


Leading edge



A roundup of C++20 and C++23

Kris van Rens




(source: Wikipedia, CC-BY-SA-3.0)

Frederick P. Brooks Jr.

April 19, 1931 – November 17, 2022

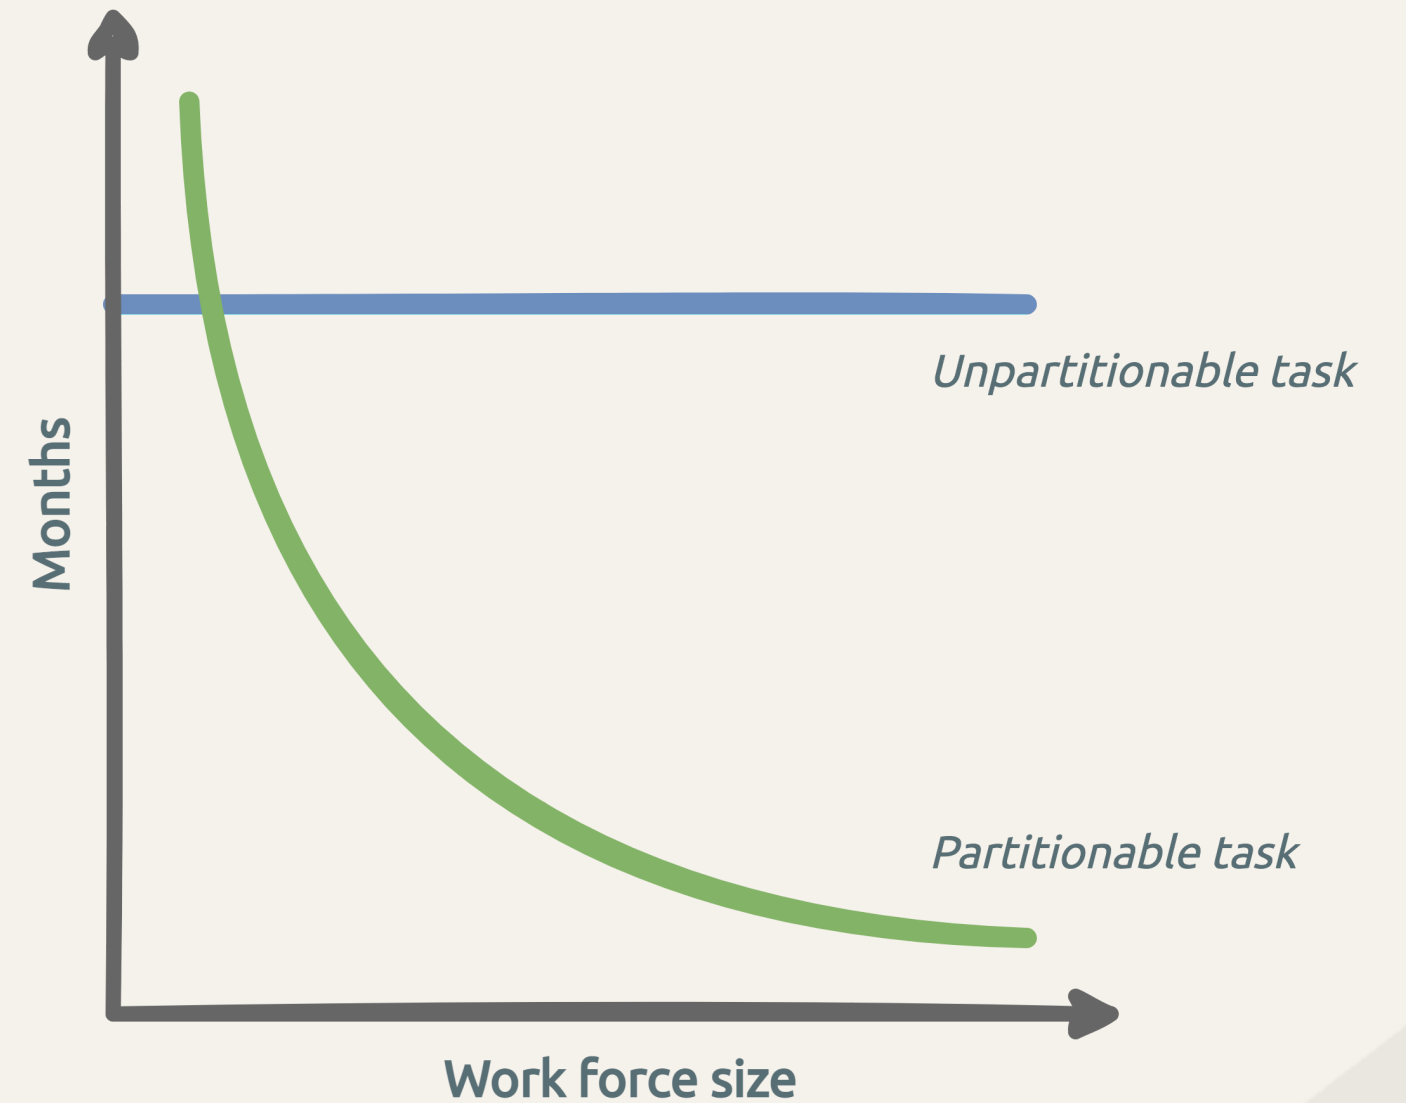
- Led the development of **IBM System/360**
- Teaching assistant for **Ken Iverson (of APL)**
- Coined term “**Computer Architecture**”
- Invented “**Brooks’s Law**”
- Moved **byte size from 6 to 8 bits**
- **IEEE John von Neumann Medal (1993)**
- **ACM Turing Award (1999)**



(source: Wikipedia, fair use)


Brooks's Law:

“Adding manpower to a late project makes it later.”



Brooks's Law:

“Adding manpower to a late project makes it later.”



