The Art of Writing Reasonable Concurrent Code

Pre-Conference Workshop ACCU 2017

©2017

Felix Petriconi

2017-04-25

The [C++] language is too large for *anyone* to master
So *everyone* lives within a subset

Sean Parent, C++Now, 2012

Felix Petriconi

- School (UCSD Pascal, Turbo Pascal)
- Studied electrical engineering (Modula 2, Ada, C++)
- Student research assistant (1992-1996) (Turbo Pascal, C++, C)
- Freelance programmer 1996-2003 (Ericsson, Siemens-VDO, etc.)
  - Development of test software for embedded devices (Perl, C)
- Programmer and development manager 2003-today at MeVis Medical Solutions AG, Bremen, Germany
  - Development of medical devices in the area of mammography and radio therapy (C++, Ruby, Python)
- Programming activities:
  - Blog editor of ISO C++ website
  - Active member of C++ User Group Bremen
  - Contributor to Sean Parent's concurrency library
  - Member of ACCU conference committee
- Married with Nicole, having three children, living near Bremen, Germany

Why am I here?

I saw how we developed multi threaded code in the past.
I saw how easy it is to make mistakes.
I saw and still see how difficult it is to maintain this code.

I watched recordings from Sean Parent's talks about "Better Code".
I was impressed.
I wanted to learn more.

I'm collaborating in his open source project for a new library.
I'm continuously learning there a lot.
I care about sharing my knowledge, here at the ACCU conference.

Why are you here?

What is your motivation to be here?

Problems from my domain

- Loading of huge images blocks UI
- Storing of files blocks UI
- Re-coding of huge images takes very long
- DB accesses takes too long
- ...

Why do we have to talk about concurrency?

The free lunch is over!

Herb Sutter, 2005[1]

[1] The Free Lunch Is Over: A Fundamental Turn Toward Concurrency in Software
http://www.gotw.ca/publications/concurrency-ddj.htm

The free lunch is over

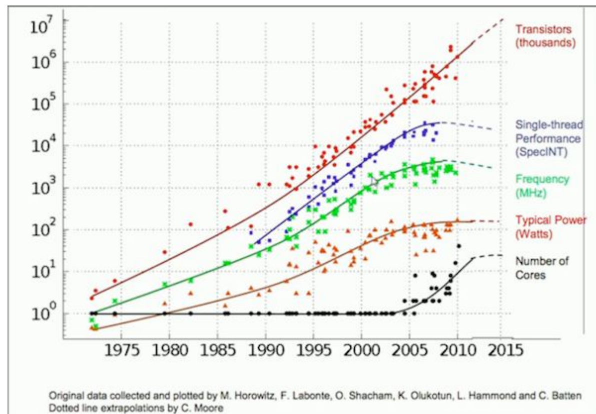

Desktop Compute Power

8-core 3.5GHz (Sandy Bridge + AMD Radeon 6950)

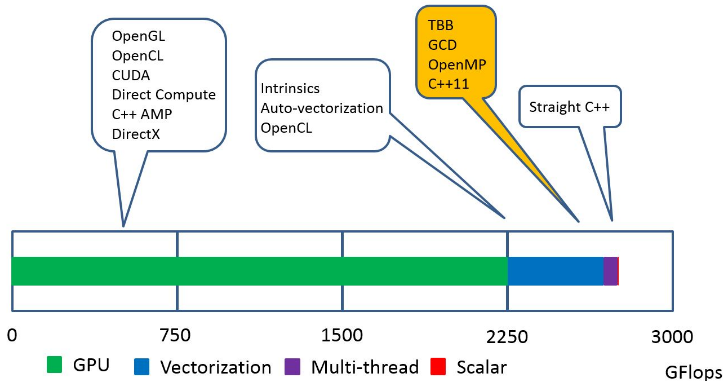

That's what we are targeting for!

Amdahl's Law

0% Synchronization

$$S(N) = \frac{1}{(1-P)+\frac{P}{N}}$$
$$P = 0$$

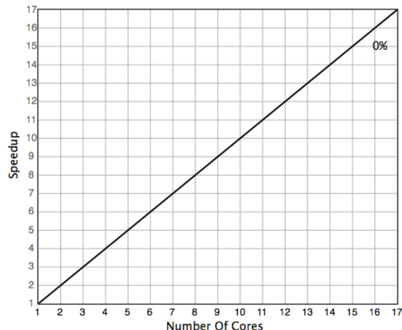

Amdahl's Law

90% Synchronization

$S(N) = \frac{1}{(1-P)+\frac{P}{N}}$
$P = 0.9$

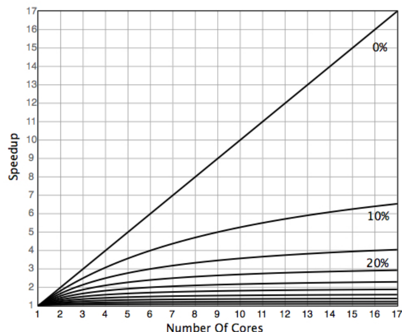

Outline Futures

Futures
- Why Futures?
- Introduction
- C++ Standard - Futures
  - Exceptions
  - Deficiencies
- Boost - Futures
  - Deficiencies
  - Future Continuation
  - Future Join
- stlab - Futures
  - Executors
  - Error Recovery
  - Join
  - Splits
- Exercise 1

Why Futures?

