

An Adventure in Race Conditions

Prepared for ACCU 2019

©2019

Felix Petriconi
felix@petriconi.net

2019-04-10

- ▶ Started with C++ 1994
- ▶ Programmer and development manager since 2003 at MeVis Medical Solutions AG, Bremen, Germany
 - ▶ Development of medical devices in the area of mammography and breast cancer therapy (C++, Ruby)
- ▶ Programming activities:
 - ▶ Blog editor of ISO C++ website
 - ▶ Active member of C++ User Group Bremen
 - ▶ Contributor to slab's concurrency library
 - ▶ Member of ACCU conference committee
- ▶ Married with Nicole, having three children, living near Bremen, Germany
- ▶ Other interests: Classic film scores, composition

Being wrong isn't a bad thing like they teach you in school.
It is an opportunity to learn something.

Richard Feynman

Why I am here?

- ▶ I like being a programmer
- ▶ I like sharing my experience
- ▶ I like to learn from you

The Adventure

Felix Petriconi
felix@petriconi.net

Why I am here?

Why are you here?

Motivation

Pr my domain



Why I am here?

Why are you here?

Me

Probably from my domain

Why are you here?

Display of radiological images for breast cancer detection and diagnosis

Why I am here?

Why are you here?


Motivation

Problem from my domain



Problem from my domain


- ▶ 3D Mammography 16bit grayscale images of 2048*2560, 50-90 slices
- ▶ Display up to 30fps cine mode on 5MP displays
- ▶ Only lossless compression is allowed
- ▶ JPEG 2000 decompression is too slow while decompression for display
- ▶ Re-compression into proprietary format
- ▶ Initial approach used 2 user threads
- ▶ 620'000 slices / day $\equiv \sim 9h$ processing time



¹Mammography image from <http://www.dclunie.com/>

Problem became recently more challenging

- ▶ 3D Mammography 16bit grayscale images of 3328*4096, 50-90 slices
- ▶ Users expected same performance for display
- ▶ Some improvements were necessary...