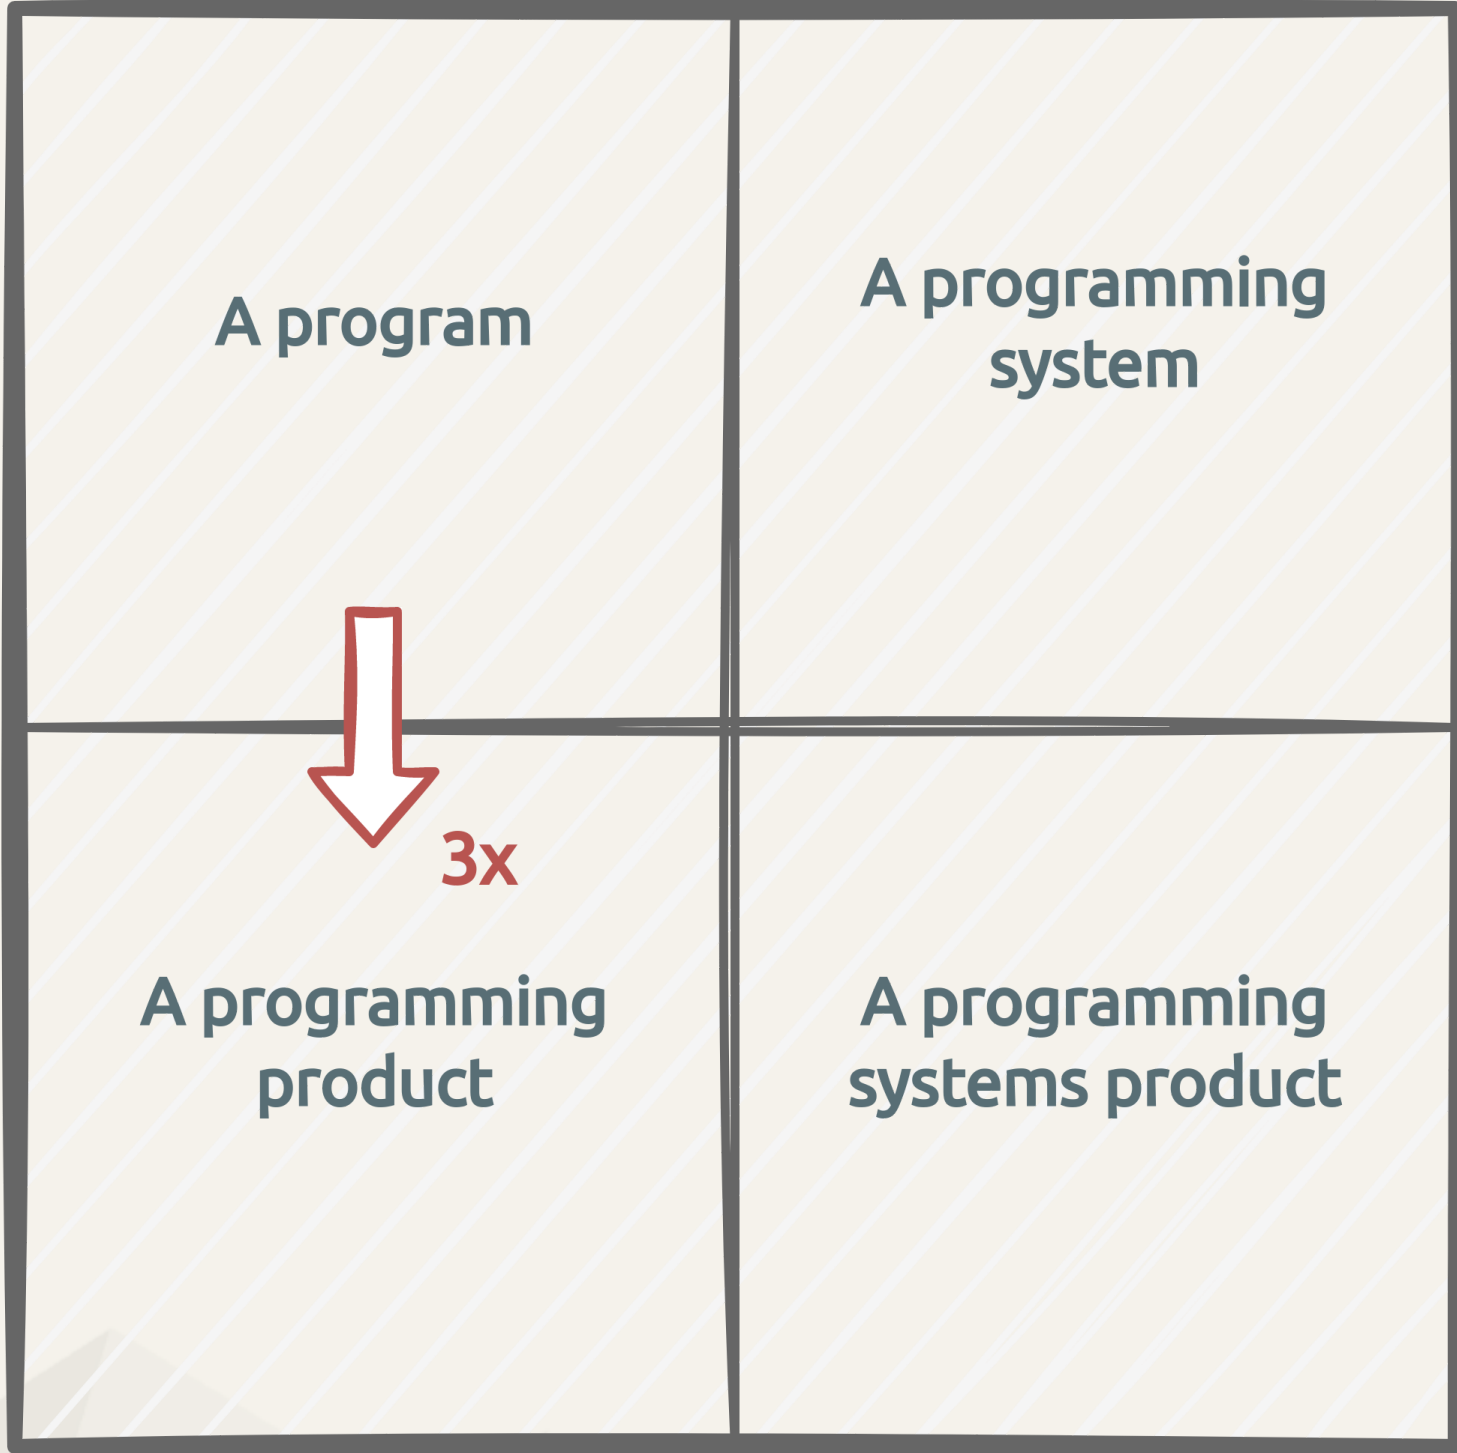
A program

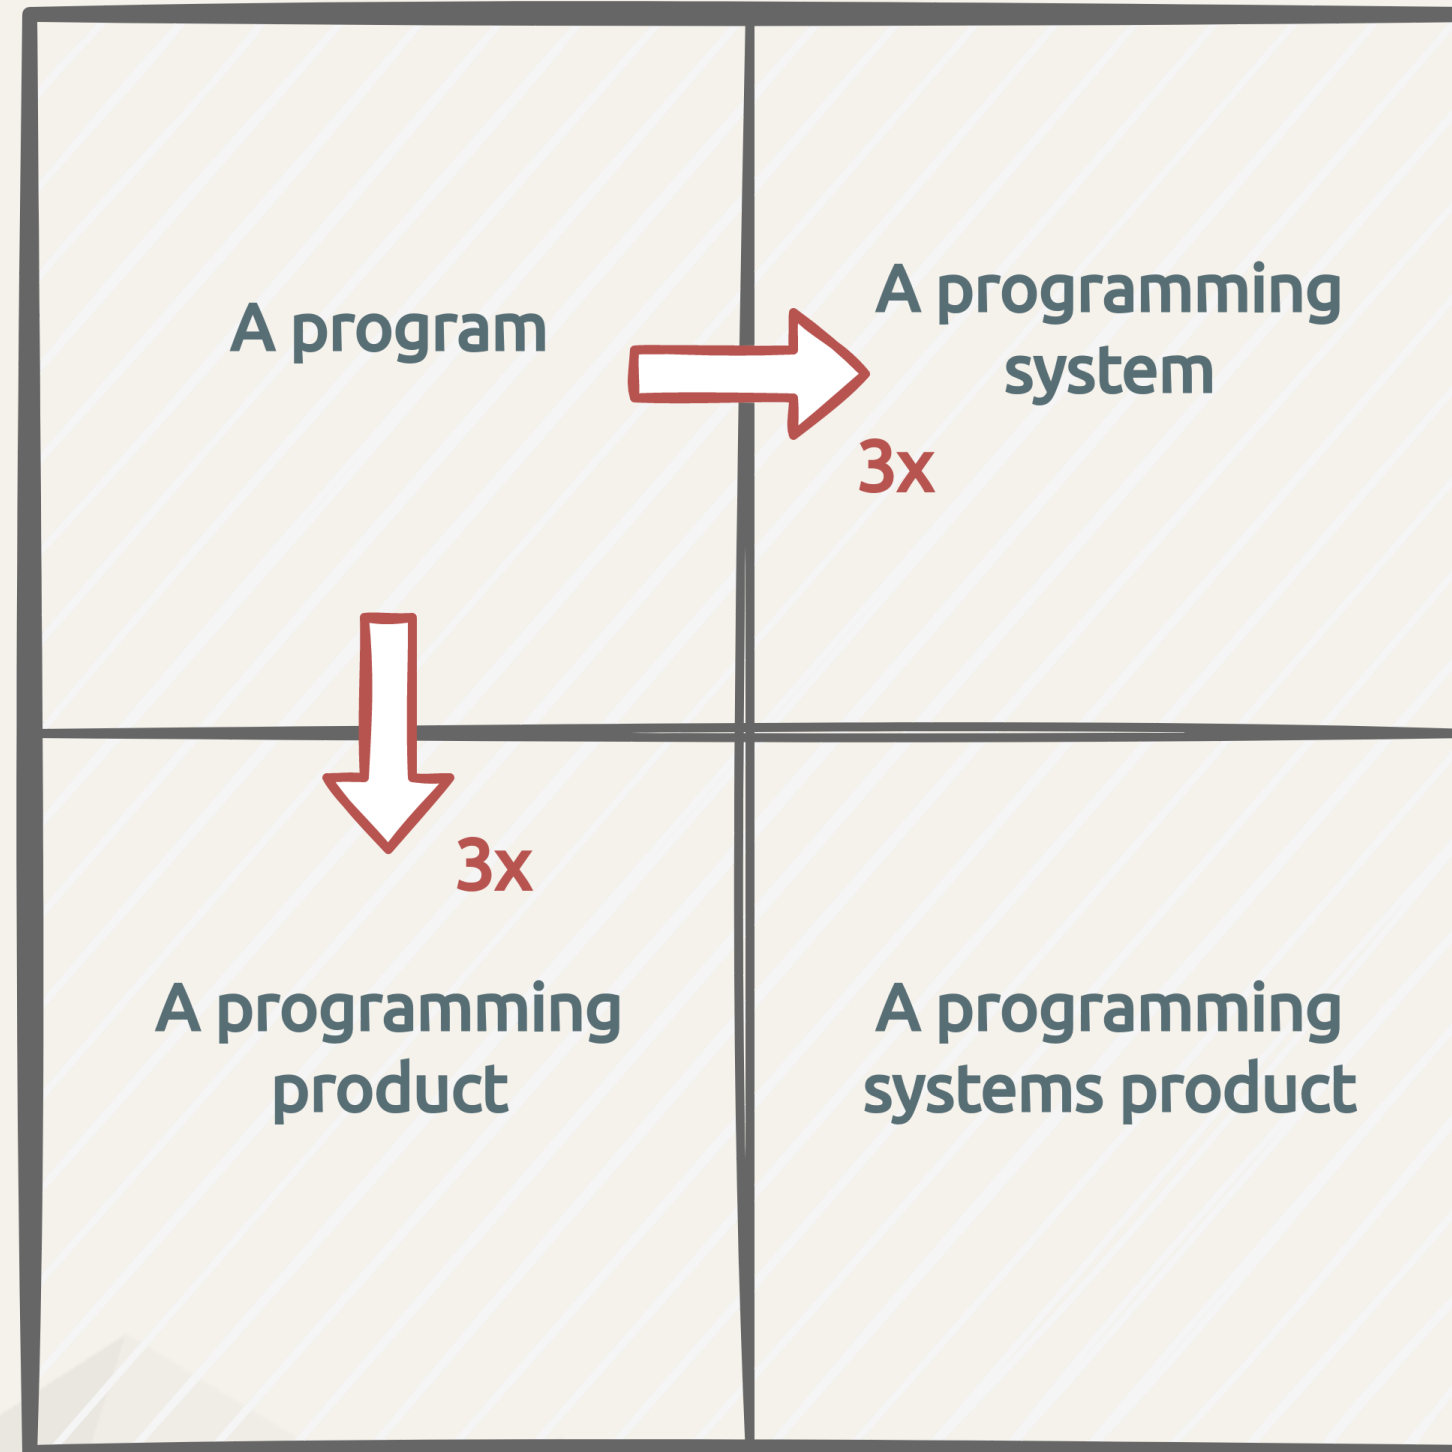
**A programming
system**

**A programming
product**

**A programming
systems product**

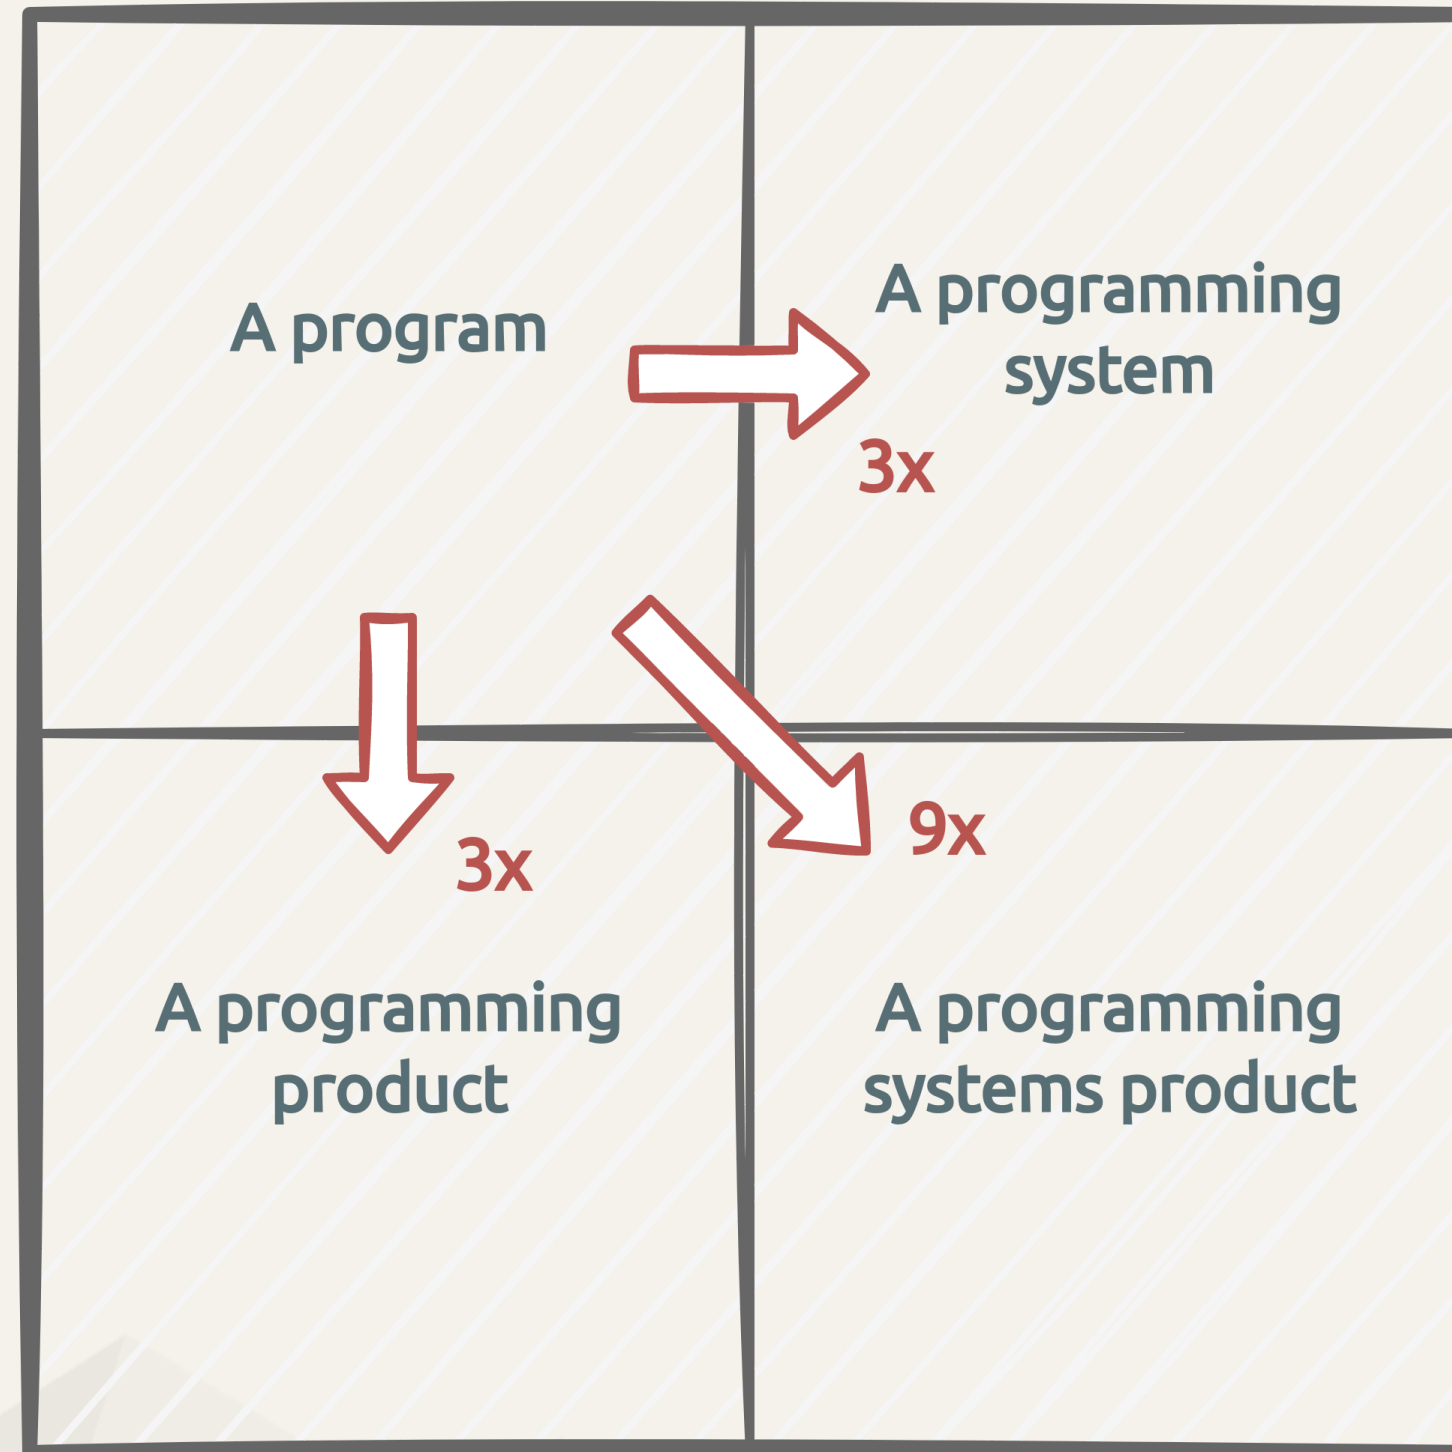
*+ generalization +
tests + docs +
maintenance*





*+ interfaces +
system integration*

*+ generalization +
tests + docs +
maintenance*



*+ interfaces +
system integration*

*+ generalization +
tests + docs +
maintenance*

What's ahead?

- Leading Edge C++
- C++20 by example
- An outlook to C++23

A little bit about me




kris@vanrens.org

Leading Edge C++



How C++ is cooked up

- C++ is an **ISO-standardized language** (language + library),
- ISO WG21 comprises of **23 study groups** and a total over 150 members.



C++20 – not just any update

- **Concepts**
- **Modules**
- **Coroutines**
- **Ranges library**
- Three-way comparison '**operator <=>**'
- Array views with **std::span**
- Advanced text formatting library `<format>`
- New thread support primitives and `jthread`
- Calendars and time zones in `<chrono>`
- ...