Futures
Why Futures?
Introduction
C++ Standard - Futures
Boost - Futures
stlab - Futures
Exercise 1

Why using futures?
Aren't threads, mutex, atomics great?
They are great tools "to shot yourself into the foot!"
It is so easy

- having race conditions
- having dead locks
- wasting CPU cycles through contention

Do you program your application in assembly?
Only if it absolute time critical.
Then don't use tools from the level of assembly!

Future Introduction

- Futures provide a mechanism to separate a function from its result
- After the function is called the result appears "magically" in the future
- A future is a token to the result of a function
- Added with C++11
- Futures, resp. promises where invented 1977/1978 by Daniel P. Friedman, David Wise, Henry Baker and Carl Hewitt

Futures
Why Futures?
Introduction
C++ Standard - Futures
Boost - Futures
stlab - Futures
Exercise 1

C++ Standard - Futures

```cpp
#include <future>
#include <iostream>

using namespace std;

int main() {

  auto getTheAnswer = [] {
    this_thread::sleep_for(chrono::milliseconds(815));
    return 42;
  };

  future<int> f2 = async(launch::async, getTheAnswer);

  // Do other stuff, getting the answer may take longer
  cout << f2.get() << '\n';         // access the value
}
```

Output

42

Futures
Why Futures?
Introduction
C++ Standard - Futures
Exceptions
Deficiencies
Boost - Futures
stlab - Futures
Exercise 1

C++ Standard - Futures - Exceptions

```
#include <future>
#include <iostream>
#include <exception>

using namespace std;

int main() {
  auto getTheAnswer = [] {
    throw runtime_error("Bad things happened: Vogons appeared!");
    return 42;
  };

  future<int> f2 = async(launch::async, getTheAnswer);

  // Do other stuff, getting the answer may take longer
  try {
    cout << f2.get() << '\n';  // try accessing the value
                               // rethrows the stored exception
  }
  catch (const runtime_error& ex) {
    cout << ex.what() << '\n';
  }
}
```

Output

Bad things happened: Vogons appeared!

Futures
Why Futures?
Introduction
C++ Standard - Futures
Exceptions
Deficiencies
Boost - Futures
stlab - Futures
Exercise 1

C++11/14 Future Deficiencies

- No continuation – .then() ✗*
- No join – .when_all() and .when_any() ✗*
- No split – continuation in different directions ✗
- No cancellation (but can be modelled) ✗
- No progress monitoring (except ready) ✗
- No custom executor ✗
- Blocks on destruction (may even blocks until termination of used thread) ✗
- .get() has two problems:
  1. One thread resource is consumed which increases contention and possibly causing a deadlock ✗
  2. Any subsequent non-dependent calculations on the task are also blocked ✗
- Don't behave as a regular type ✗

∗ Comes with C++17(TS)

Futures
Why Futures?
Introduction
C++ Standard - Futures
Exceptions
Deficiencies
Boost - Futures
stlab - Futures
Exercise 1

Boost - Futures

- Continuation – .then() ✓
- Join – .when_all() and .when_any() ✓
- No split – continuation in different directions ✗
- No cancellation (but can be modelled) ✗
- No progress monitoring (except ready) ✗
- Custom executor ✓
- Blocks on destruction (may even blocks until termination of used thread) ✗
- .get() has two problems:
  1. One thread resource is consumed which increases contention and possibly causing a deadlock ✗
  2. Any subsequent non-dependent calculations on the task are also blocked ✗
- Don't behave as a regular type ✗

Futures
Why Futures?
Introduction
C++ Standard - Futures
Boost - Futures
Deficiencies
Future Continuation
Future Join
stlab - Futures
Exercise 1

Future Continuation

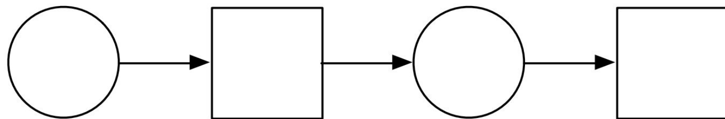

Futures
Why Futures?
Introduction
C++ Standard - Futures
Boost - Futures
Deficiencies
Future Continuation
Future Join
stlab - Futures
Exercise 1