Parallel Data Access - Basics

The Adventure

Felix Petriconi
felix@petriconi.net

Thread Basics

Start


Let's Improve

1st Correction

Parallel Data Read-Only Access

Thread Basics


Start
Let's Improve
1st Correction



Parallel Data Read-Only Access

Thread Basics

Start
Let's Improve
1st Correction



Thread Basics

Start

Let's Improve

1st Correction

- ▶ Atomic
- ▶ Mutex
- ▶ Semaphore
- ▶ Memory Fence
- ▶ Transactional Memory

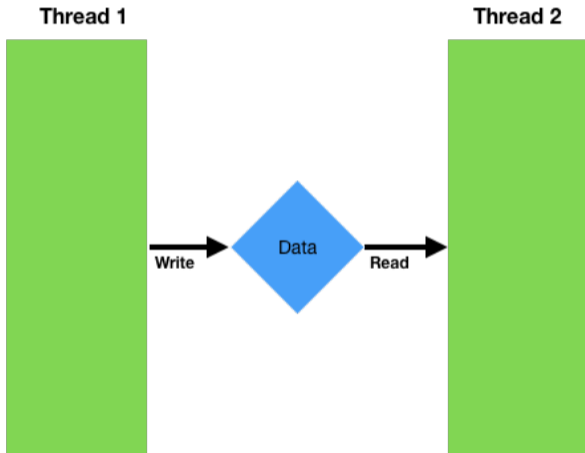
Parallel Data Read/Write Access

The Adventure

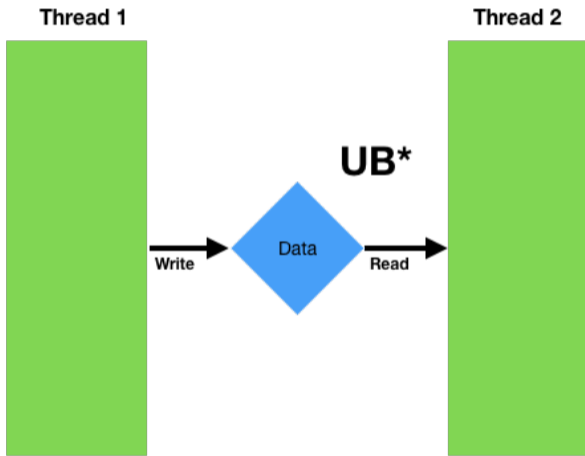
Felix Petriconi
felix@petriconi.net

Thread Basics

Start
Let's Improve
1st Correction



Parallel Data Read/Write Access



* Undefined Behavior

Code Example - Undefined Behaviour I

```
1 #include <iostream>
2 #include <thread>
3
4 using namespace std;
5
6 int main() {
7     int value = 42;
8
9     auto t1 = thread{ [&value]{ ++value; } };
10    auto t2 = thread{ [&value]{ value *= 2; } };
11
12    t1.join();
13    t2.join();
14
15    cout << value << endl;
16 }
```

Output

Possible results on my machine: 43, 84, 85, 86


View on a CPU

Thread Basics

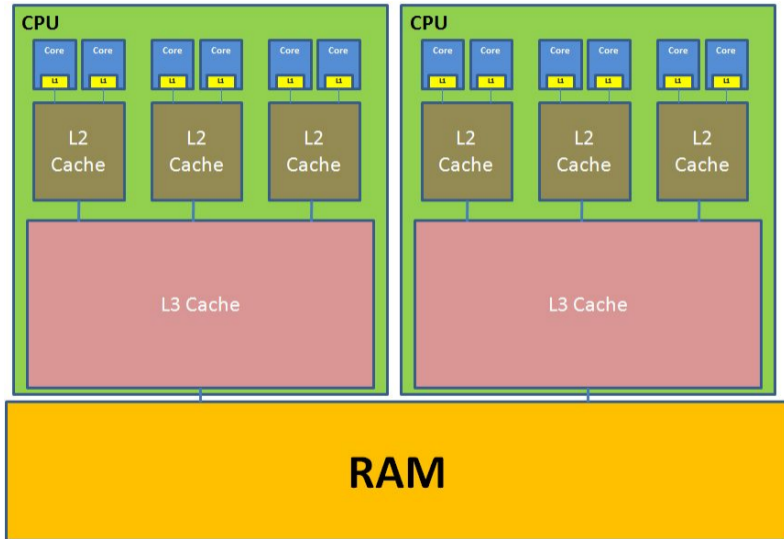
Start

Let's Improve

1st Correction



More Accurate View of a CPU



Code Example - Using atomic

The Adventure

Felix Petriconi
felix@petriconi.net

Thread Basics

Start

Let's Improve

1st Correction

```
1 #include <atomic>
2 #include <iostream>
3 #include <thread>
4
5 using namespace std;
6
7 int main() {
8     atomic_int value{42};
9
10    auto t1 = thread{ [&value]{ ++value; } };
11    auto t2 = thread{ [&value]{ value = value * 2; } };
12
13    t1.join();
14    t2.join();
15
16    cout << value << endl;
17 }
```

Output

Possible results on my machine: 84, 85, 86

Code Example - Using mutex

Thread Basics

Start

Let's Improve

1st Correction

```
1
2 int main() {
3     int value{42};
4     mutex m;
5     auto t1 = thread{ [&]{
6         unique_lock block{m};
7         ++value;
8     } };
9     auto t2 = thread{ [&]{
10        unique_lock block{m};
11        value *= 2;
12    } };
13
14    t1.join(); t2.join();
15
16    cout << value << endl;
17 }
```

Output

Possible results: 85, 86

Beware of your compiler - Undefined Behaviour II

Thread Basics

Start

Let's Improve

1st Correction

```
1 #include <thread>
2 using namespace std;
3
4 int x = 0, y = 1;
5 int test = x;
6
7 void setTest() {
8     test = y;
9 }
10
11 int main() {
12     thread run(setTest);
13
14     while (test == 0) { }
15
16     run.join();
17     return 0;
18 }
```