Compiler support for C++20/23

There are **only a few compilers** supporting C++20 in full

C++20

	Decent	Full
GCC	10	13
Clang	10	16
MSVC	19.25	19.30

(MS VS2019 ships MSVC 19.29 in v16.11.2)

C++23

	Decent	Full
GCC	13	-
Clang	16	-
MSVC	19.35*	-

(: mostly just the standard library)*

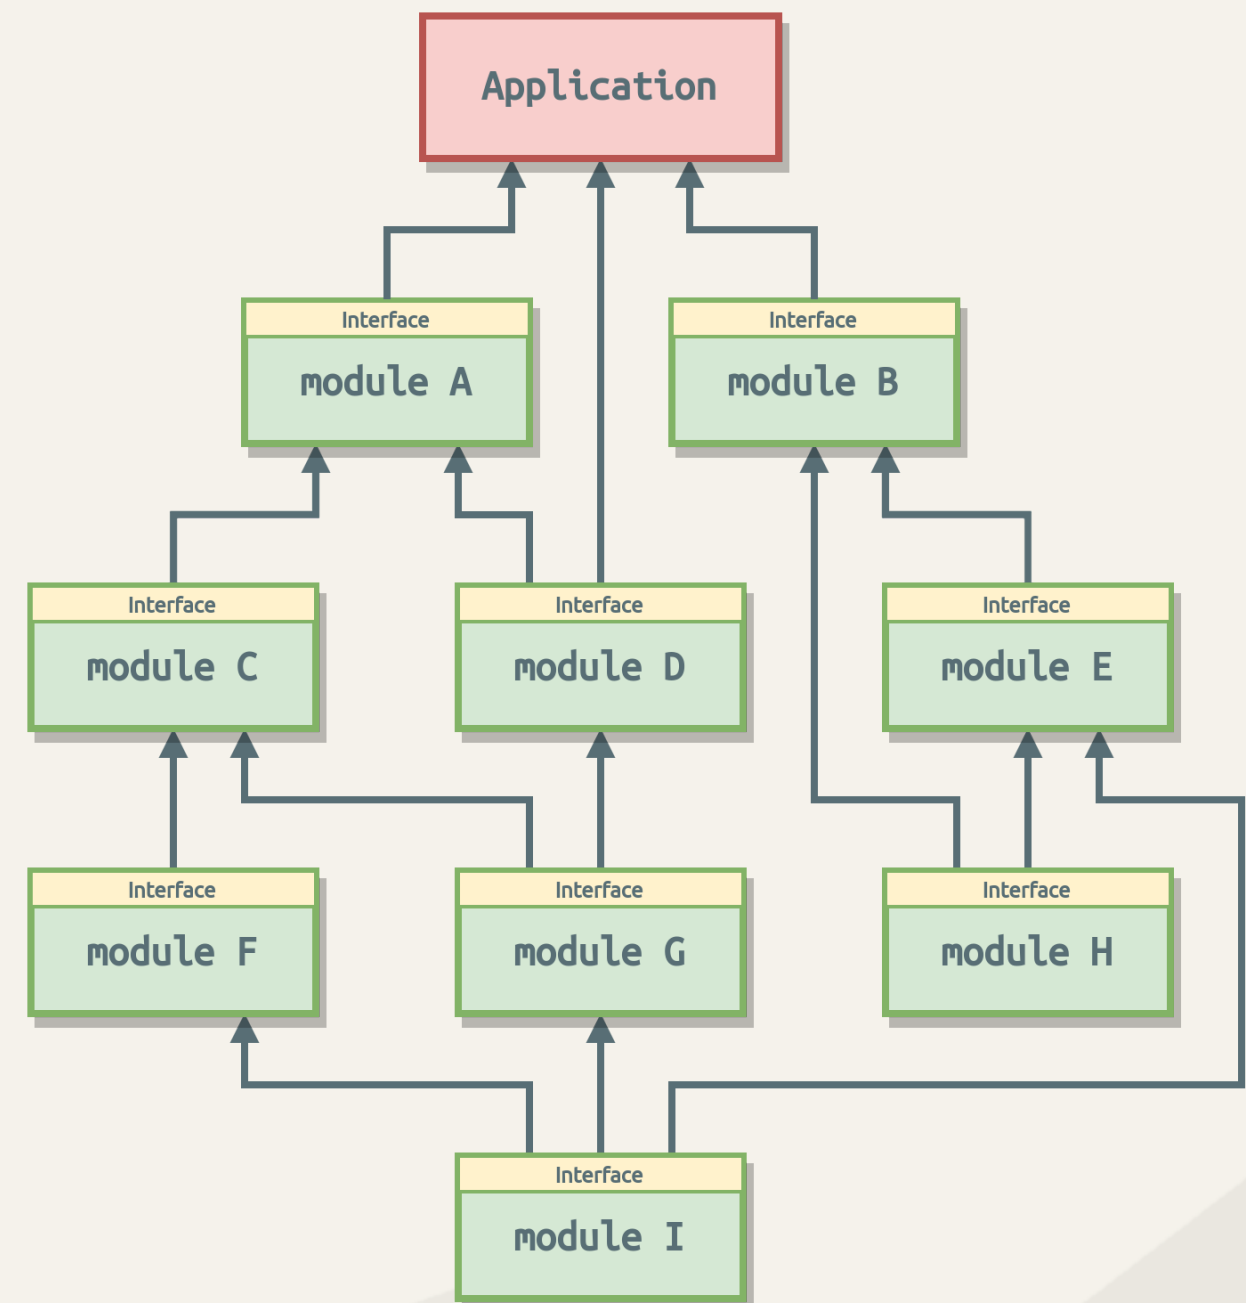
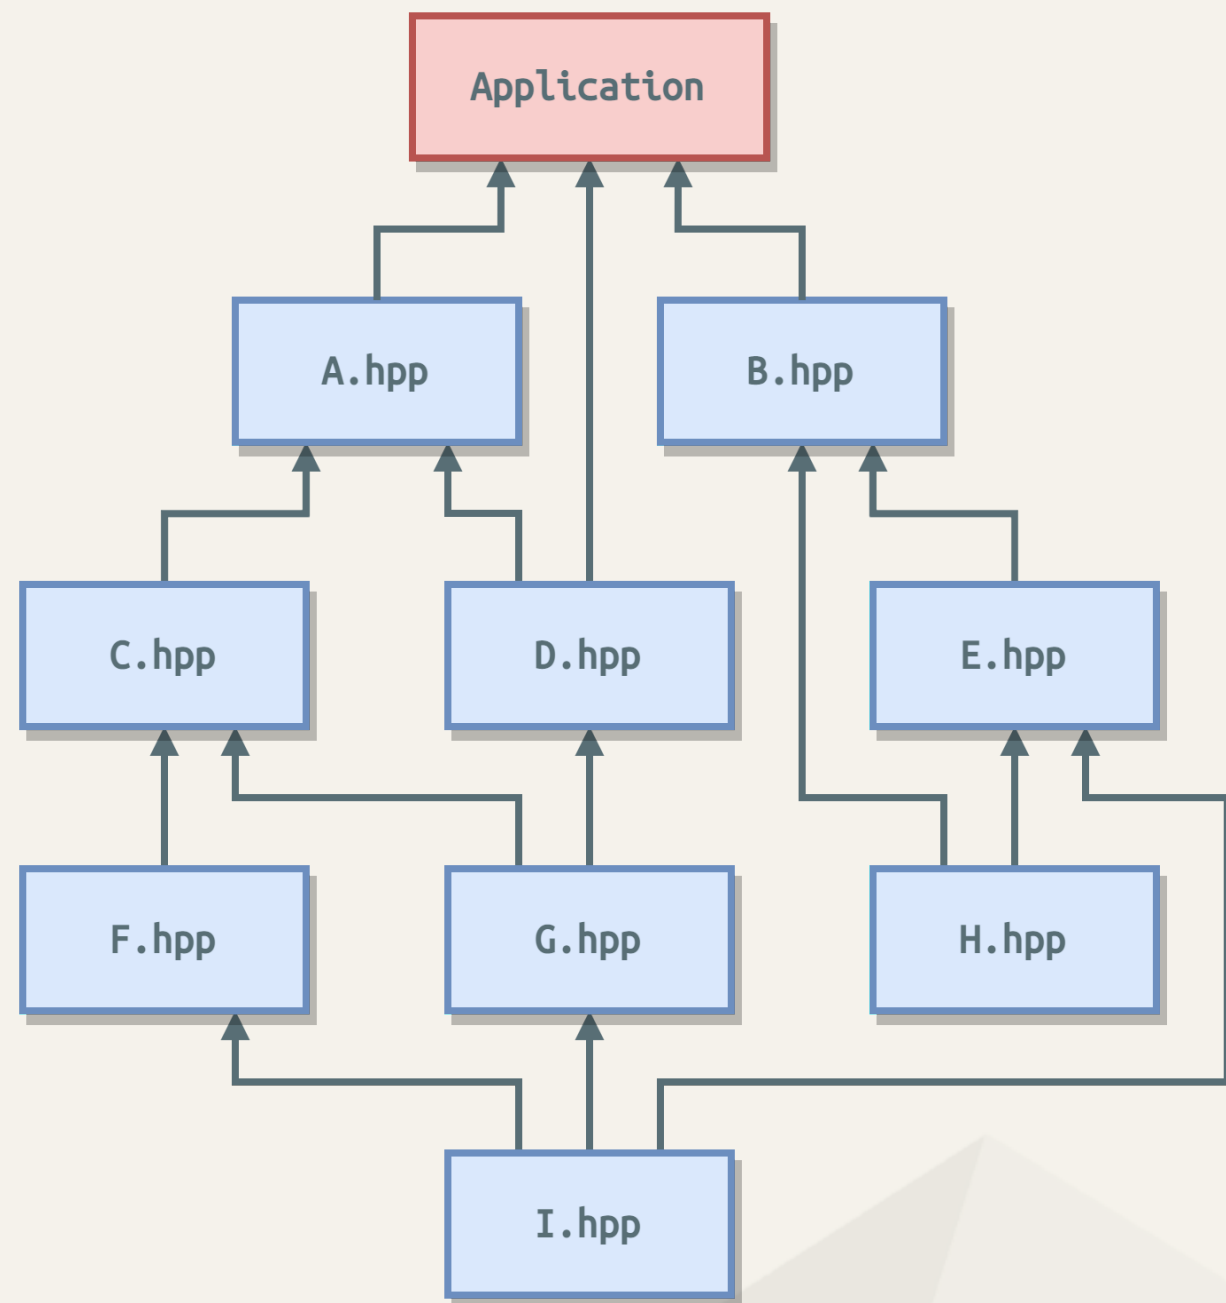
(MS VS2022 ships MSVC 19.33 in v17.4.3)