C++17(TS) / Boost - Continuation

```
#include <iostream>
#include <boost/thread/future.hpp>

using namespace std;

int main() {
  boost::future<int> answer = boost::async([]{ return 42; });

  boost::future<void> done = answer.then(
    [](boost::future<int> a) { std::cout << a.get() << '\n';} );

  // do something else
  done.wait();       // waits until future done is fulfilled
}
```

Output

42

Futures
Why Futures?
Introduction
C++ Standard - Futures
Boost - Futures
Deficiencies
Future Continuation
Future Join
stlab - Futures
Exercise 1

Future Join

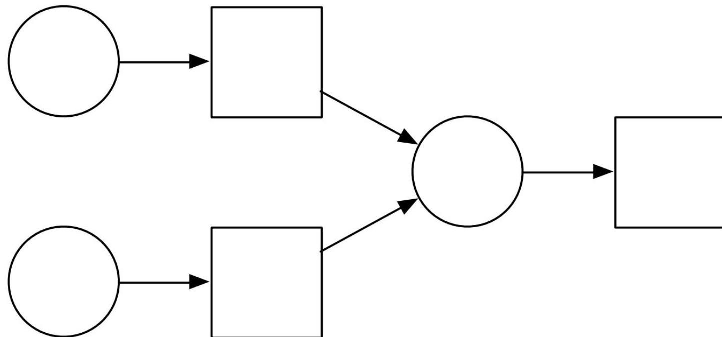

C++17(TS) / Boost - Join

```
#include <iostream>
#include <boost/thread/future.hpp>

using namespace std;

int main() {
  auto a = boost::async([]{ return 40; });
  auto b = boost::async([]{ return 2; });

  auto answer = boost::when_all(std::move(a), std::move(b)).then(
    [](auto f) {
      auto t = f.get();
      return get<0>(t).get() + get<1>(t).get();
    });

  // wait for the something else
  cout << answer.get() << '\n';
}
```

Output

42

Futures
Why Futures?
Introduction
C++ Standard - Futures
Boost - Futures
Deficiencies
Future Continuation
Future Join
stlab - Futures
Exercise 1

C++17(TS) / Boost - Join

```
#include <iostream>
#include <boost/thread/future.hpp>

using namespace std;

int main() {
  auto a = boost::async([]{ return 40; });
  auto b = boost::async([]{ return 2; });

  auto answer = boost::when_all(std::move(a), std::move(b)).then(
    [](auto f) {
      auto t = f.get();
      return get<0>(t).get() + get<1>(t).get();
    });

  // wait for the something else
  cout << answer.get() << '\n';
}
```

What is the type of f?

f is a future tuple of futures:

`future<tuple<future<int>, future<int>>>`

Futures

Why Futures?

Introduction

C++ Standard - Futures

Boost - Futures

Deficiencies

Future Continuation

Future Join

stlab - Futures

Exercise 1

stlab - Futures

stlab::future

Source: `https://github.com/stlab/libraries`
Documentation: `http://www.stlab.cc/libraries`

Futures
Why Futures?
Introduction
C++ Standard - Futures
Boost - Futures
stlab - Futures
Executors
Error Recovery
Join
Splits
Exercise 1

stlab - Futures

- Continuation – .then() ✓
- Join – .when_all() and .when_any() ✓
- Split – continuation in different directions ✓
- Cancellation ✓
- No progress monitoring (except ready), more planned ✗
- Custom executor ✓
- Do not block on destruction ✓
- Behave as a regular type ✓
- Additional dependencies:
  - C++14: boost (optional, variant)
  - C++17: none

Futures
Why Futures?
Introduction
C++ Standard - Futures
Boost - Futures
stlab - Futures
Executors
Error Recovery
Join
Splits
Exercise 1

stlab::future

```cpp
#include <stlab/future.hpp>
#include <stlab/default_executor.hpp>
#include <iostream>
using namespace std;
int main() {
  auto getTheAnswer = [] {
    this_thread::sleep_for(chrono::milliseconds(815));
    return 42;
  };
  stlab::future<int> f =
    stlab::async(
      stlab::default_executor,// default_executor
                              // uses platfrom thread pool on Win/OSX
                              // uses stlab thread pool on other OS
      getTheAnswer
    );

  while (!f.get_try()) {      // does not block
    // Do other stuff, getting the answer may take longer :-)
  }

  cout << f.get_try().value() << '\n'; // access the value
                           // throws exception .value() if not ready
}
```

Output

42

Futures
Why Futures?
Introduction
C++ Standard - Futures
Boost - Futures
stlab - Futures
Executors
Error Recovery
Join
Splits
Exercise 1

stlab::future - Exceptions

```cpp
#include <stlab/future.hpp>
#include <stlab/default_executor.hpp>
#include <iostream>
#include <exception>

int main() {
  auto getTheAnswer = [] {
    throw std::runtime_error("Bad thing happened: Vogons appeared!");
    return 42;
  };
  auto f = stlab::async(stlab::default_executor, getTheAnswer);

  try {
    while (!f.get_try()) { // try accessing the value
                           // may rethrow a stored exception
      // Do other stuff, getting the answer may take longer
    }

    std::cout << f.get_try().value() << '\n';
  }
  catch (const std::runtime_error& ex) {
    std::cout << ex.what() << '\n';
  }
}
```

Output

Bad things happened: Vogons appeared!