<https://godbolt.org/z/g7ZEXL>

Beware of your compiler - Undefined Behaviour II

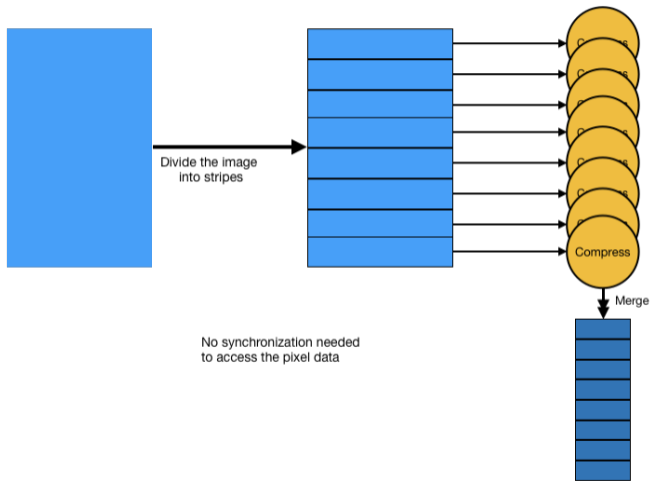
```
1 #include <thread>
2 using namespace std;
3
4 int x = 0, y = 1;
5 int test = x;    // test is not an atomic
6
7 void setTest() {
8     test = y;    // access to test is not synchronized
9 }
10
11 int main() {
12     thread run(setTest);
13
14     while (test == 0) { } // Since the access to test is not synchronized,
15                          // the compiler can assume, that test is only
16     run.join();          // changed by this thread, so it optimizes it
17     return 0;            // away and the program never terminates.
18 }
```

<https://godbolt.org/z/g7ZEXL>

Image Compression Strategy

Thread Basics

- Start
- Let's Improve
- 1st Correction



```
1 struct CompressContext{} ctx;           // Holds source and target pixel
2 void compress(CompressContext&) {}     // compresses a single stripe
3 void merge(CompressContext&) {}       // merges all compressed stripes
4
5 int main() {
6     const int ThreadNumber = 2;
7     vector<thread> threads{ThreadNumber};
8
9     for (auto& item : threads)
10        item = thread{ []{ compress(ctx); } };
11
12    for (auto& item : threads)
13        item.join();
14
15    merge(ctx);
16 }
```



- ▶ User threads are very expensive
- ▶ We have seen that it could take in seldom cases up to 1s to start a user thread
- ▶ Starting of more user threads than available cores leads to oversubscription. This leads to expensive context switches.

Usage of ThreadPool

```
1  const int TaskNumber{16};
2  atomic_int to_do{TaskNumber};
3
4
5
6  for (int i = 0; i < TaskNumber; ++i)
7      stlab::default_executor( // thread pool from stlab/concurrency
8          [&]() {
9              compress(data);
10             --to_do;
11         });
12
13
14
15  while (to_do != 0)
16
17
18  merge(data);
19 }
```

Usage of ThreadPool

```
1  const int TaskNumber{16};
2  atomic_int to_do{TaskNumber};
3  mutex block;
4  condition_variable cv;
5
6  for (int i = 0; i < TaskNumber; ++i)
7      stlab::default_executor( // thread pool from stlab/concurrency
8          [&]() {
9              compress(ctx);
10             --to_do;
11             cv.notify_one();
12         });
13
14     unique_lock lock{block};
15     while (to_do != 0)
16         cv.wait(lock);
17
18     merge(ctx);
19 }
```

The Adventure Begins

The Adventure

Felix Petriconi
felix@petriconi.net

Thread Basics

Start

Let's Improve

1st Correction



1st Race



```
1  const int TaskNumber{16};
2  atomic_int to_do{TaskNumber};
3  mutex block;
4  condition_variable cv;
5
6  for (int i = 0; i < TaskNumber; ++i)
7      stlab::default_executor( // thread pool from stlab/concurrency
8          [&]() {
9              compress(ctx);
10             --to_do;
11             cv.notify_one();
12         });
13
14  unique_lock lock{block};
15  while (to_do != 0)
16      cv.wait(lock);
17
18  merge(ctx);
19 }
```