C++20 by example

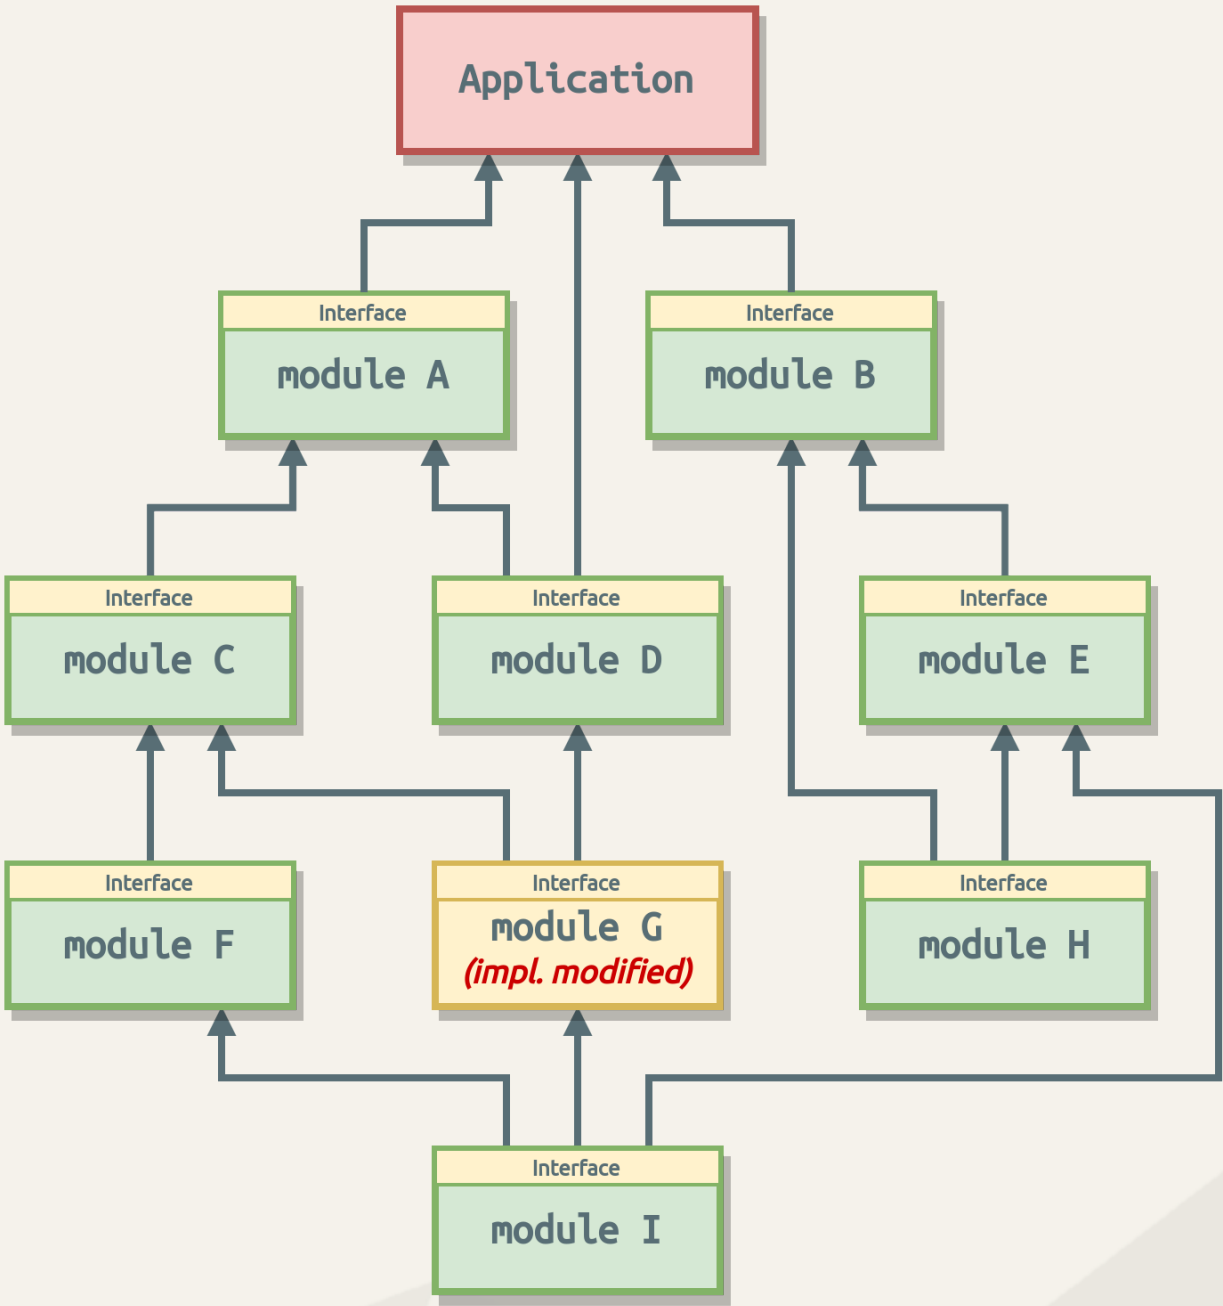
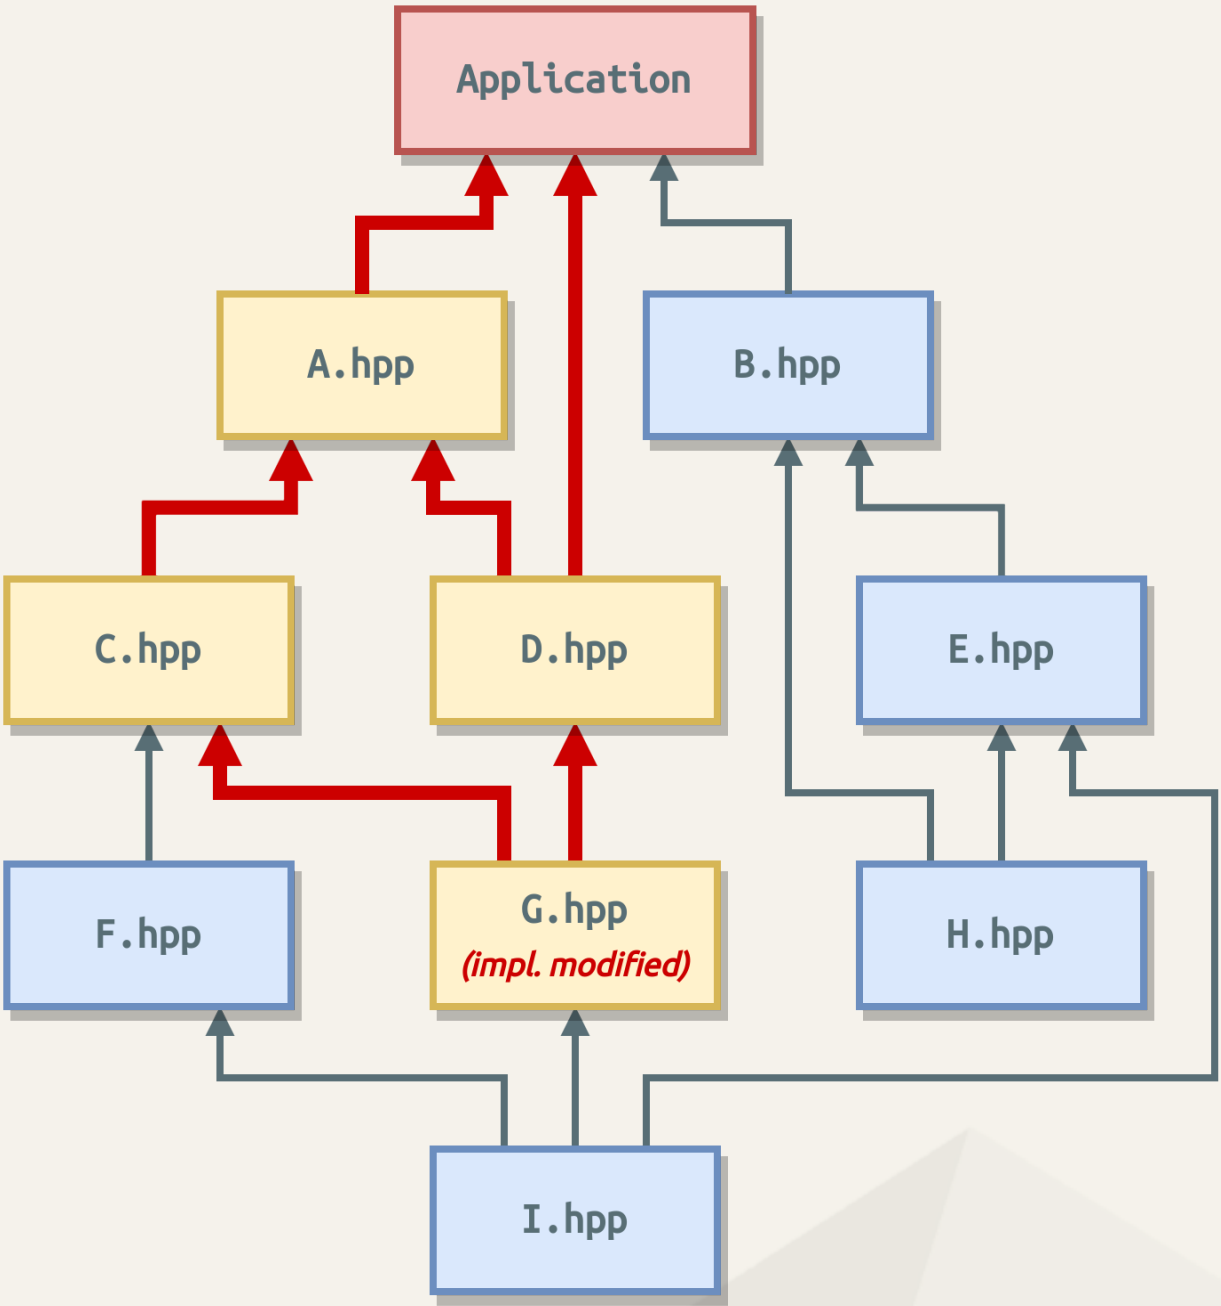


Modules

The problem modules solve



The problem modules solve



The simplest ever module

fraction.cppm (or .ixx):

```
1 export module fraction;
2
3 template<typename Type>
4 concept arithmetic = std::integral<Type> || std::floating_point<Type>;
5
6 template<arithmetic Type>
7 export class Fraction final {
8 public:
9     // ...
10 };
11
12 template<typename Type>
13 export std::ostream& operator<<(std::ostream& os, const Fraction<Type>& f) {
14     os << "[" << f.numerator() << "/" << f.denominator() << "]";
15     return os;
16 }
```



main.cpp:

```
1 import fraction;
2 import std; // Proposed in C++23.
3
4 int main() {
5     Fraction f{22, 7};
6     std::cout << "Fraction:" << f << '\n';
7 }
```



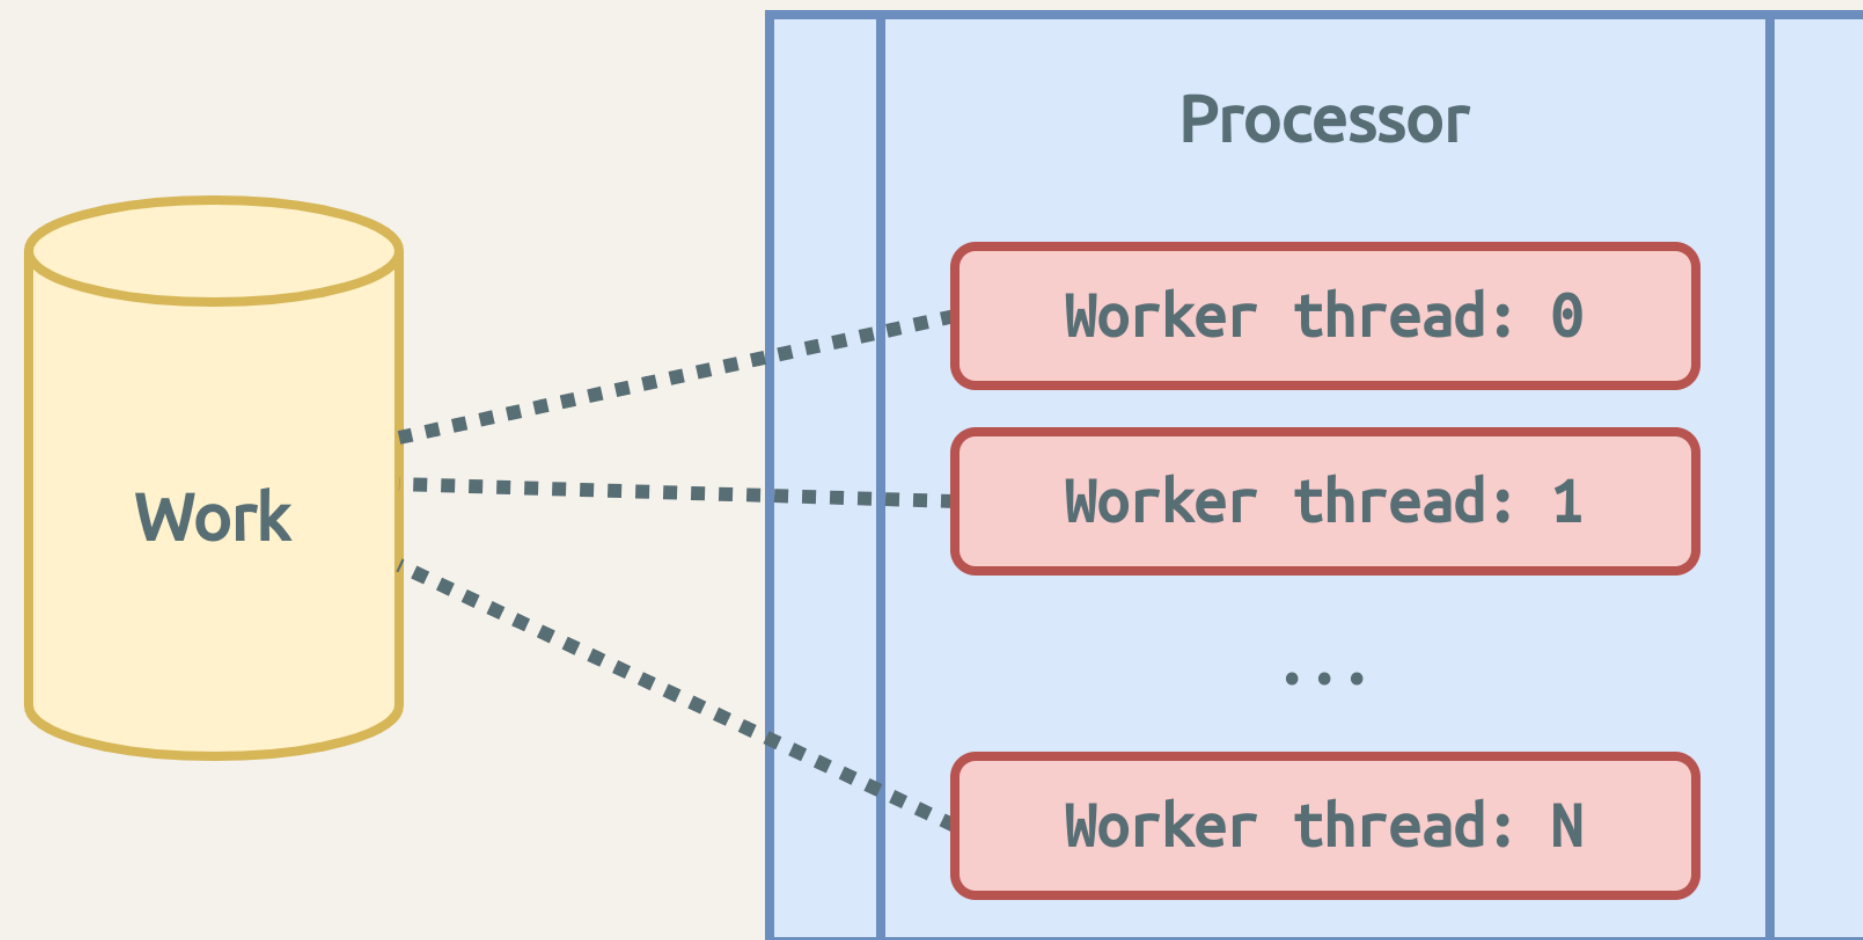
Not all build systems integrate support for **modules** still.. 😞

(there is no consensus on the extension .ixx/.cppm AFAIK)