Futures
Why Futures?
Introduction
C++ Standard - Futures
Boost - Futures
stlab - Futures
Executors
Error Recovery
Join
Splits
Exercise 1

stlab::future - Continuation

```
#include <stlab/future.hpp>
#include <stlab/default_executor.hpp>
#include <iostream>

int main() {
  auto answer =
    stlab::async(stlab::default_executor, []{ return 42; } );

  stlab::future<void> done = answer.then(
    [](int a)                    // pass by value and not by future
    {
      std::cout << a << '\n';
    });

  while (!done.get_try()) {
    // do something in the meantime
  }
}
```

Output

42

Futures
Why Futures?
Introduction
C++ Standard - Futures
Boost - Futures
stlab - Futures
Executors
Error Recovery
Join
Splits
Exercise 1

Executors

- Executors are needed to customize where the task shall be executed
- Executors can be general thread pools, serial queues, main queues, dedicated task groups, etc.

Futures
Why Futures?
Introduction
C++ Standard - Futures
Boost - Futures
stlab - Futures
Executors
Error Recovery
Join
Splits
Exercise 1

stlab::future - Continuation with Custom Executor

```
 1 #include <stlab/future.hpp>
 2 #include <stlab/default_executor.hpp>
 3 #include <iostream>
 4 #include <QLineEdit>
 5 #include "QtScheduler.h"
 6
 7 int main() {
 8   QLineEdit theAnswerEdit;
 9
10   auto answer =
11     stlab::async(stlab::default_executor, []{ return 42; } );
12
13   stlab::future<void> done = answer.then(
14     QtScheduler(),                     // different scheduler
15     [&](int a) { theAnswerEdit.setValue(a); }// here update in main
        thread
16   );
17
18   while (!done.get_try()) {
19     // do something in the meantime
20   }
21 }
```


Futures
Why Futures?
Introduction
C++ Standard - Futures
Boost - Futures
stlab - Futures
Executors
Error Recovery
Join
Splits
Exercise 1

stlab::future - Custom Executor

Futures
Why Futures?
Introduction
C++ Standard - Futures
Boost - Futures
stlab - Futures
Executors
Error Recovery
Join
Splits
Exercise 1

- In boost, executors derive from a common base class
- In stlab the executors must provide
  `template <typename F> void operator()(F f)`
- Let's build exemplary a custom executor for the Qt GUI, that allows to perform updates in the Qt main event loop

stlab::future - Custom Executor - Qt

```
#include <QApplication>
#include <Event>

class QtExecutor
{
  using result_type = void;

  class ExecutorEvent : public QEvent
  {
  };

public:
  template <typename F>
  void operator()(F f) {
    auto event = new ExecutorEvent(std::move(f));
    QApplication::postEvent(event->receiver(), event);
  }
};
```


Futures
Why Futures?
Introduction
C++ Standard - Futures
Boost - Futures
stlab - Futures
Executors
Error Recovery
Join
Splits
Exercise 1

stlab::future - Custom Executor - Qt cont. I

```
1  #include <QApplication>
2  #include <Event>
3
4  class QtExecutor
5  {
6    using result_type = void;
7
8    class EventReceiver;
9
10   class ExecutorEvent : public QEvent
11   {
12     std::function<void()> _f;
13     std::unique_ptr<EventReceiver> _receiver;
14
15   public:
16     ExecutorEvent(std::function<void()> f)
17       : QEvent(QEvent::User)
18       , _f(std::move(f))
19       , _receiver(new EventReceiver()) {
20       _receiver()->moveToThread(QApplication::instance()->thread());
21     }
22
23     void execute() { _f(); }
24
25     QObject *receiver() const { return _receiver.get(); }
26   };
27
28 public:
29 };
```