1st Race

```
1  const int TaskNumber{16};
2  atomic_int to_do{TaskNumber};
3  mutex block;
4  condition_variable cv;
5
6  for (int i = 0; i < TaskNumber; ++i)
7      stlab::default_executor( // thread pool from stlab/concurrency
8          [&]() {
9              compress(ctx);
10             --to_do;           // Even if the shared variable is atomic,
11             cv.notify_one();  // it must be modified under the mutex in
12             });              // order to correctly publish the
13                             // modification to the waiting thread.
14  unique_lock lock{block};
15  while (to_do != 0)
16      cv.wait(lock);
17
18  merge(ctx);
19 }
```

https://en.cppreference.com/w/cpp/thread/condition_variable

1st Correction

```
1  const int TaskNumber{16};
2  atomic_int to_do{TaskNumber};
3  mutex block;
4  condition_variable cv;
5
6  for (int i = 0; i < TaskNumber; ++i)
7      stlab::default_executor( // thread pool from stlab/concurrency
8          [&]() {
9              compress(ctx);
10             --to_do;
11             cv.notify_one();
12         });
13
14  unique_lock lock{block};
15  while (to_do != 0)
16      cv.wait(lock);
17
18  merge(ctx);
19 }
```

1st Correction

```
1  const int TaskNumber{16};
2  atomic_int to_do{TaskNumber};
3  mutex block;
4  condition_variable cv;
5
6  for (int i = 0; i < TaskNumber; ++i)
7      stlab::default_executor(
8          [&]() {
9              compress(ctx);
10             {
11                 unique_lock guard{block};
12                 --to_do;
13             }
14             cv.notify_one();
15         });
16
17     unique_lock lock{block};
18     while (to_do != 0)
19         cv.wait(lock);
20
21     merge(ctx);
22 }
```

2nd Race



```
1  const int TaskNumber{16};
2  atomic_int to_do{TaskNumber};
3  mutex block;
4  condition_variable cv;
5
6  for (int i = 0; i < TaskNumber; ++i)
7      stlab::default_executor(
8          [&]() {
9              compress(ctx);
10             {
11                 unique_lock guard{block};
12                 --to_do;
13             }
14             cv.notify_one();
15         });
16
17     unique_lock lock{block};
18     while (to_do != 0)
19         cv.wait(lock);
20
21     merge(ctx);
22 }
```


Main Thread

```
{
  int TaskNumber{4};
  int to_do{tasks};
  mutex block;
  condition_variable cv;

  unique_t lock(block);
  while (to_do != 0) {
    cv.wait(lock);
  }
}
-conditon_variable()
```

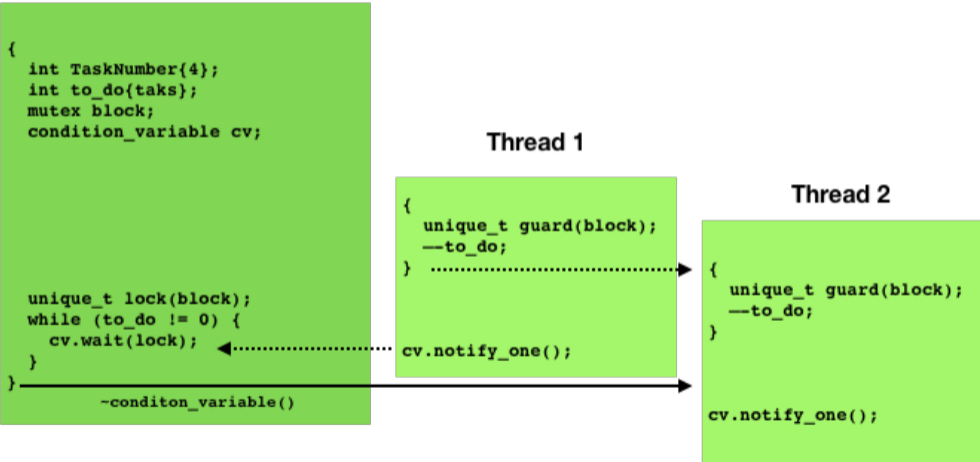
Thread 1

```
{
  unique_t guard(block);
  --to_do;
} .....
cv.notify_one();
```