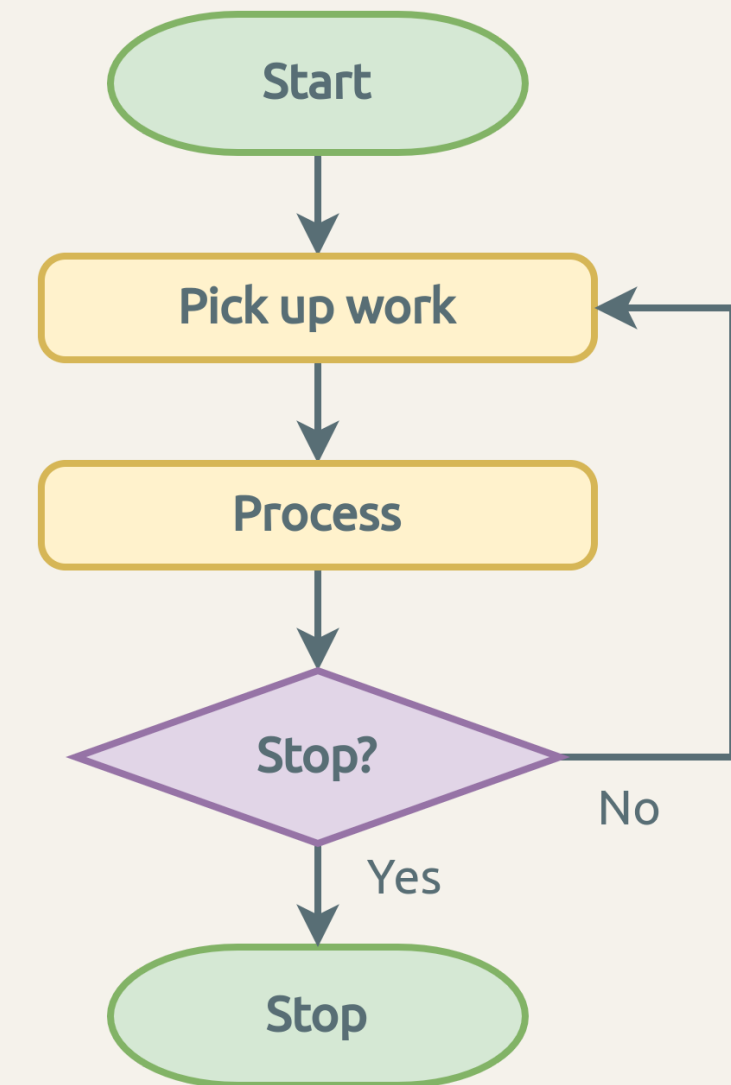
Concurrency

(with threads)

Example overview



Worker flow



Version 1: Setup + polling for work

```
1 template<std::size_t NumWorkers>
2 class Processor {
3     std::vector<std::jthread> workers_;
4
5     void worker(std::stop_token stop_token);
6
7 public:
8     Processor() {
9         for (std::size_t i = 0; i < NumWorkers; i++) {
10             workers_.emplace_back(
11                 std::bind_front(&Processor::worker, this));
12         }
13     }
14
15     ~Processor() {
16         for (auto& worker : workers_) {
17             worker.request_stop();
18         }
19     }
20 };
```



```
1 template<std::size_t NumWorkers>
2 void Processor<NumWorkers>::worker(std::stop_token stop_token) {
3     while (!stop_token.stop_requested()) {
4         //
5         // ..poll for work and process it..
6         //
7     }
8 }
```



```
1 int main() {
2     Processor<4> proc;
3
4     // ..exit when work is done..
5 }
```



Version 2: Properly waiting for work

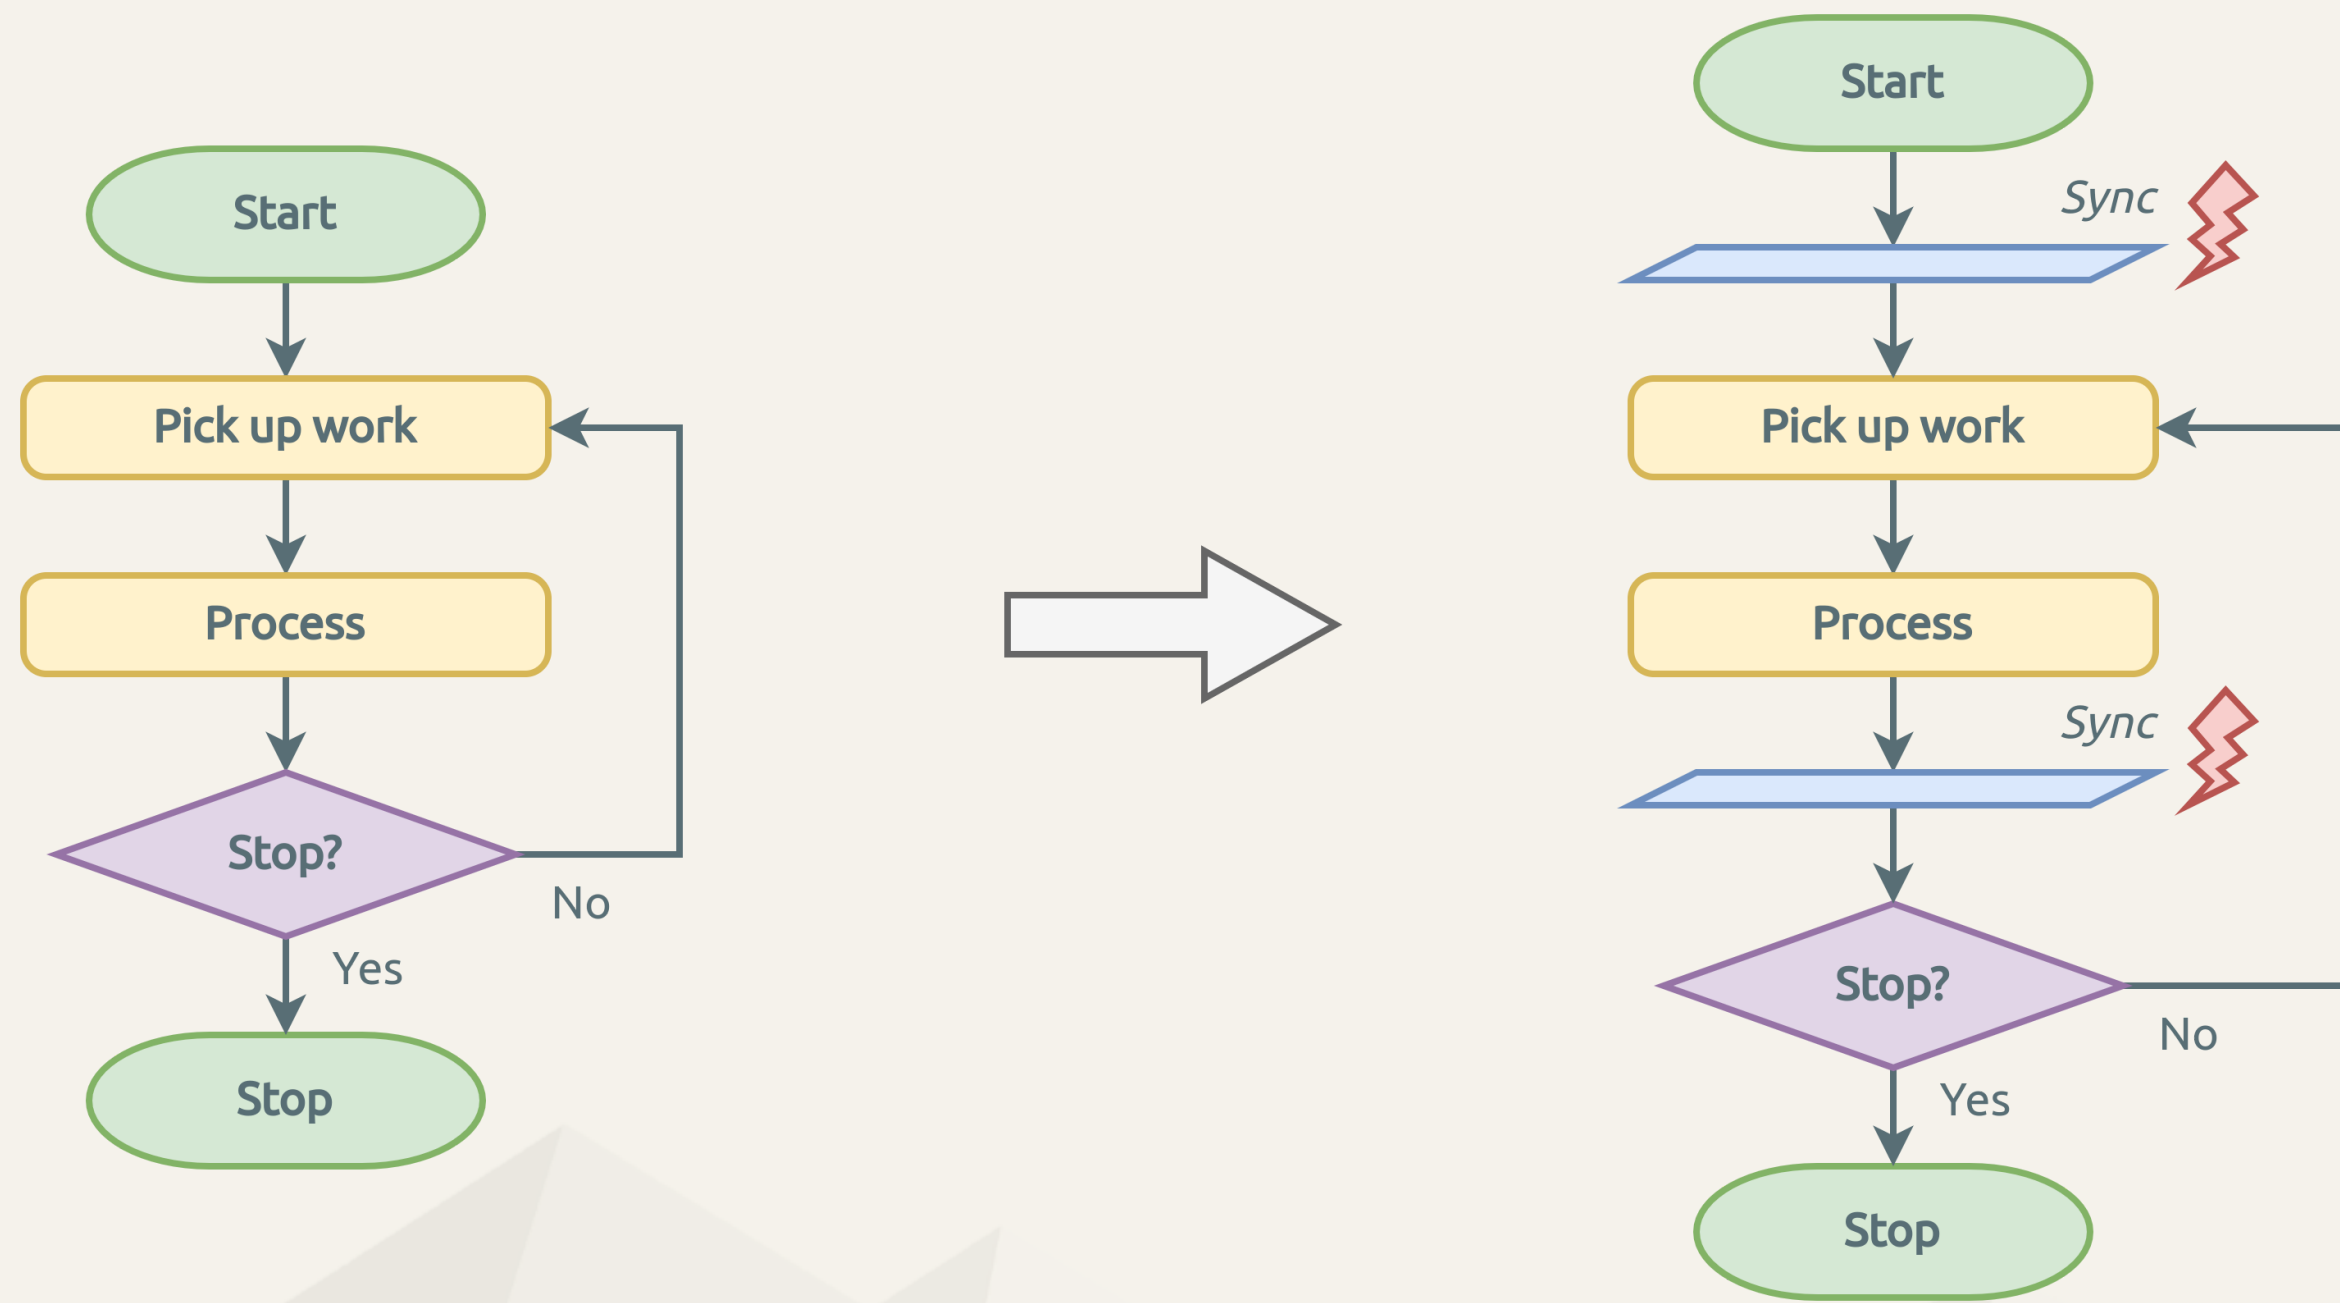
```
1 template<std::size_t NumWorkers>
2 class Processor {
3     std::vector<std::jthread> workers_;
4     std::mutex work_mutex_;
5     std::condition_variable_any work_cv_;
6
7     void worker(std::stop_token stop_token);
8
9 public:
10    Processor();
11
12    ~Processor() {
13        for (auto& worker : workers_) {
14            worker.request_stop();
15        }
16    }
17 };
```



```
1 template<std::size_t NumWorkers>
2 void Processor<NumWorkers>::worker(std::stop_token stop_token) {
3     while (!stop_token.stop_requested()) {
4         if (auto task{work_queue.get()}; task) {
5             //
6             // ..process work..
7             //
8         } else {
9             std::unique_lock lock{work_mutex_};
10            work_cv_.wait(lock, stop_token, [&]{ return !work_queue.empty(); });
11        }
12    }
13 }
```