Futures
Why Futures?
Introduction
C++ Standard - Futures
Boost - Futures
stlab - Futures
Executors
Error Recovery
Join
Splits
Exercise 1

stlab::future - Custom Executor - Qt cont. II

```
1  #include <QApplication>
2  #include <Event>
3
4  class QtExecutor
5  {
6    class ExecutorEvent : public QEvent
7    {
8      QObject *receiver() const { return _receiver.get(); }
9    };
10
11   class EventReceiver : public QObject
12   {
13   public:
14     bool event(QEvent *event) override {
15       auto myEvent = dynamic_cast<ExecutorEvent*>(event);
16       if (myEvent) {
17         myEvent->execute();
18         return true;
19       }
20       return false;
21     }
22   };
23
24 public:
25   template <typename F>
26   void operator()(F f) {
27     auto event = new ExecutorEvent(std::move(f));
28     QApplication::postEvent(event->receiver(), event);
29   }
30 };
```


Futures
Why Futures?
Introduction
C++ Standard - Futures
Boost - Futures
stlab - Futures
Executors
Error Recovery
Join
Splits
Exercise 1

stlab::future - Error Recovery

```
int main() {
  auto getTheAnswer = [] {
    throw std::runtime_error("Bad thing happened: Vogons appeared");
    std::cout << "I have got the answer\n"; return 42;
  };
  auto handleTheAnswer = [](int v) {
    if (v == 0) std::cout << "We have a problem!\n";
    else std::cout << "The answer is " << v << '\n';
  };

  auto f = stlab::async(stlab::default_executor, getTheAnswer)
    .recover([](stlab::future<int> result) {
      if (result.error()) {
        std::cout << "Listen to Vogon poetry!\n";
        return 0;
      }
      return result.get_try().value();
  }).then(handleTheAnswer);

  while (!f.get_try());
}
```

Output

Listen to Vogon poetry!
We have a problem!

Futures
Why Futures?
Introduction
C++ Standard - Futures
Boost - Futures
stlab - Futures
Executors
Error Recovery
Join
Splits
Exercise 1

stlab::future - Join

```
 1 #include <stlab/future.hpp>
 2 #include <stlab/default_executor.hpp>
 3 #include <iostream>
 4
 5 using namespace stlab;
 6
 7 int main() {
 8   auto a = async(default_executor,[]{ return 40; });
 9   auto b = async(default_executor,[]{ return 2; });
10
11   auto answer = when_all(
12     default_executor,
13     [](int x, int y) { return x + y; },
14     a, b);
15
16   while (!answer.get_try()) {
17   // wait for something else
18   }
19   std::cout << answer.get_try().value() << '\n';
20 }
```

Output

42

Futures
Why Futures?
Introduction
C++ Standard - Futures
Boost - Futures
stlab - Futures
Executors
Error Recovery
Join
Splits
Exercise 1

future - Split

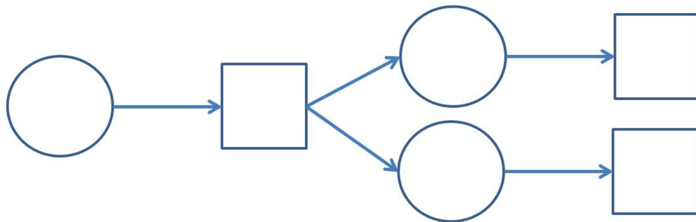

Futures
Why Futures?
Introduction
C++ Standard - Futures
Boost - Futures
stlab - Futures
Executors
Error Recovery
Join
Splits
Exercise 1

stlab::future - Split

```
#include <stlab/future.hpp>
#include <stlab/default_executor.hpp>
#include <iostream>

using namespace stlab;

int main() {
  auto answer = async(default_executor,[]{ return 42; });

  auto dent = answer.then([](int a) {
    std::cout << "Tell the answer " << a << " Arthur Dent\n";
  });

  auto marvin = answer.then([](int a) {
    std::cout << "May the answer " << a << " shear up Marvin\n";
  });

  while (!dent.get_try() && !marvin.get_try()) {
  // wait for  something else
  }
}
```

Output

Tell the answer May the answer 42 Arthur Dent

42 shear up Marvin ⇛ **Race condition by using std::cout**

Futures

Why Futures?

Introduction

C++ Standard - Futures

Boost - Futures

stlab - Futures

Executors

Error Recovery

Join

Splits

Exercise 1

Exercise 1

Change the application in a way that

- using Start does not block the UI,
- it is possible to cancel the running operation,
- it is possible to restart it.

Futures
Why Futures?
Introduction
C++ Standard - Futures
Boost - Futures
stlab - Futures
Exercise 1

Conclusion

Futures are a great concept to structure the code so that it runs with minimal contention.

After a single execution the graph cannot be used any more.

