



PROGRAMMING IN C++

Jülich Supercomputing Centre

8 – 12 May 2023 | Sandipan Mohanty | Forschungszentrum Jülich, Germany

Day 5

Type erasure

TYPE ERASURE TECHNIQUE

```
1  auto f(int i) -> PolyVal;
2  void elsewhere() {
3      std::vector<PolyVal> v;
4      v.push_back(1);
5      v.push_back(2.0);
6      v.push_back("Green"s);
7
8      for (auto&& elem : v) {
9          func1(elem);
10     }
11     PolyVal X = f(i);
12 }
```

- Polymorphic behaviour attained using a class hierarchy and virtual functions...
 - is extensible by simply inheriting from the `Base` type and overriding the virtual functions
 - But, it has “reference symantics”, so that we can not return those polymorphic objects by value from functions
 - Built in types can not be accommodated into the same hierarchy
- `variant` provides a solution to the two problems above, but we need to commit to a fixed number of polymorphic types in the problem, from the outset
- `std::any` is a library provided facility for type erasure

TYPE ERASURE TECHNIQUE

```
1 void func1(int x);
2 void func1(double x);
3 void func1(std::string x);
4 auto f(int i) -> PolyVal;
5 void elsewhere() {
6     std::vector<PolyVal> v;
7     v.push_back(1);
8     v.push_back(2.0);
9     v.push_back("Green"s);
10
11     for (auto&& elem : v) {
12         func1(elem);
13     }
14     PolyVal X{3.141};
15     // func1(X) should go to func1(double)
16     X = PolyVal{"some string"s};
17     // func1(X) should now go to func1(string)
18     X = f(i);
19     // func1(X) should redirect according to what
20     // polymorphic value f happens to return
21 }
```

- We want a type `PolyVal`, so that we can store different types of entities in it
- A uniform container of `PolyVal` should be able to hold values of different types
- When a certain instance is used, it should still be able to behave according to the value it is currently holding.
- We should be able to copy a `PolyVal` object using normal copy construction or copy assignment in such a way that the copy of a `PolyVal` storing a `Triangle` would still behave as a `Triangle`

TYPE ERASURE TECHNIQUE

```
1  class PolyVal {
2      struct Internal {
3          virtual ~Internal() noexcept = default;
4          virtual auto clone() const -> std::unique_ptr<Internal> = 0;
5          virtual void func1_() const = 0;
6      };
7      template <class T>
8      struct Wrapped : public Internal // continued...
9
10 public:
11     template <class T>
12     PolyVal(const T& var) : ptr{ std::make_unique<Wrapped<T>>(var) } {}
13     PolyVal(const PolyVal& other) : ptr { other.ptr->clone() } {}
14 private:
15     std::unique_ptr<Internal> ptr;
16 };
```

-
- Make a normal class with an internal class with virtual functions defining the desired interface, and another internal wrapper class template deriving from the internal base
 - Give the outer class one template constructor (unrestrained here to isolate the TE technique)

TYPE ERASURE TECHNIQUE

```
1  class PolyVal {
2      struct Internal {
3          virtual ~Internal() noexcept = default;
4          virtual auto clone() const -> std::unique_ptr<Internal> = 0;
5          virtual void func1_() const = 0;
6      };
7      template <class T>
8      struct Wrapped : public Internal // continued...
9
10 public:
11     template <class T>
12     PolyVal(const T& var) : ptr{ std::make_unique<Wrapped<T>>(var) } {}
13     PolyVal(const PolyVal& other) : ptr { other.ptr->clone() } {}
14 private:
15     std::unique_ptr<Internal> ptr;
16 };
```

- Let the class contain a smart pointer to this base, but initialize that member using a class template which inherits from the internal base.
- Implement a copy constructor for `PolyVal` by using a virtual `clone()` function for the internal class
- Use the template constructor to create a wrapped object containing a copy of the input parameter

TYPE ERASURE TECHNIQUE

```
1  class PolyVal {
2      template <class T>
3      struct Wrapped : public Internal {
4          Wrapped(T ex) : obj{ex} {}
5          ~Wrapped() noexcept override {}
6          auto clone() const -> std::unique_ptr<Internal> override
7          {
8              return std::make_unique<Wrapped>(obj);
9          }
10         void func1_() const override { func1(obj); }
11         T obj;
12     };
13 };
```