Add worker synchronization points



Version 3: Add thread start synchronization point

```
1 template<std::size_t NumWorkers>
2 class Processor {
3     std::vector<std::jthread> workers_;
4     std::latch workers_start_{NumWorkers};
5     std::mutex work_mutex_;
6     std::condition_variable_any work_cv_;
7
8     void worker(std::stop_token stop_token);
9
10 public:
11     Processor();
12     ~Processor();
13 };
```

A `std::latch` is a single-use synchronization point

```
1 template<std::size_t NumWorkers>
2 void Processor<NumWorkers>::worker(std::stop_token stop_token) {
3
4     // All threads will synchronize here:
5     workers_start_.arrive_and_wait();
6
7     while (!stop_token.stop_requested()) {
8         if (auto task{work_queue.get()}; task) {
9             //
10            // ..process work..
11            //
12        } else {
13            std::unique_lock lock{work_mutex_};
14            work_cv_.wait(lock, stop_token, [&]{ return !work_queue.empty(); });
15        }
16    }
17 }
```


Version 4: Add post-work synchronization point

```
1 template<std::size_t NumWorkers>
2 class Processor {
3     std::vector<std::jthread> workers_;
4     std::latch workers_start_{NumWorkers};
5     std::barrier<> workers_sync_{NumWorkers};
6     std::mutex work_mutex_;
7     std::condition_variable_any work_cv_;
8
9     void worker(std::stop_token stop_token);
10
11 public:
12     Processor();
13     ~Processor();
14 };
```

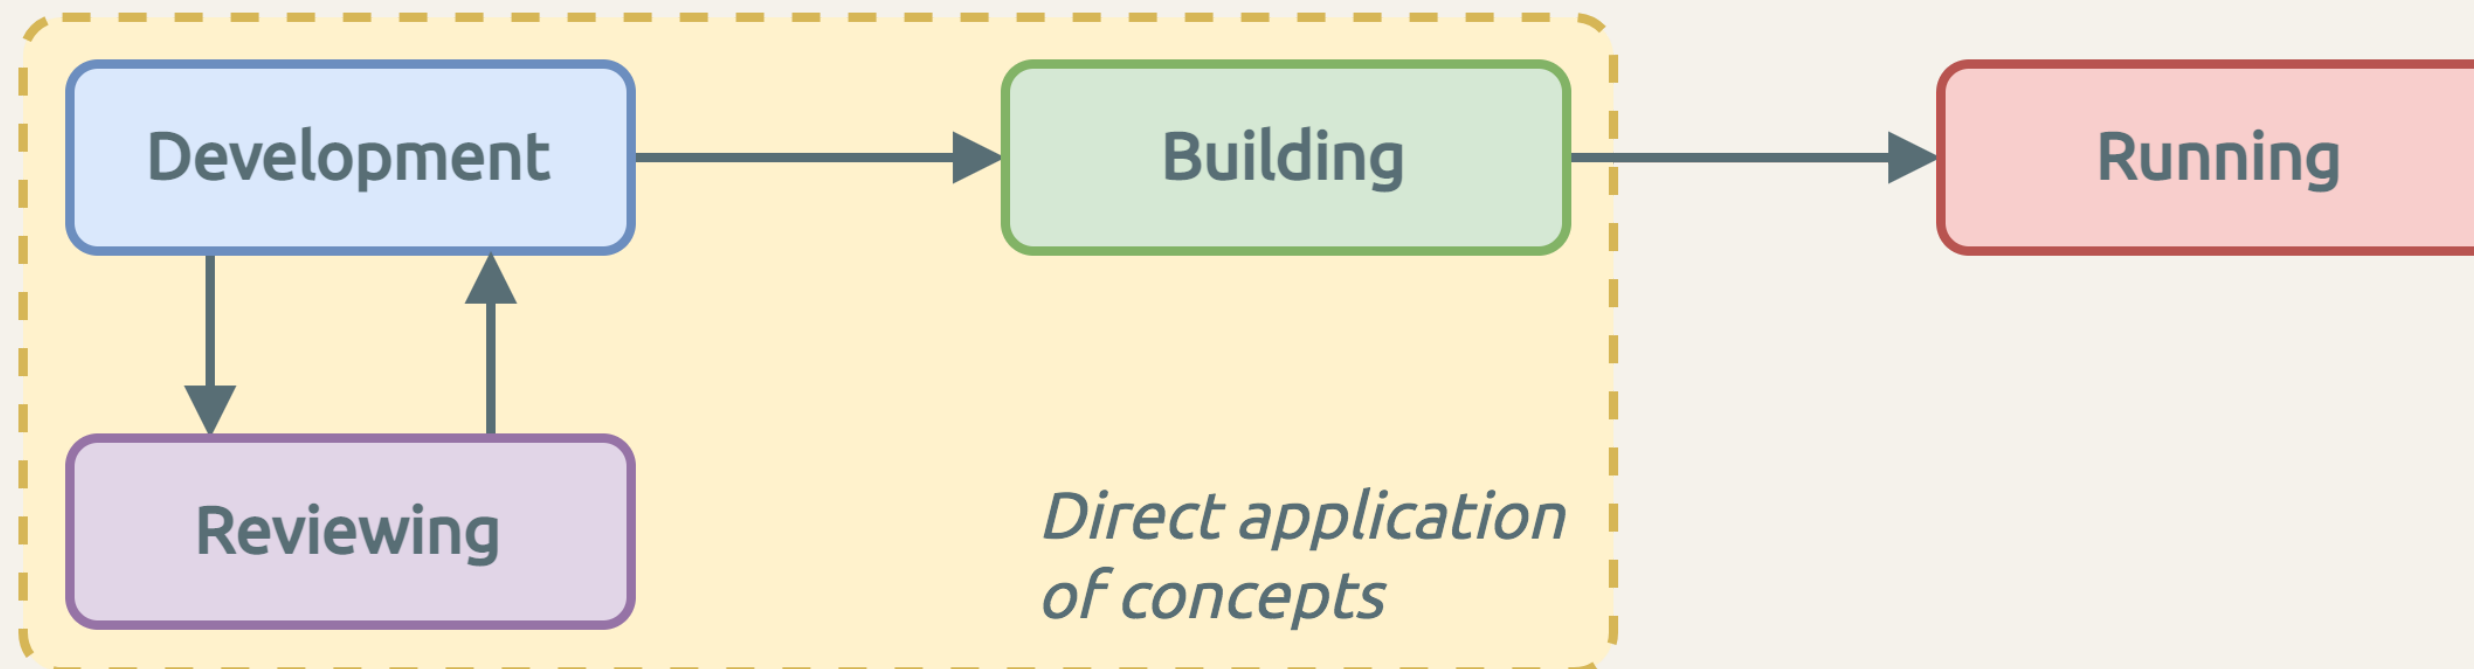
A `std::barrier` is a multi-use synchronization point

```
1 template<std::size_t NumWorkers>
2 void Processor<NumWorkers>::worker(std::stop_token stop_token) {
3     workers_start_.arrive_and_wait();
4
5     while (!stop_token.stop_requested()) {
6         if (auto task{work_queue.get()}; task) {
7             //
8             // ..process work..
9             //
10
11             workers_sync_.arrive_and_wait();
12
13             //
14             // ..post-process work..
15             //
16         } else {
17             std::unique_lock lock{work_mutex_};
18             work_cv_.wait(lock, stop_token, [&]{ return !work_queue.empty(); });
19         }
20     }
21 }
```

Compile-time validation with **concepts**

What **concepts** are about..

- Concepts are **named constraints**, predicates evaluated at compile-time,
- Concepts are **part of template interfaces** to improve *equational reasoning*,
- Concepts can be used to impose **syntactic and semantic constraints**.



Imposing type constraints

Use case: constrain class template type parameter.

```
SensorReading<float> r1; // OK.  
SensorReading<int> r2; // Compiler error.
```

Pre-C++20

```
1 #include <type_traits>  
2  
3 template<typename T>  
4 class SensorReading {  
5     static_assert(std::is_floating_point_v<T>,  
6                 "T must be a floating point type");  
7 public:  
8     // ...  
9 };
```



C++20

```
1 #include <concepts>  
2  
3 template<std::floating_point T>  
4 class SensorReading {  
5 public:  
6     // ...  
7 };
```



Constraining (member-)functions

Use case: conditionally enable member function.

```
SensorReading<double> r1; // Has 'guess_snr'.  
SensorReading<float> r2; // Has no 'guess_snr'.  
  
double snr1 = r1.guess_snr(); // OK.  
double snr2 = r2.guess_snr(); // Compiler error!
```

Pre-C++20

```
1 #include <type_traits>  
2  
3 template<typename T>  
4 class SensorReading {  
5 public:  
6     template<typename U = T,  
7             typename = std::enable_if_t<std::is_same_v<double, U>>>  
8     double guess_snr() const;  
9 };
```



C++20

```
1 #include <concepts>  
2  
3 template<std::floating_point T>  
4 class SensorReading {  
5 public:  
6     double guess_snr() const requires std::same_as<double, T>;  
7 };
```



Defining custom concepts

Example: concept to test for type equivalence.

```
1 class Sensor {
2     // ...
3 };
4
5 class SpecificSensor : public Sensor {
6     // ...
7 };
8
9 class MeasurementSystem {
10 public:
11     operator Sensor() const { // Type conversion function.
12         // ...
13     }
14 };
```



```
1 #include <concepts>
2
3 template<typename T>
4 concept sensor_like = std::same_as<Sensor, T>
5                     || std::derived_from<T, Sensor>
6                     || std::convertible_to<T, Sensor>;
7
8 template<sensor_like S>
9 void process(const S& s);
```



```
1 Sensor          s1; // std::same_as      --> true.
2 SpecificSensor  s2; // std::derived_from --> true.
3 MeasurementSystem s3; // std::convertible_to --> true.
4
5 process(s1);
6 process(s2);
7 process(s3);
```



All types can be used as a 'Sensor'

requires expressions

Requires expressions can be used to **easily define powerful constraints.**

Testing for get()

```
1 template<typename T>
2 concept has_get = requires (T t) { t.get(); };
3 // ~~~~~ Requires expression.
4
5 template<typename Input> requires has_get<Input>
6 void process(const Input& i) {
7     const auto value{i.get()};
8     // ...
9 }
```

```
1 class Sensor {
2 public:
3     int get() const;
4 };
5
6 class SomeType {
7 public:
8     // No 'get' function here..
9 };
```

```
1 Sensor x1;
2 SomeType x2;
3
4 process(x1); // OK.
5 process(x2); // Compiler error.
```

requires expressions

Example: arithmetic key/value mapping type definition.

```
1 template<typename T>
2 concept key_value_store = requires (T x, int i) {
3     { x.get()      };
4     { x.get()      } noexcept;
5     { x.get()      } -> std::same_as<std::pair<int, std::string>>;
6     { x.get_val(i) } -> std::same_as<std::string>;
7     { x + x        } -> std::same_as<T>;
8 } && std::equality_comparable<T>;
9
10 template<key_value_store Store>
11 void process(const Store& s) {
12     const auto& [key, value] = s.get();
13
14     // ...
15 }
```



```
1 class Mapping {
2 public:
3     using KeyValue = std::pair<int, std::string>;
4
5     [[nodiscard]] KeyValue get() const noexcept;
6     [[nodiscard]] std::string get_val(int index) const;
7
8     Mapping operator+(const Mapping&);
9
10    bool operator==(const Mapping&) const = default;
11 };
```



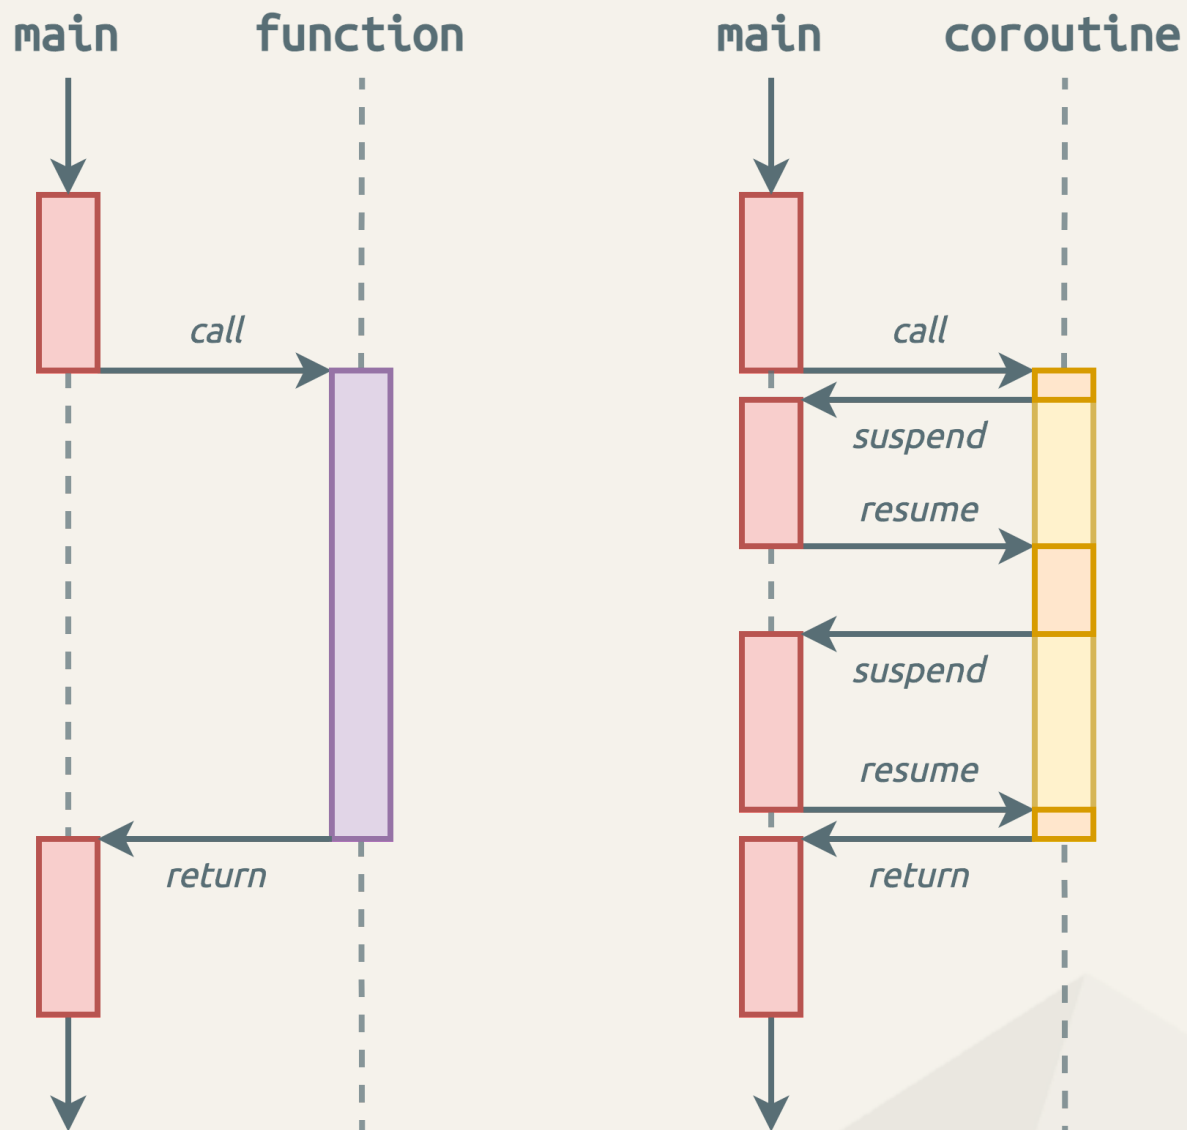
```
1 Mapping m1, m2;
2
3 process(m1 + m2);
```



Coroutines

What are coroutines?