Thread 2

```
{
  unique_t guard(block);
  --to_do;
}

cv.notify_one();
```



2nd Correction



```
1  const int TaskNumber{16};
2  atomic_int to_do{TaskNumber};
3  mutex block;
4  condition_variable cv;
5
6  for (int i = 0; i < TaskNumber; ++i)
7      stlab::default_executor(
8          [&]() {
9              compress(ctx);
10             {
11                 unique_lock guard{block};
12                 --to_do;
13             }
14             cv.notify_one();
15         });
16
17     unique_lock lock{block};
18     while (to_do != 0)
19         cv.wait(lock);
20
21     merge(ctx);
22 }
```

2nd Correction

```
1  const int TaskNumber{16};
2  atomic_int to_do{TaskNumber};
3  mutex block;
4  condition_variable cv;
5
6  for (int i = 0; i < TaskNumber; ++i)
7      stlab::default_executor(
8          [&]() {
9              compress(ctx);
10             {
11                 unique_lock guard{block};
12                 --to_do;
13                 cv.notify_one();
14             }
15         });
16
17     unique_lock lock{block};
18     while (to_do != 0)
19         cv.wait(lock);
20
21     merge(ctx);
22 }
```

Adding Error Handling

```
1 struct CompressContext{} ctx;
2 bool compress(CompressContext&) // true when OK, false when failed
3 void merge(CompressContext&) {}
4
5 int main() {
6     const int TaskNumber{16};
7     int to_do{TaskNumber};
8     atomic_bool abort{false};
9     mutex block;
10    condition_variable cv;
11
12    for (int i = 0; i < TaskNumber; ++i)
13        stlab::default_executor(
14            [&]() {
15                if (abort) return;
16                auto do_abort = !compress(ctx);
17                {
18                    unique_lock guard{block};
19                    --to_do;
20                    abort = do_abort || abort;
21                    cv.notify_one();
22                }
23            });
```

Adding Error Handling

```
1  int to_do{TaskNumber};
2  atomic_bool abort{false};
3  mutex block;
4  condition_variable cv;
5
6  for (int i = 0; i < TaskNumber; ++i)
7      stlab::default_executor(
8          [&]() {
9              if (abort) return;
10             auto do_abort = !compress(ctx);
11             {
12                 unique_lock guard{block};
13                 --to_do;
14                 abort = do_abort || abort;
15                 cv.notify_one();
16             }
17         });
18
19  unique_lock lock{block};
20  while (to_do != 0 && !abort)
21      cv.wait(lock);
22
23  merge(ctx);
```

3rd Race

```
1  const int TaskNumber{16};
2  int to_do{TaskNumber};
3  atomic_bool abort{false};
4  mutex block;
5  condition_variable cv;
6
7  for (int i = 0; i < TaskNumber; ++i)
8      stlab::default_executor(
9          [&]() {
10             if (abort) return;
11             auto do_abort = !compress(ctx);
12             {
13                 unique_lock guard{block};
14                 --to_do;
15                 abort = do_abort || abort;
16                 cv.notify_one();
17             }
18         });
19
20 unique_lock lock{block};
21 while (to_do != 0 && !abort)
22     cv.wait(lock);
```

3rd Race

```
1  const int TaskNumber{16};           // None of these variables exist
2  int to_do{TaskNumber};             // when there is an early exit
3  atomic_bool abort{false};          // with abort is been set. So the
4  mutex block;                       // other running tasks must not
5  condition_variable cv;             // use it further.
6
7  for (int i = 0; i < TaskNumber; ++i)
8      stlab::default_executor(
9          [&]() {
10             if (abort) return;
11             auto do_abort = !compress(ctx);
12             {
13                 unique_lock guard{block};
14                 --to_do;
15                 abort = do_abort || abort;
16                 cv.notify_one();
17             }
18         });
19
20     unique_lock lock{block};
21     while (to_do != 0 && !abort)
22         cv.wait(lock);
```

```
1 struct CompressContext{} ctx;
2 bool compress(CompressContext&)
3 { return true; }
4 void merge(CompressContext&) {}
5
6 struct ProcessContext
7 {
8     mutex block;
9     condition_variable cv;
10    int to_do = 0;
11    atomic_bool abort{false};
12};
```