- The internal wrapper should store an object of the template parameter type
- It should provide copy, clone etc.
- It should redirect function calls in our original requirement to free functions

TYPE ERASURE TECHNIQUE

```
1  class PolyVal {
2      template <class T>
3      struct Wrapped : public Internal {
4          Wrapped(T ex) : obj{ex} {}
5          ~Wrapped() noexcept override {}
6          auto clone() const -> std::unique_ptr<Internal> override
7          {
8              return std::make_unique<Wrapped>(obj);
9          }
10         void func1_() const override { func1(obj); }
11         T obj;
12     };
13 };
```

- As long as those free functions exist for a type `F`, it will be possible to create objects of `PolyVal` type from type `F`

Exercise 1.1:

`examples/TE/PolyVal.cc` contains the code corresponding to the slides shown here. Verify that we achieve our purpose of having a copyable object preserving polymorphic behaviour. Add a function `func1()` (processing a new type) into the mix, and extend the existing setup.

Exercise 1.2:

Sequences of objects with polymorphic behaviour is a frequently occurring programming problem. We have seen one example before, with a vector of `unique_ptr<Shape>`, filled with newly created instances of types inherited from `Shape`, such as `Circle`, `Triangle` etc. The problem can be solved in many alternative ways. `examples/polymorphic` contains 4 subdirectories with different approaches to the geometric object example. (i) Inheritance with virtual functions (ii) `std::variant` with visitors (iii) Using `std::any` (iv) Custom type erasure.

VALARRAY

```
1  #include <valarray>
2
3  void varray_ops()
4  {
5      std::valarray V1(0., 1000000UL);
6      std::valarray<double> V2;
7      v2.resize(1000000UL, 0.);
8      auto x = exp(-V1 * V1) * sin(V2);
9      if (x.sum() < 100.0) {
10         //
11     }
12 }
```

- Another dynamic array type
- Mostly intended for numeric operations
- Expression template based whole array math operations
- Algorithms through `std::begin(v)` etc., instead of own member functions
- Bizarre constructor with different convention compared to any other container in the STL.

NUMERIC ALGORITHMS

```
1  #include <numeric>
2
3  using std::reduce;
4  using std::transform_reduce;
5
6  auto res = reduce(v.begin(), v.end());
7  auto res = reduce(v.begin(), v.end(), init);
8  auto res = reduce(v.begin(), v.end(),
9      init, std::plus<double>{});
10 auto res = transform_reduce(
11     u.begin(), u.end(),
12     v.begin(), init);
13 auto res = transform_reduce(
14     u.begin(), u.end(),
15     v.begin(), init, reduce_op, transf_op);
16 auto res = transform_reduce(
17     std::execution::par,
18     u.begin(), u.end(),
19     v.begin(), init, reduce_op, transf_op);
```

- Algorithms focused on numeric calculations are in the `numeric` header
- Given `b`, `e` as iterators in a range V ,
 $\text{reduce}(b, e) : \sum_{i=b}^e V_i$
- $\text{transform_reduce}(b, e, f) : \sum_{i=b}^e f(V_i)$
- $\text{adjacent_difference}(b, e) :$
 $\{V_b, (V_{b+1} - V_b), (V_{b+2} - V_{b+1}), \dots\}$
- Parallel versions also in the library
- To run the numeric operations in parallel,
use the parallel execution policy

SPAN

```
1  using std::span;
2  using std::transform_reduce;
3  using std::plus;
4  using std::multiplies;
5  auto compute(span<const double> u,
6              span<const double> v) -> double
7  {
8      return transform_reduce(
9          u.begin(), u.end(),
10         v.begin(), 0., plus<double>{},
11         multiplies<double>{});
12 }
13
14 void elsewhere(double* x, double* y,
15               unsigned N)
16 {
17     return compute(span(x, N), span(y, N));
18 }
```

- Non-owning view type for a contiguous range
- No memory management
- Numeric operations can often be expressed in terms of existing arrays in memory, irrespective of how they got there and who cleans up after they expire
- `span` is designed to be that input for such functions
- Cheap to copy: essentially a pointer and a size
- STL container like interface

Exercise 1.3:

`examples/spans` is a directory containing the `compute` function as shown here. Notice how this function is used directly using C++ array types as arguments instead of `spans`, and indirectly when we only have pointers.

