value(compute());
16     });
17
18     std::cout << "Result: " << fut.get() << "\n";
19     t.join();
20     return 0;
21 }
22
```

std::promise sets a value (or failure)

std::future retrieves it on demand

Using `std::async`

- Launches tasks asynchronously and returns a `std::future`
- Can run immediately in a new thread or be deferred until needed
- Simplifies asynchronous programming by handling thread management

```
1 #include <future>
2 #include <iostream>
3
4 int computeSquare(int x) {
5     // Simulate heavy computation
6     std::this_thread::sleep_for(std::chrono::seconds(1));
7     return x * x;
8 }
9
10 int main() {
11     // Launch computeSquare asynchronously. Using std::launch::async forces immediate
        execution
12     std::future<int> fut = std::async(std::launch::async, computeSquare, 5);
13     std::cout << "Square: " << fut.get() << std::endl;
14     return 0;
15 }
16
```

Different launch policies available in `std::async`

```
1 #include ...
2
3 int computeCube(int x) { ... }
4
5 int main() {
6     // Using deferred launch: execution is postponed until get() is called
7     std::future<int> futDeferred = std::async(std::launch::deferred, computeCube, 3);
8     std::cout << "Deferred result: " << futDeferred.get() << std::endl;
9
10    // Using async launch: task is executed immediately in a new thread
11    std::future<int> futAsync = std::async(std::launch::async, computeCube, 3);
12    std::cout << "Async result: " << futAsync.get() << std::endl;
13
14    // Using default launch policy: implementation defined behavior
15    std::future<int> futDefault = std::async(computeCube, 3);
16    std::cout << "Default launch result: " << futDefault.get() << std::endl;
17
18    return 0;
19 }
20
```

Policies:

- `std::launch::deferred`: Execution is delayed until `get()` is called
- `std::launch::async`: Execution starts immediately in a separate thread
- Default policy: The decision is left to the implementation

Synchronization primitives (mutexes, condition variables, ...)

Synchronization primitives overview

- C++ provides several mechanisms to coordinate concurrent operations:
 - Mutual exclusion: `std::mutex`, `std::lock_guard`, `std::unique_lock`
 - Condition variables: `std::condition_variable`
 - Atomic operations: `std::atomic`
- Choose the appropriate primitive based on the required control and performance
- Proper synchronization is key to ensuring thread safety and avoiding data races

std::mutex

```
1  #include <mutex>
2  #include <thread>
3  #include <iostream>
4
5  std::mutex mtx;
6  int counter = 0;
7
8  void increment() {
9      mtx.lock();
10     {
11         ++counter; // protected section
12     }
13     mtx.unlock();
14 }
15
16 int main() {
17     std::thread t1(increment);
18     std::thread t2(increment);
19
20     t1.join();
21     t2.join();
22
23     std::cout << "Final counter value: " << counter << std::endl;
24     return 0;
25 }
26
```

- Manual locking with `mtx.lock()` and unlocking with `mtx.unlock()`
- Be cautious with exceptions to avoid deadlocks

There is a way to ensure that mutex will be unlocked...

std::mutex and std::lock_guard

```
1  #include <mutex>
2
3  std::mutex mtx;
4  int counter = 0;
5
6  void increment() {
7      std::lock_guard<std::mutex> lock(mtx);
8      ++counter; // protected section
9  }
10
```

- std::mutex: exclusive lock
- std::lock_guard: RAII wrapper—locks on construction, unlocks on destruction

std::unique_lock and std::condition_variable

```
1  #include <mutex>
2  #include <condition_variable>
3
4  std::mutex mtx;
5  std::condition_variable cv;
6  bool ready = false;
7
8  void worker() {
9      std::unique_lock<std::mutex> lock(mtx);
10     cv.wait(lock, []{ return ready; });
11     // proceed once ready == true
12 }
13
14 void notifier() {
15     {
16         std::lock_guard<std::mutex> lock(mtx);
17         ready = true;
18     }
19     cv.notify_one();
20 }
21
```

`std::unique_lock`: provides flexible locking mechanisms such as deferred locking, timed locking, and manual unlocking/relocking, unlike `lock_guard` which locks immediately and strictly scopes the lock.

`std::condition_variable`: allows one or more threads to wait until a particular condition is met. It's typically used with `unique_lock` to enable the thread to wait and then be notified.

Lock-free primitives for simple data types. Avoid mutex overhead for simple counters and flags Supported operations list:

- store: assign a new value
- load: retrieve the current value
- exchange: replace the value and obtain the old value
- compare_exchange_weak/compare_exchange_strong: atomically compare and set
- fetch_add: add to the value and return the previous value
- fetch_sub: subtract from the value and return the previous value
- fetch_and: perform bitwise AND and return the previous value
- fetch_or: perform bitwise OR and return the previous value
- fetch_xor: perform bitwise XOR and return the previous value

Atomic operations usage example

```
1  #include <atomic>
2  #include <thread>
3  #include <iostream>
4
5  std::atomic<int> counter(0);
6
7  void increment() {
8      for (int i = 0; i < 100000; ++i) {
9          counter.fetch_add(1, std::memory_order_relaxed);
10     }
11 }
12
13 int main() {
14     std::thread t1(increment);
15     std::thread t2(increment);
16
17     t1.join();
18     t2.join();
19
20     std::cout << "Final counter value: " << counter.load() << std::endl;
21     return 0;
22 }
23
```

`std::atomic` provides lock-free operations. Use `fetch_add` method and `load` to operate on atomic variables safely

Best practices and recommendations

Best practices and recommendations

- Always join or detach threads before destruction
- Prefer RAI wrappers (`std::lock_guard`, `std::unique_lock`)
- Minimize shared state; prefer message passing or futures
- Be cautious with detached threads, manage lifetimes carefully
- Consider higher-level thread pools (e.g., `std::async`)

When and where to use `std::thread` in real world scenarios

- Fine-grained control over threading
- Building custom schedulers or thread pools
- Interfacing directly with OS thread APIs
- Performance critical sections where you avoid scheduler overhead
- When introducing 3rdparty library is an overkill

Thank You!

- cppreference.com: <https://en.cppreference.com/w/cpp/thread>