Futures
Why Futures?
Introduction
C++ Standard - Futures
Boost - Futures
stlab - Futures
Exercise 1

Outline Channels

Channel Motivation

Channel - Stateless Process
- Channel - Split
- Channel - Join
- Exercise 2

Channel Stateful Process
- Exercise 3

Channel Introduction

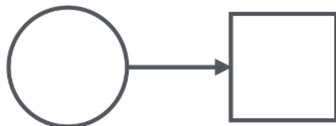

- Each change triggers a notification to the sink values
- Channels allow the creation of persistent execution graphs
- This is also known as reactive programming model
- First published by Tony Hoare 1978

Channel - Stateless Process

```cpp
#include <stlab/channel.hpp>
#include <stlab/default_executor.hpp>
#include <iostream>
int main() {
  stlab::sender<int> send;         // sending part of the channel
  stlab::receiver<int> receiver; // receiving part of the channel
  std::tie(send, receiver) =      // combining both to a channel
    stlab::channel<int>(stlab::default_executor);

  auto printer =
    [](int x){ std::cout << x << '\n'; }; // stateless process

  auto printer_process =
    receiver | printer;           // attaching process to the receiving
                                  // part
  receiver.set_ready();           // no more processes will be attached
                                  // process starts to work
  send(1); send(2); send(3);      // start sending into the channel

  int end; std::cin >> end;       // simply wait to end application
}
```

Output

1
2
3

Channel - Split

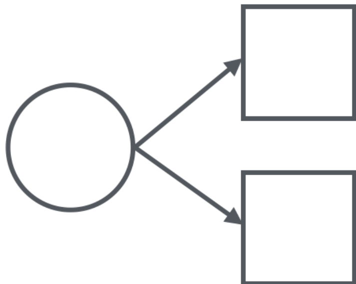

Channel Motivation

Channel - Stateless Process

Channel - Split

Channel - Join

Exercise 2

Channel Stateful Process

Channel - Split Process

```
using namespace stlab;
int main() {
  sender<int> send;
  receiver<int> receiver;
  std::tie(send, receiver) = channel<int>(default_executor);

  auto printerA = [](int x){ printf("Process A %d\n", x); };
  auto printerB = [](int x){ printf("Process B %d\n", x); };

  auto printer_processA = receiver | printerA;
  auto printer_processB = receiver | printerB;

  receiver.set_ready();          // no more processes will be attached
                                 // process may start to work
  send(1); send(2); send(3);
  int end; std::cin >> end;
}
```

Output

Process A 1
Process B 1
Process A 2
Process B 2
Process B 3
Process A 3

Channel Motivation
Channel - Stateless Process
Channel - Split
Channel - Join
Exercise 2
Channel Stateful Process

Channel - Join

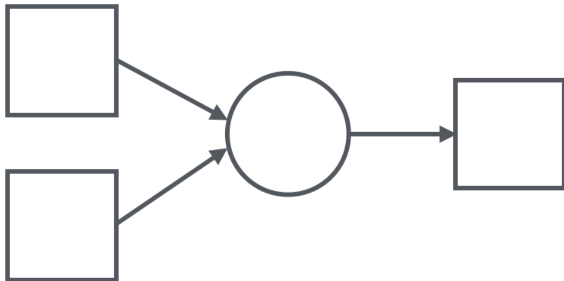

Channel Motivation
Channel - Stateless Process
Channel - Split
Channel - Join
Exercise 2
Channel Stateful Process

Channel - Joined Processes

```
using namespace stlab;

int main() {
  sender<int> sendA, sendB;
  receiver<int> receiverA, receiverB;
  std::tie(sendA, receiverA) = channel<int>(default_executor);
  std::tie(sendB, receiverB) = channel<int>(default_executor);

  auto printer = [](int x, int y){ printf("Process %d %d\n", x, y); };

  auto printProcess = join(default_executor, printer,
    receiverA, receiverB);

  receiverA.set_ready();
  receiverB.set_ready();

  sendA(1); sendA(2); sendB(3); sendA(4); sendB(5); sendB(6);

  int end; std::cin >> end;
}
```

Output

Process 1 3
Process 2 5
Process 4 6

Channel operators

Beside `join()` there are:

- `zip()` The process takes the passed values in a round-robin manner, starting with the result from the first receiver.
- `merge()` The process takes the values in an arbitrary order.

Channel Motivation

Channel - Stateless Process

Channel - Split

Channel - Join

Exercise 2

Channel Stateful Process

Exercise 2

Create a process chain with

- the inputs
  - one `int` input
  - one `std::string` input
  - one `double` input
- all inputs are joined to a process that concatenates all the results into a string and
- the result is split into
  - one process that prints the result into console,
  - one process that stores the result into a file
- show with two value triplets, that the implementation works
- don't use any synchronization primitive

Conclusion

- Stateless processes (from the point of view of the channel) have a 1:1 relationship from input to output

Channel Motivation
Channel - Stateless Process
Channel - Split
Channel - Join
Exercise 2
Channel Stateful Process

Channel Stateful Process - Motivation

- Some problems need a processor with state
- Some problems have an n : m relationship from input to output
- The picture becomes more complicated with states:
  - When to proceed?
  - How to handle situations when less than expected values come downstream?