RANGES

```
1 std::vector v{ 1, 2, 3, 4, 5, 6, 7, 8, 9 };
2 // before std::ranges we did this...
3 std::reverse(v.begin(), v.end());
4 std::rotate(v.begin(), v.begin() + 3, v.end());
5 std::sort(v.begin(), v.end());
```

```
1 std::vector v{ 1, 2, 3, 4, 5, 6, 7, 8, 9 };
2 namespace sr = std::ranges;
3 sr::reverse(v);
4 sr::rotate(v, v.begin() + 3);
5 sr::sort(v);
```

- The `<ranges>` header defines a set of algorithms taking “ranges” as inputs instead of pairs of iterators
- A `range` is a **concept** : something with `sr::begin()` , which returns an entity which can be used to iterate over the elements, and `sr::end()` which returns a sentinel which is equality comparable with an iterator, and indicates when the iteration should stop.
- `sr::sized_range` : the range knows its size in constant time
- `input_range` , `output_range` etc. based on the iterator types
- `borrowed_range` : a type such that its iterators can be returned without the danger of dangling.
- `view` is a range with constant time copy/move/assignment

USING RANGES FROM STD OR FROM RANGE-V3

```
1 // cxx220ranges
2 #include <version>
3 #ifdef __cpp_lib_ranges
4     #include<ranges>
5     namespace sr = std::ranges;
6     namespace sv = sr::views;
7 #elif __has_include (<range/v3/all.hpp>)
8     #include<range/v3/all.hpp>
9     namespace sr = ranges;
10    namespace sv = sr::views;
11    #warning Using ranges-v3 3rd party library
12 #else
13 #error No suitable header for C++20 ranges was found!
14 #endif
```

- The C++20 `<ranges>` library is based on the open source `range-v3` library. Parts of the `range-v3` library were adopted for C++20, more might be added in C++23.
- Even if the standard library shipping with some compilers do not have many features of `<ranges>`, one can start using them, with a redirecting header, which makes use of another standard library feature
- Including `<version>` results in the definition of library feature test macros, which can be used to choose between different header files

Our examples are actually written using a redirecting header as shown here. Compilation with GCC uses the compiler's own version. Compilation with Clang uses the `range-v3` version.

FUN WITH RANGES AND VIEWS

```
1 // examples/ranges0.cc
2 #include <ranges>
3 #include <span>
4 auto sum(std::ranges::input_range auto&& seq) {
5     std::iter_value_t<decltype(seq)> ans{};
6     for (auto x : seq) ans += x;
7     return ans;
8 }
9 auto main() -> int
10 {
11     //using various namespaces;
12     cout << "vector    : " << sum(vector( { 9,8,7,2 } )) << "\n";
13     cout << "list      : " << sum(list( { 9,8,7,2 } )) << "\n";
14     cout << "valarray : " << sum(valarray({ 9,8,7,2 } )) << "\n";
15     cout << "array    : "
16         << sum(array<int,4>({ 9,8,7,2 } )) << "\n";
17     cout << "array    : "
18         << sum(array<string, 4>({ "9"s,"8"s,"7"s,"2"s } )) << "\n";
19     int A[]{1,2,3};
20     cout << "span(built-in array) : " << sum(span(A)) << "\n";
21 }
```

FUN WITH RANGES AND VIEWS

- The `ranges` library gives us many useful concepts describing sequences of objects.
- The function template `sum` in `examples/ranges0.cc` accepts any input range, i.e., some entity whose iterators satisfy the requirements of an `input_iterator`.
- Notice how we obtain the value type of the range
- Many STL algorithms have `range` versions in C++20. They are functions like `sum` taking various kinds of ranges as input.
- The range concept is defined in terms of
 - the existence of an iterator type and a sentinel type.
 - the iterator should behave like an iterator, e.g., allow `++it` `*it` etc.
 - it should be possible to compare the iterators with other iterators or with a sentinel for equality.
 - A `begin()` function returning an iterator and an `end()` function returning a sentinel

FUN WITH RANGES AND VIEWS

```
1 // examples/iota.cc
2 #include <ranges>
3 #include <iostream>
4 auto main() -> int {
5     namespace sv = std::views;
6     for (auto i : sv::iota(1UL)) {
7         if ((i+1) % 10000UL == 0UL) {
8             std::cout << i << ' ';
9             if ((i+1) % 100000UL == 0UL)
10                 std::cout << '\n';
11             if (i >= 100000000UL) break;
12         }
13     }
14 }
```

- All containers are ranges, but not all ranges are containers
- `std::string_view` is a perfectly fine range. Has iterators with the right properties. Has `begin()` and `end()` functions. It does not own the contents, but “ownership” is not part of the idea of a range.
- We could take this further by creating `views` which need not actually contain any objects of the sequence, but simply fake it when we dereference the iterators.
- Example: the standard view
`std::views::iota(integer)` gives us an infinite sequence of integers starting at a given value.