3rd Correction

```
1  const int TaskNumber{16};
2  auto pctx = make_shared<ProcessContext>();
3  pctx->to_do = TaskNumber;
4  for (int i = 0; i < TaskNumber; ++i)
5      stlab::default_executor(
6          [_weakContext = weak_ptr<ProcessContext>(pctx)] {
7              auto p = _weakContext.lock();
8              if (!p || p->abort)
9                  return;
10             auto do_abort = !compress(ctx);
11             {
12                 unique_lock guard{p->block};
13                 --p->to_do;
14                 p->abort = do_abort || p->abort;
15                 p->cv.notify_one();
16             }
17         });
18
19     unique_lock lock{pctx->block};
20     while (pctx->to_do != 0 && !pctx->abort)
21         pctx->cv.wait(lock);
22
23     merge(ctx);
```

Conclusion

The Adventure


Felix Petriconi
felix@petriconi.net

Thread Basics

Start

Let's Improve

1st Correction



²Geni - photo by user:geni, CC BY-SA 4.0
<https://commons.wikimedia.org/w/index.php?curid=71925797>

- ▶ It is easy to get a CPU with more cores.
- ▶ It is hard to write concurrent code correct.
- ▶ It is even harder to use low synchronization primitives correctly.

Try to use high level abstractions like

- ▶ Future
- ▶ Channel
- ▶ Actor
- ▶ ...

Example with boost futures

```
1 #include <vector>
2 #define BOOST_THREAD_PROVIDES_FUTURE
3 #define BOOST_THREAD_PROVIDES_FUTURE_CONTINUATION
4 #define BOOST_THREAD_PROVIDES_FUTURE_WHEN_ALL_WHEN_ANY
5 #include <boost/thread/future.hpp>
6
7 using std::vector;
8
9 struct CompressContext{} ctx;
10 bool compress(CompressContext&)
11 { return true; }
12 void merge(CompressContext&) {}
13
14 int main() {
15     vector<boost::future<void>> tasks{16};
16
17     for (auto& f : tasks)
18         f = boost::async( []{ compress(ctx); } );
19
20     auto done = boost::when_all(tasks.begin(), tasks.end())
21         .then([](auto) { merge(ctx); });
22 }
```

Example with stlab futures

```
1 struct CompressContext{} ctx;
2 bool compress(CompressContext&)
3 { return true; }
4 void merge(CompressContext&) {}
5
6 int main() {
7     size_t TaskNumber{16};
8
9     vector<stlab::future<void>> tasks{TaskNumber};
10
11     for (auto& task : tasks)
12         task = stlab::async(stlab::default_executor,
13                             [] { compress(ctx); });
14
15     auto done = stlab::when_all(stlab::default_executor,
16                                []{ merge(ctx); }, make_pair(tasks.begin(), tasks.end()) );
17
18     stlab::blocking_get(done);
```