Coroutines are **functions that can suspend execution to be resumed later**



Simple example

```
async_worker<void> read() {  
    float value = co_await async_read();  
    // ...  
}
```

```
std::future<float> async_read() {  
    // ..read 'value' from some source..  
    co_return value;  
}
```

async_read requires implementation of `std::coroutine_traits`:

```
template<typename T, typename... Args>  
struct std::coroutine_traits<std::future<T>, Args...> {  
    // ...  
};
```

Basic implementation of a coroutine return type based on `std::future`:

```
1  template<typename T, typename... Args>
2  struct std::coroutine_traits<std::future<T>, Args...> {
3      struct promise_type : std::promise<T> {
4          promise_type() = default;
5
6          std::future<T> get_return_object() noexcept {
7              return this->get_future();
8          }
9
10         void unhandled_exception() noexcept {
11             this->set_exception(std::current_exception());
12         }
13
14         static std::suspend_never initial_suspend() noexcept {
15             return {};
16         }
17
18         static std::suspend_never final_suspend() noexcept {
19             return {};
20         }
21
22         template<std::convertible_to<T> U>
23         void return_value(U&& value) noexcept(std::is_nothrow_constructible_v<T, decltype(std::forward<U>(value))>) {
24             this->set_value(std::forward<U>(value));
25         }
26     };
27 };
```



Associated type `promise_type`

Each coroutine has a `promise_type`:

```
1 struct promise_type {
2
3     ***** Mandatory interface *****
4     promise_type() = default; // Option (A) for c'tor.
5     promise_type(args...);    // Option (B) for c'tor.
6
7     auto get_return_object();
8     void unhandled_exception();
9     auto initial_suspend();
10    auto final_suspend() noexcept;
11
12    ***** Optional interface *****
13    void return_value(auto expr); // Option (A) for result.
14    void return_void();          // Option (B) for result.
15
16    auto yield_value(auto expr);
17    auto get_return_object_on_allocation_failure();
18    auto await_transform(auto expr);
19    void* operator new(std::size_t size, args...);
20    void operator delete(void* ptr, std::size_t size);
21 };
```



Customizable interfaces:

- Implement asynchronous tasks,
- Implement *awaitable types*,
- Implement *generators*,
- ...

No standard library support yet 😞

Range-enabled coroutine return type generator:

(Code @  Compiler Explorer)

```
struct coro_deleter {
    template<typename Promise>
    void operator()(Promise* promise) const noexcept {
        if (auto handle = std::coroutine_handle<Promise>::from_promise(*promise); handle) {
            handle.destroy();
        }
    }
};
```

```
template<typename T>
using promise_ptr = std::unique_ptr<T, coro_deleter>;
```

```
template<typename T>
class [[nodiscard]] generator {
public:
    using value_type = std::remove_reference_t<T>;
    using reference = std::conditional_t<std::is_reference_v<T>, T, const value_type&>;
    using pointer = const value_type*;
```

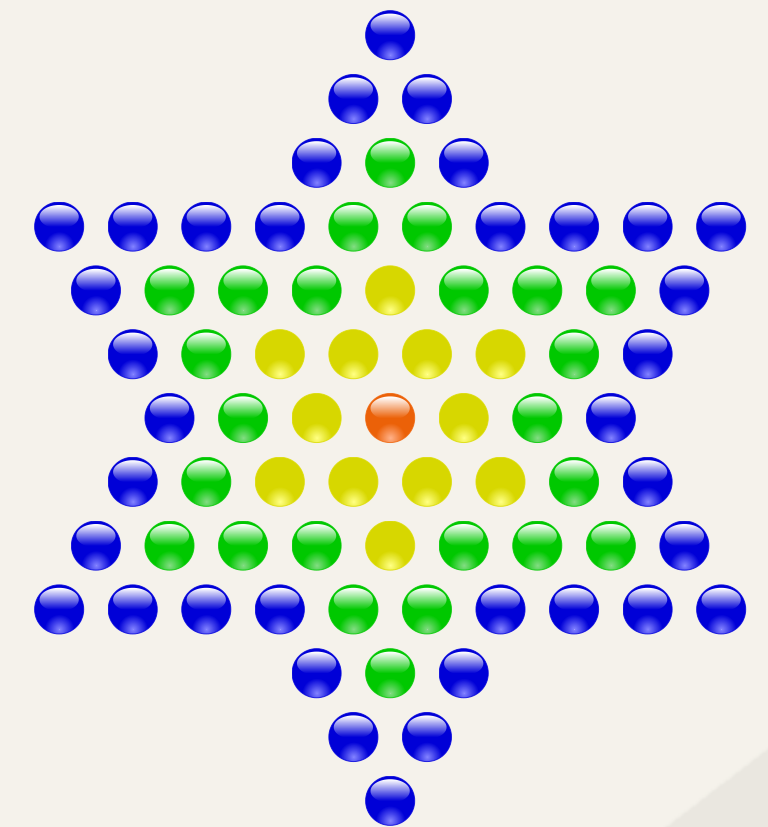
```
    struct promise_type {
        pointer value;

        static std::suspend_always initial_suspend() noexcept {
            return {};
        }
    }
```

```
    static std::suspend_always final_suspend() noexcept {
```

Star numbers:

1 13 37 73 121 181 253 337 433 ..



(source: [Wikipedia](#), CC BY-SA 4.0)

Improving algorithms using **ranges**

Ranges

A “range” is an **abstraction over iterators**, enabling a *functional style*.

Classic algorithms

```
1 std::vector vec{60, 4, 84, 14, 79, 51, 93, 25, 59};  
2  
3 std::sort(vec.begin(), vec.end());
```



We supply a pair of iterators.

C++20 ranges

```
1 std::vector vec{60, 4, 84, 14, 79, 51, 93, 25, 59};  
2  
3 std::ranges::sort(vec);
```

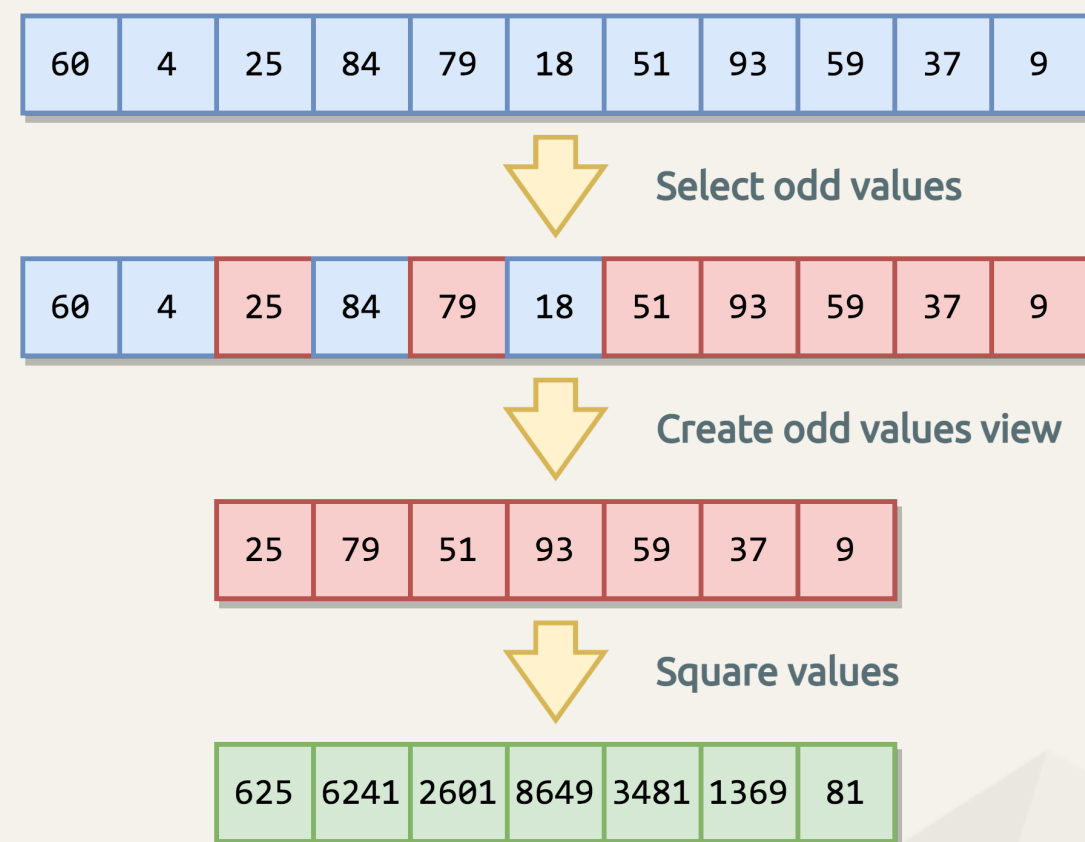


We supply a single range.

One of the main motivations for ranges is **correctness at construction**.

Range pipelines and views

A range view is a **lightweight (sub-)range representation** for ad-hoc processing.



```
1 const std::vector vec{60, 4, 84, 14, 79, 51, 93, 25, 59, 37, 9};
2
3 const auto is_odd = [](int value) -> bool { return value % 2; };
4 const auto square = [](int value) -> int { return value * value; };
5 const auto print = [](int value) { std::cout << value << '\n'; };
6
7 // Shorthand names to remove noise (advice: don't apply 'using namespace').
8 namespace rg = std::ranges;
9 namespace rv = std::views;
```

Business logic:

```
rg::for_each(vec | rv::filter(is_odd) | rv::transform(square), print);
```

Ranges, views and projections

Range views have *reference semantics* and use **lazy evaluation**.

```
1 enum class EventType { Error, Info };
2
3 struct Event {
4     unsigned long id{};
5     EventType     type{EventType::Info};
6     std::string   message;
7 };
8
9 const std::vector<Event> system_log = {
10     {0, EventType::Info, "motor1: start"},
11     {1, EventType::Info, "motor2: start"},
12     {2, EventType::Info, "motor1: OK" },
13     {3, EventType::Error, "motor2: FAIL" },
14     {4, EventType::Info, "motor1: stop" },
15     // ...
16 };
```



```
1 namespace rg = std::ranges;
2 namespace rv = std::views;
3
4 void process_log(rg::input_range auto log); // C++20 abbreviated function syntax.
5
6 const auto is_info = [](const Event& e) { return (e.type == EventType::Info); };
7 const auto is_error = [](const Event& e) { return (e.type == EventType::Error); };
```



```
1 for (auto& e : system_log | rv::filter(is_error)) {
2     std::cerr << std::format("Error event: ({}: '{}')\n", e.id, e.message);
3 }
4
5 process_log(system_log); // Feed the whole range.
6 process_log(system_log | rv::take_while(is_info)); // Feed up to first non-info.
7 process_log(system_log | rv::take(3)); // Feed first three items.
8 process_log(system_log | rv::drop(2) | rv::take(3)); // Feed third to fifth item.
```



Ranges, views and projections

Projections can be used in algorithms to **transform their behavior**.

```
1 enum class EventType { Error, Info };
2
3 struct Event {
4     unsigned long id{};
5     EventType     type{EventType::Info};
6     std::string   message;
7 };
8
9 std::vector<Event> system_log = {
10 {0, EventType::Info, "motor1: start"},
11 {1, EventType::Info, "motor2: start"},
12 {2, EventType::Info, "motor1: OK" },
13 {3, EventType::Error, "motor2: FAIL" },
14 {4, EventType::Info, "motor1: stop" },
15 // ...
16 };
```



Projections for sort transformation:

```
1 namespace rg = std::ranges;
2
3 rg::sort(system_log);           // Default: take default comparison op.
4 rg::sort(system_log, {}, &Event::id); // Sort using field 'id'.
5 rg::sort(system_log, {}, &Event::message); // Sort using field 'message'.
```



Many other algorithms take projections.