FUN WITH RANGES AND VIEWS

```
1 // examples/iota.cc
2 #include <ranges>
3 #include <iostream>
4 auto main() -> int {
5     namespace sv = std::views;
6     for (auto i : sv::iota(1UL)) {
7         if ((i+1) % 10000UL == 0UL) {
8             std::cout << i << ' ';
9             if ((i+1) % 100000UL == 0UL)
10                 std::cout << '\n';
11             if (i >= 100000000UL) break;
12         }
13     }
14 }
```


- All containers are ranges, but not all ranges are containers
- `std::string_view` is a perfectly fine range. Has iterators with the right properties. Has `begin()` and `end()` functions. It does not own the contents, but “ownership” is not part of the idea of a range.
- We could take this further by creating `views` which need not actually contain any objects of the sequence, but simply fake it when we dereference the iterators.
- Example: the standard view
`std::views::iota(integer)` gives us an infinite sequence of integers starting at a given value.

FUN WITH RANGES AND VIEWS

```
1 // examples/iota.cc
2 #include <ranges>
3 #include <iostream>
4 auto main() -> int {
5     namespace sv = std::views;
6     for (auto i : sv::iota(1UL)) {
7         if ((i+1) % 10000UL == 0UL) {
8             std::cout << i << ' ';
9             if ((i+1) % 100000UL == 0UL)
10                 std::cout << '\n';
11             if (i >= 100000000UL) break;
12         }
13     }
14 }
```

- All containers are ranges, but not all ranges are containers
- `std::string_view` is a perfectly fine range. Has iterators with the right properties. Has `begin()` and `end()` functions. It does not own the contents, but “ownership” is not part of the idea of a range.
- We could take this further by creating `views` which need not actually contain any objects of the sequence, but simply fake it when we dereference the iterators.
- Example: the standard view
`std::views::iota(integer)` gives us an infinite sequence of integers starting at a given value.

FUN WITH RANGES AND VIEWS



- All containers are ranges, but not all ranges are containers
- `std::string_view` is a perfectly fine range. Has iterators with the right properties. Has `begin()` and `end()` functions. It does not own the contents, but “ownership” is not part of the idea of a range.
- We could take this further by creating `views` which need not actually contain any objects of the sequence, but simply fake it when we dereference the iterators.
- Example: the standard view `std::views::iota(integer)` gives us an infinite sequence of integers starting at a given value.

FUN WITH RANGES AND VIEWS

```
1  #include <ranges>
2  #include <iostream>
3  auto main() -> int {
4      namespace sv = std::views;
5      for (auto i : sv::iota(1UL) ) {
6          if ((i+1) % 10000UL == 0UL) {
7              std::cout << i << ' ';
8              if ((i+1) % 100000UL == 0UL)
9                  std::cout << '\n';
10             if (i >= 100000000UL) break;
11         }
12     }
13 }
```

- All containers are ranges, but not all ranges are containers
- `std::string_view` is a perfectly fine range. Has iterators with the right properties. Has `begin()` and `end()` functions. It does not own the contents, but “ownership” is not part of the idea of a range.
- We could take this further by creating `views` which need not actually contain any objects of the sequence, but simply fake it when we dereference the iterators.
- Example: the standard view
`std::views::iota(integer)` gives us an infinite sequence of integers starting at a given value.