Image Compression Processing Comparison

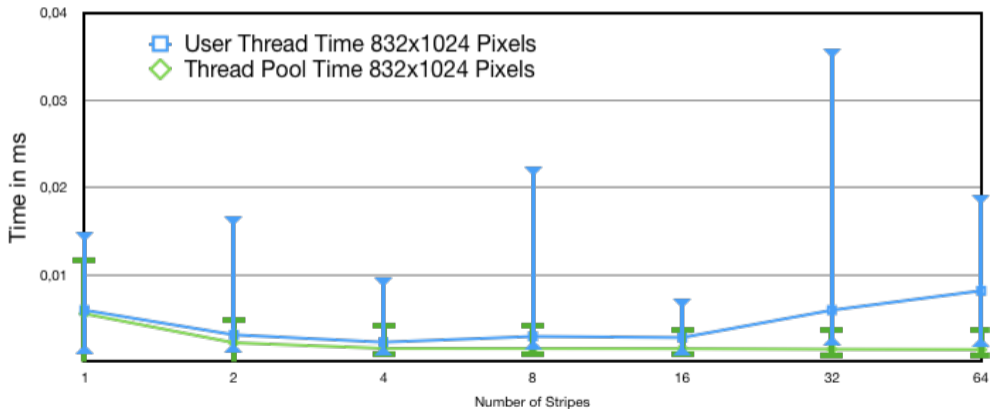





Image Compression Processing Comparison



- ▶ Try to break down your problem into small parts that can be solved without any synchronization.
- ▶ Whenever it is possible prefer high level abstractions over low level synchronization primitives.
- ▶ Try to think in parallel. Have in mind that
 - ▶ any operation can be interrupted at any time,
 - ▶ e.g. between any lines or even within a single line.
 - ▶ This is true for hidden code, e.g. destructors too.

- ▶ My family, who supports me in my work on the concurrency library and this conference.
- ▶ Sean Parent, who taught me over time lots about concurrency and abstraction. He gave me the permission to use whatever I needed from his presentations for my own.
- ▶ My company MeVis Medical Solutions AG, who give me the possibility to be here.
- ▶ The C++ UserGroup in Bremen, where I can test my sessions.
- ▶ All contributors to the slab library.

- ▶ Concurrency library <https://github.com/stlab/libraries>
- ▶ Documentation <http://stlab.cc/libraries>
- ▶ Communicating Sequential Processes by C. A. R. Hoare
<http://usingcsp.com/cspbook.pdf>
- ▶ The Theory and Practice of Concurrency by A.W. Roscoe <http://www.cs.ox.ac.uk/people/bill.roscoe/publications/68b.pdf>
- ▶ Towards a Good Future, C++ Standard Proposal by Felix Petriconi, David Sankel and Sean Parent <http://open-std.org/JTC1/SC22/WG21/docs/papers/2017/p0676r0.pdf>
- ▶ A Unified Futures Proposal for C++ by Bryce Adelstein Lelbach, et al
<http://www.open-std.org/jtc1/sc22/wg21/docs/papers/2018/p1054r0.html>

Software Principles and Algorithms

- ▶ Elements of Programming by Alexander Stepanov, Paul McJones, Addison Wesley
- ▶ From Mathematics to Generic Programming by Alexander Stepanov, Daniel Rose, Addison Wesley

Concurrency and Parallelism

- ▶ HPX <http://stellar-group.org/libraries/hpx/>
- ▶ C++CSP <https://www.cs.kent.ac.uk/projects/ofa/c++csp>
- ▶ CAF_C++ Actor Framework <http://actor-framework.org/>
- ▶ C++ Concurrency In Action by Anthony Williams, Manning, 2nd Edition

- ▶ Goals for better code by Sean Parent:
<http://sean-parent.stlab.cc/papers-and-presentations>
- ▶ Goals for better code by Sean Parent: Concurrency:
<https://youtu.be/au0xX4h8SCI?t=16354>
- ▶ Future Ruminations by Sean Parent <http://sean-parent.stlab.cc/2017/07/10/future-ruminations.html>
- ▶ CppCast with Sean Parent <http://cppcast.com/2015/06/sean-parent/>
- ▶ Thinking Outside the Synchronization Quadrant by Kevlin Henney:
<https://vimeo.com/205806162>
- ▶ Inside Windows 8 Thread Pool <https://channel9.msdn.com/Shows/Going+Deep/Inside-Windows-8-Pedro-Teixeira-Thread-pool>

Reference

Reference

Further listening and
viewing

Contact



stlab::future

Source: <https://github.com/stlab/libraries>

Documentation: <https://www.stlab.cc/libraries>

Thank's for your attention!

- ▶ Mail: `felix@petriconi.net`
- ▶ GitHub: `https://github.com/FelixPetriconi`
- ▶ Web: `https://petriconi.net`
- ▶ Twitter: `@FelixPetriconi`

Q & A

- ▶ Mail: `felix@petriconi.net`
- ▶ GitHub: `https://github.com/FelixPetriconi`
- ▶ Web: `https://petriconi.net`
- ▶ Twitter: `@FelixPetriconi`

Feedback is always welcome!