Channel - Stateful Process Signature


```
#include <stlab/channel.hpp>

using process_state_scheduled =
  std::pair<process_state, std::chrono::system_clock::time_point>;

struct process_signature
{
    void await(T... val);

    U yield();

    process_state_scheduled state() const;

    void close();                          // optional

    void set_error(std::exception_ptr); // optional
};
```

Stateful Process Signature - await

```cpp
#include <stlab/channel.hpp>

using process_state_scheduled =
  std::pair<process_state, std::chrono::system_clock::time_point>;

struct process_signature
{
    void await(T... val);

    U yield();

    process_state_scheduled state() const;

    void close();                          // optional

    void set_error(std::exception_ptr); // optional
};
```

The `await` method is called on the process whenever a new value was received from upstream. The type T stands here for any semi regular or move-only type. The number of arguments depends on the number of attached upstream sender. Potential state changes from awaitable to yieldable should happen while this method is invoked.

Stateful Process Signature - yield

```
#include <stlab/channel.hpp>

using process_state_scheduled =
  std::pair<process_state, std::chrono::system_clock::time_point>;

struct process_signature
{
    void await(T... val);

    U yield();

    process_state_scheduled state() const;

    void close();                            // optional

    void set_error(std::exception_ptr); // optional
};
```

The `yield` method is called on the process whenever the `process_state_scheduled.first` is `process_state::yield` or a timeout was provided with the recent call to `state()` and that has elapsed.

Stateful Process Signature - state

```cpp
#include <stlab/channel.hpp>

using process_state_scheduled =
  std::pair<process_state, std::chrono::system_clock::time_point>;

struct process_signature
{
    void await(T... val);

    U yield();

    process_state_scheduled state() const;

    void close();                          // optional

    void set_error(std::exception_ptr); // optional
};
```

This method must return the current state of the process. Typical return values are `await_forever` and `yield_immediate`. By explicit using the second part of the return type, one can set a possible timeout. Subsequent calls *without* an intermittent `await()`, `close()`, or `yield()` must return the same values. Otherwise the result is undefined.

Stateful Process Signature - close

```
#include <stlab/channel.hpp>

using process_state_scheduled =
  std::pair<process_state, std::chrono::system_clock::time_point>;

struct process_signature
{
    void await(T... val);

    U yield();

    process_state_scheduled state() const;

    void close();                       // optional

    void set_error(std::exception_ptr); // optional
};
```

The optional `close()` method is called on the process whenever the process state is `await_forever` and the incoming queue went dry. As well it is called when an exception is thrown while calling `await()` or `yield()` and no `set_error()` is available.

Stateful Process Signature - set_error

```cpp
#include <stlab/channel.hpp>

using process_state_scheduled =
  std::pair<process_state, std::chrono::system_clock::time_point>;

struct process_signature
{
    void await(T... val);

    U yield();

    process_state_scheduled state() const;

    void close();                          // optional

    void set_error(std::exception_ptr); // optional
};
```

The method `set_error()` is optional. It is called if either on calling `await()` or `yield()` an exception was thrown. The pointer of the caught exception is passed. In case that the process does not provide this method, `close()` is called instead of.

Channel - Stateful Process Example

```
#include <stlab/channel.hpp>
#include <stlab/default_executor.hpp>
#include <iostream>
using namespace stlab;

struct adder
{
};

int main() {
  sender<int> send;
  receiver<int> receiver;
  std::tie(send, receiver) = channel<int>(default_executor);

  auto calculator = receiver | adder{} |
    [](int x) { std::cout << x << '\n'; };

  receiver.set_ready();

  while (true) {
    int x;
    std::cin >> x;
    send(x);
  }
}
```

Channel - Stateful Process Example cont.

```
struct adder
{
  int _sum = 0;
  process_state_scheduled _state = await_forever;

  void await(int x) {
    _sum += x;
    if (x == 0) {
      _state = yield_immediate;
    }
  }

  int yield() {
    int result = _sum;
    _sum = 0;
    _state = await_forever;
    return result;
  }

  auto state() const { return _state; }
};

int main() {
  auto calculator = receiver | adder{} |
    [](int x) { std::cout << x << '\n'; };
  while (true) {
    int x;
    std::cin >> x;
    send(x);
  }
}
```

Exercise 3

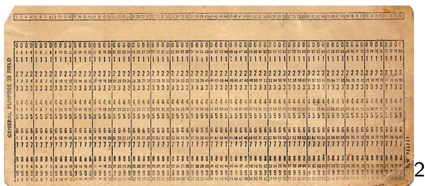

[2]

A process which inputs cards of eighty characters and outputs their text, tightly packed into lines of 125 characters each.