The Range-v3 library

The ranges library is **fully implemented in Range-v3** and can be used today.

```
1 enum class EventType { Error, Info };
2
3 struct Event {
4     unsigned long id{};
5     EventType     type{EventType::Info};
6     std::string   message;
7 };
8
9 const std::vector<Event> system_log = {
10     {0, EventType::Info, "motor1: start"},
11     {1, EventType::Info, "motor2: start"},
12     {2, EventType::Info, "motor1: OK" },
13     {3, EventType::Error, "motor2: FAIL" },
14     {4, EventType::Info, "motor1: stop" },
15     // ...
16 };
```



```
1 namespace rg = ranges;
2 namespace rv = ranges::views; // Lazy views of ranges.
3 namespace ra = ranges::actions; // Eager in-place range modifications.
4
5 void process_log(rg::input_range auto log);
6
7 const auto is_error = [](const Event& e) { return (e.type == EventType::Error); };
```



Use an intermediate for modification:

```
1 process_log(system_log | rv::take(50)
2               | rv::filter(is_error)
3               | rg::to<std::vector>() // <-- Create a mutable intermediate.
4               | ra::sort(std::less<>{}, &Event::id));
```



More notable *stuff*

Array view type `std::span`

A **reference value type** to any contiguous sequence of values

Can be used on `std::vector`, `std::array` and C-style arrays (including mutation)

```
1 void process_region(std::span<int> region) {  
2     for (int &value : region) {  
3         // Round up to the nearest decade.  
4         value = static_cast<int>(10.0f * std::ceil(value / 10.0f));  
5     }  
6 }
```



```
1 std::array<int, 8> data{/* ... */};  
2  
3 // Process values with index 3..6 (i.e. items 4..7):  
4 process_region({std::begin(data) + 3, 4});  
5  
6 // Process the whole array:  
7 process_region({data});
```



`std::span` can also be fixed-size by using `std::span<Type, Size>`

String formatter `std::format`

```
float    real1{3.141592654f};  
float    real2{2.718281828f};  
int      natural{82649};  
std::string name{"Gideon Mantell"};
```

Desired print result:

```
Values: [3.141593, 003, 2.7183, 82649, 0x142d9, Gideon Mantell]
```

I/O streams

```
1 // Print and format using I/O streams:  
2 std::cout << "Values: ["  
3     << std::setprecision(7) << real1 << ", "  
4     << std::setw(3) << std::setfill('0') << std::setprecision(0) << real1 << ", "  
5     << std::setprecision(5) << real2 << ", "  
6     << natural << ", "  
7     << std::hex << "0x" << natural << ", "  
8     << name << "]\n";
```



`std::format`

```
1 // Print using I/O streams, formatting with std::format:  
2 std::cout << std::format("Values: [{:.7}, {:03.0}, {:.5}, {}, 0x{:x}, {}]\n",  
3     real1, real1, real2, natural, natural, name);
```



Default comparisons and **operator**<=>

```
1 struct DbEntry {  
2     std::uint64_t key;  
3     Uuid         uuid;  
4  
5     struct Data {  
6         std::uint32_t id;  
7         std::string  value;  
8     } data;  
9 };
```



```
1 std::map<std::string, DbEntry> entries;  
2 entries["alpha"] = DbEntry{  
3     .key = 123UL,  
4     .uuid = "3d5d476e-f80c-4dec-a308-27cd190399e3",  
5     .data = {.id = 0U,  
6             .value = "Io"}};  
7 entries["beta"] = DbEntry{  
8     .key = 456UL,  
9     .uuid = "0f090b6f-2d63-4a25-8438-5b211baf7d8e",  
10    .data = {.id = 1U,  
11            .value = "Ganymede"}};
```



```
const bool result = (entries["alpha"] < entries["beta"]);
```

```
error: no match for 'operator<'  
(operand types are 'std::map<std::__cxx11::basic_string<char>,  
                    DbEntry>::mapped_type' {aka 'DbEntry'} and  
                    'std::map<std::__cxx11::basic_string<char>,  
                    DbEntry>::mapped_type' {aka 'DbEntry'})  
|   const bool result1 = (entries["alpha"] < entries["beta"]);
```

Implement custom **operator**<

OR

use C++20 **default comparison ops**

Default comparisons and operator<=>

```
1 struct DbEntry {  
2     std::uint64_t key;  
3     Uuid         uuid;  
4  
5     struct Data {  
6         std::uint32_t id;  
7         std::string  value;  
8  
9         auto operator<=>(const Data&) const = default;  
10    } data;  
11  
12    auto operator<=>(const DbEntry&) const = default;  
13};
```



```
1 std::map<std::string, DbEntry> entries;  
2 entries["alpha"] = DbEntry{  
3     .key = 123UL,  
4     .uuid = "3d5d476e-f80c-4dec-a308-27cd190399e3",  
5     .data = {.id = 0U,  
6             .value = "Io"}};  
7 entries["beta"] = DbEntry{  
8     .key = 456UL,  
9     .uuid = "0f090b6f-2d63-4a25-8438-5b211baf7d8e",  
10    .data = {.id = 1U,  
11            .value = "Ganymede"}};
```



```
const bool result = (entries["alpha"] < entries["beta"]);
```

OK

Three-way comparison

```
1 #include <compare>  
2  
3 const auto result = (entries["beta"] <=> entries["gamma"]);  
4  
5 if (std::is_eq(result)) {  
6     // ...  
7 } else if (std::is_gt(result)) {  
8     // ...  
9 } else if (std::is_lt(result)) {  
10    // ...  
11 }
```



Designated initializers

```
1 struct SystemId {  
2     Uuid      uuid;  
3     std::string name;  
4     std::string alias;  
5     bool      verified{};  
6 };  
7  
8 void verify(SystemId&& id);
```



Pre-C++20

```
1 verify({"465ff086-92c8-43a9-854e-f41f13bac804",  
2         "Quantum vector shape interceptor",  
3         "QVSI",  
4         true});
```



C++20

```
1 verify({.uuid      = "465ff086-92c8-43a9-854e-f41f13bac804",  
2          .name      = "Quantum vector shape interceptor",  
3          .alias     = "QVSI",  
4          .verified  = true});
```



Omitted fields are zero-initialized.
Initialization order must be correct.

Initialization in range-based for

Pre-C++20

```
1 auto data = receive(); // Outside the scope of the loop..
2 for (auto& chunk : data.chunks) {
3     if (!error_correct(chunk)) {
4         log_error("Failed to correct chunk #{}", chunk.id);
5     }
6 }
```



```
1 std::size_t index = 0; // Outside the scope of the loop..
2 for (const auto& e : collection) {
3     register_element(e, index);
4     ++index;
5 }
```



C++20

```
1 for (auto data = receive(); auto& chunk : data.chunks) {
2     if (!error_correct(chunk)) {
3         log_error("Failed to correct chunk #{}", chunk.id);
4     }
5 }
```



```
1 for (std::size_t index = 0; const auto& e : collection) {
2     register_element(e, index);
3     ++index;
4 }
```



Method contains for STL containers

Pre-C++20

```
1 PartId lookup(std::string_view name) {  
2     static CacheMap cache;  
3  
4     auto result = cache.find(name);  
5     if (result == cache.end()) {  
6         auto id = lookup_expensive(name);  
7         cache[name] = id;  
8         return id;  
9     }  
10  
11     return result->second;  
12 }
```



C++20

```
1 PartId lookup(std::string_view name) {  
2     static CacheMap cache;  
3  
4     if (!cache.contains(name)) {  
5         auto id = lookup_expensive(name);  
6         cache[name] = id;  
7         return id;  
8     }  
9  
10     return cache[name];  
11 }
```



Immediate functions with **constexpr**

```
1 #include <array>
2
3 using Lut = std::array<int, 10>;
4
5 constexpr Lut init_lut() { // Always evaluated at compile-time when called.
6     Lut result;
7
8     int a{0}, b{1};
9     for (auto &value : result) {
10         value = a;
11         a = b;
12         b += value;
13     }
14
15     return result;
16 }
17
18 // constexpr will ensure compile-time initialization of the variable:
19 constexpr Lut lut{init_lut()}; // 0, 1, 1, 2, 3, 5, 8, 13, 21, 34
20
21 int main() {
22     return lut[8];
23 }
```




Binary output

```
1 main:
2     mov     eax, DWORD PTR lut[rip+32]
3     ret
4 lut:
5     .long  0
6     .long  1
7     .long  1
8     .long  2
9     .long  3
10    .long  5
11    .long  8
12    .long  13
13    .long  21
14    .long  34
```

ASM

An outlook to C++23



C++23 in one slide

C++23: not nearly as huge as C++20, but still has a lot to offer!

- Standard library modules (**import** std),
- Result type `std::expected`,
- Multidimensional array view `std::mdspan`,
- Proper printing support with `std::print`,
- Associative containers `std::flat_map` and `std::flat_set`,
- Coroutine generator `std::generator`,
- Stack trace evaluation with `std::stacktrace`,
- Lots of range library extensions.

Multidimensional array view `mdspan`

```
1 namespace rg = std::ranges;
2 namespace rv = std::ranges::views;
3
4 const auto data = rv::iota('a') | rv::take(12) | rg::to<std::vector>; // 'a'..'l'
5
6 auto view1 = std::mdspan{data.data(), 3, 4}; // Single plane, 3x4 elements.
7
8 for (std::size_t x{0}; x != view1.extent(0); x++) {
9     for (std::size_t y{0}; y != view1.extent(1); y++) {
10         std::print(" {} ", view1[x, y]);
11     }
12     std::print("\n");
13 }
14
15 auto view2 = std::mdspan{data.data(), 3, 2, 2}; // Three planes, 2x2 elements.
16
17 for (std::size_t x{0}; x != view2.extent(0); x++) {
18     std::print("Plane {}: \n", x);
19     for (std::size_t y{0}; y != view2.extent(1); y++) {
20         for (std::size_t z{0}; z != view2.extent(2); z++) {
21             std::print(" {} ", view2[x, y, z]);
22         }
23         std::print("\n");
24     }
25 }
```

view1:

```
a b c d
e f g h
i j k l
```

view2:

```
Plane 0:
a b
c d
Plane 1:
e f
g h
Plane 2:
i j
k l
```


Result type `std::expected`

Exceptions

```
1 float read_temperature() {  
2     // ..may throw an exception!..  
3 }  
4  
5 void display_temperature() {  
6     const float t = read_temperature();  
7     std::print("Temperature: {} °C", t);  
8 }  
9  
10 display_temperature();
```



Exception handling?

C++23 `std::expected`

```
1 using ReadResult = std::expected<float, std::string>;  
2  
3 ReadResult read_temperature() {  
4     // ...  
5 }  
6  
7 void display_temperature() {  
8     if (const auto t = read_temperature(); t) {  
9         std::print("Temperature: {} °C", t.value());  
10    } else {  
11        std::print("Error ({})", t.error());  
12    }  
13 }  
14  
15 display_temperature();
```



Coroutine generator `std::generator`

Standard library support for **synchronous coroutine generators**

```
1 #include <cstdint>
2 #include <generator>
3 #include <print>
4 #include <ranges>
5
6 std::generator<std::uint64_t> star(std::uint64_t n = 0L) {
7     while (true) {
8         n++;
9         co_yield 6 * n * (n - 1) + 1;
10    }
11 }
12
13 int main() {
14     for (auto i : star() | std::views::take(20)) {
15         std::print("{} ", i);
16     }
17     std::print("\n");
18 }
```

```
1 13 37 73 121 181 253 337 433 541 661 793 937 1093 1261 1441 1633 1837 2053 2281
```



Range extensions

Many useful range views have been added, including **range conversion**

```
1 namespace rg = std::ranges;
2 namespace rv = std::ranges::views;
3
4 using namespace std::literals;
5 std::vector brands = {"Gibson"sv, "Fender"sv, "Ibanez"sv, "Martin"sv};
6 std::vector types  = {"Les Paul"sv, "Telecaster"sv, "PGM"sv, "D-45"sv};
7
8 const auto result = rv::zip(brands, types) | rg::to<std::vector>();
9
10 for (std::size_t index{0}; auto [a, b] : result) {
11     std::print("{}: {} {}\n", index++, a, b);
12 }
```

```
0: Gibson Les Paul
1: Fender Telecaster
2: Ibanez PGM
3: Martin D-45
```



std::flat_set and std::flat_map

“Flat storage” associated containers (**faster iteration + efficient memory use**)

```
1 enum class Brand { Gibson, Fender, Ibanez, Martin };
2
3 using KeyValue = std::pair<Brand, std::string_view>;
4 const std::initializer_list<KeyValue> list{{Brand::Gibson, "Les Paul"sv },
5                                             {Brand::Fender, "Telecaster"sv},
6                                             {Brand::Ibanez, "PGM"sv      },
7                                             {Brand::Martin, "D-45"sv      }};
8
9 std::map<Brand, std::string_view>      map1;
10 std::flat_map<Brand, std::string_view> map2;
11
12 for (const auto& kv : list) {
13     map1.insert(kv); // Insertion is cheaper for std::map.
14     map2.insert(kv); // Insertion in std::flat_map requires copyability + movability.
15 }
16
17 // Verify order-equality:
18 for (auto kv1 = map1.cbegin(); const auto& [key2, value2] : map2) {
19     const auto& [key1, value1] = *kv1++;
20     assert((key1 == key2) && (value1 == value2));
21 }
```



End

Thank you 😊



 github.com/krisvanrens

All emoji in this presentation are part of the [Twemoji set](#), licensed under [CC-BY 4.0](#).
All other images are mine, unless specified otherwise.