- Write one process - unpack - that collect 80 chars in a bunch and yields them one after the other
- Write one process - pack - that packs 125 chars and yields them.
- Concatenate unpack - pack as a process chain.
- In a next step write one process - filter - that drops all newlines from the stream
- Concatenate now unpack - filter - pack as a process chain.

[2] By Ventriloquist - Own work, CC BY-SA 3.0,

Process Analysis

- Are there performance or usability problems?
- Identify the overall critical part
- Disassemble this part into individual processes
- Chain the processes with futures or channels

Use case example I

Problem within our mammography application:

- Medical device shall open every case in < 1 s
- Loading of patient data and first images takes about 0.6 s
- Reading of additional data structures (CAD[3]reports) may take more than 0.4 s
- Direct access to any CAD report might be required
- If the user skips this case and advances to the next one, outstanding load operations should be cancelled or at least be ignored

[3] **C**omputer **A**ided **D**etection

Exercise 4

Improve the application that the UI is always responsible

- On Reset the reports are newly read
- If one presses 1 or 2 while the reset is running, the reports shall be displayed as soon as they become available.

Next Steps?

High level concurrency sessions at ACCU 2017:

- Thinking Outside the Synchronisation Quadrant by Kevlin Henney (Wed.)
- Coroutines in Python by Robert Smallshire (Thur.)
- Coroutines in C++ by Dominic Robinson (Fr.)
- Concurrency / Coroutines by Anthony Williams (Sat.)

Synchronization Motivation

Why do we have to synchronize?
Because we have to ensure sequential consistency.
What synchronization mechanism do you know?

- *Synchronization primitives (mutex, atomic, memory fence, ...)*
- *Guaranteed sequential access*

Synchronization with Mutex

```
1  template <typename K, typename V>
2  class registry
3  {
4    map<K, V> _data;
5    mutex     _guard;
6  public:
7    void insert(const K& key, const V& value) {
8      unique_lock<mutex> lock(_guard);
9      _data.insert(
10       make_pair("What is the answer?", 42)
11     );
12   }
13
14   V operator[](const K& key) {
15     unique_lock<mutex> lock(_guard);
16     return _data[key];
17   }
18 };
19
20 int main() {
21   registry<string, int> my_registry;
22   auto work = [&] { my_registry.insert("What is the answer?", 42); };
23   auto f1 = async(launch::async, work);
24   auto f2 = async(launch::async, work);
25   f1.get(); f2.get();
26   cout << "What is the answer? " << my_registry["What is the answer?"]
       << '\n';
27 }
```

Where are the problems in the code?

Synchronization with Mutex

Mutex - What would be a better name for it?

Bottleneck![4]

[4] Kevlin Henney, NDC London 2017

Synchronization without Mutex

How can the code be transformed into something without a mutex in the client code?
What is needed to perform that transformation? Which tools do we have in our tool box?

Synchronization without Mutex

```
1  template <typename K, typename V>
2  class registry
3  {
4    std::shared_ptr<map<K, V>>  _data;
5    serial_queue                _queue;
6  public:
7    void insert(K key, V val) {
8      _queue.async([_d = _data,
9                    _key = std::move(key),
10                   _val = std::move(val)] {
11         d->emplace(std::move(key), std::move(val));
12     });
13   }
14
15   future<V> operator[](K key) {
16     return _queue.async([_d = _data,
17                          _key = std::move(key)] {
18       return _d->at(key);
19     });
20   }
21 };
```

Synchronization Epilogue

So we try to avoid mutexes wherever it is possible.

Synchronization Epilogue

All computer wait at the same speed

Use case example II

Image Preparation Pipeline

- A medial device shall display multi-frame image data sets
- Each incoming data set is JPEG 2000 compressed
- The slices must be decompressed and then compressed in FELICS[5] format for fast decompression and display
- Reading and writing to disk takes a reasonable amount of time

[5] Special compression algorithm for 16bit grayscale images

Reference

- Concurrency library
  https://github.com/stlab/libraries
- Documentation http://www.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. Roscoehttp://www.cs.ox.ac.uk/people/bill.roscoe/publications/68b.pdf

Further reading I

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

Further reading II

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 coming soon)

Further viewing

- 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`
- Thinking Outside the Synchronization Quadrant by Kevlin Henney: `https://vimeo.com/205806162`

Acknowledgement

- My family, who gave me the freedom to develop over months the library, prepare this tutorial and let me travel to the ACCU.
- 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, that released me from work during the ACCU.

Contact

- Mail: felix@petriconi.net
- Web: `https://petriconi.net`
- Twitter: @FelixPetriconi