BORROWED RANGES

```
1 // examples/dangling0.cc
2 auto get_vec() {
3     std::vector v{ 2, 4, -1, 8, 0, 9 };
4     return v;
5 }
6 auto main() -> int {
7     auto v = get_vec();
8     auto iter = std::min_element(v.begin(),
9                                 v.end());
10    std::cout << "Minimum " << *iter << "\n";
11 }
```

- The `min_element` function finds the minimum element in a range and returns an iterator

Example from a CPPCon 2020 talk by Tristan Brindle.
[Link.](#)

BORROWED RANGES

```
1 // examples/dangling0.cc
2 auto get_vec() {
3     std::vector v{ 2, 4, -1, 8, 0, 9 };
4     return v;
5 }
6 auto main() -> int {
7     auto v = get_vec();
8     auto iter = sr::min_element(v);
9
10    std::cout << "Minimum " << *iter << "\n";
11 }
```

- The `min_element` function finds the minimum element in a range and returns an iterator
- The version from the ranges library takes only a range

Example from a CPPCon 2020 talk by Tristan Brindle.
[Link.](#)

BORROWED RANGES

```
1 // examples/dangling0.cc
2 auto get_vec() {
3     std::vector v{ 2, 4, -1, 8, 0, 9 };
4     return v;
5 }
6 auto main() -> int {
7
8     auto iter = sr::min_element(get_vec());
9
10    std::cout << "Minimum " << *iter << "\n";
11 }
```

- The `min_element` function finds the minimum element in a range and returns an iterator
- The version from the ranges library takes only a range
- It may be tempting to directly feed the output from a function to the algorithm. But, we would receive an iterator to a container that is already destructed, i.e., a dangling iterator. Dereferencing should therefore lead to a SEGFAULT.

Example from a CPPCon 2020 talk by Tristan Brindle.
[Link.](#)

BORROWED RANGES

```
1 // examples/dangling0.cc
2 auto get_vec() {
3     std::vector v{ 2, 4, -1, 8, 0, 9 };
4     return v;
5 }
6 auto main() -> int {
7
8     auto iter = sr::min_element(get_vec());
9
10    std::cout << "Minimum " << *iter << "\n";
11 }
```

- The `min_element` function finds the minimum element in a range and returns an iterator
- The version from the ranges library takes only a range
- It may be tempting to directly feed the output from a function to the algorithm. But, we would receive an iterator to a container that is already destructed, i.e., a dangling iterator. Dereferencing should therefore lead to a SEGFAULT.
- In reality, what happens is this!

Example from a CPPCon 2020 talk by Tristan Brindle.
Link.

```
error: no match for 'operator*' (operand type is 'std::ranges::dangling')
19 |     std::cout << "Minimum value is " << *iter << "\n";
```

BORROWED RANGES

```
1 // examples/dangling0.cc
2 auto get_vec() {
3     std::vector v{ 2, 4, -1, 8, 0, 9 };
4     return v;
5 }
6 auto main() -> int {
7
8     auto iter = sr::min_element(get_vec());
9
10    std::cout << "Minimum " << *iter << "\n";
11 }
```

- The ranges algorithms are written with overloads such that when you pass an R-value reference of a container as input, the output type is `ranges::dangling`, an empty `struct` with no operations defined.
- `iter` here will be deduced to be of type `ranges::dangling`, and hence `*iter` leads to that insightful error message.

```
error: no match for 'operator*' (operand type is 'std::ranges::dangling')
19 |     std::cout << "Minimum value is " << *iter << "\n";
```

- When the input was an L-value reference, the algorithm returning the iterator returned a valid iterator.
- Therefore: valid use cases work painlessly, and invalid ones result in actionable insights from the compiler!

BORROWED RANGES

```
1 // examples/dangling1.cc
2 static std::vector u{2, 3, 4, -1, 9};
3 static std::vector v{3, 1, 4, 1, 5};
4 auto get_vec(int c) -> std::span<int> {
5     return { (c % 2 == 0) ? u : v };
6 }
7 auto main(int argc, char* argv[]) -> int {
8     auto iter = sr::min_element(get_vec(argc));
9     // iter is valid, even if its parent span
10    // has expired.
11    std::cout << "Minimum " << *iter << "\n";
12 }
```

- Sometimes, an iterator can point to a valid element even when the “container” (imposter) has been destructed. `span`, `string_view` etc. do not own the elements in their range.
- No harm in returning real iterators of these objects, even if they are R-values. Even in this case, there is no danger of dangling.
- A `borrowed_range` is a range so that its iterators can be returned from a function without the danger of dangling, i.e.,
it is an L-value reference or
has been explicitly certified to be a borrowed range.

BORROWED RANGES

```
1 // examples/dangling1.cc
2 static std::vector u{2, 3, 4, -1, 9};
3 static std::vector v{3, 1, 4, 1, 5};
4 auto get_vec(int c) -> std::span<int> {
5     return { (c % 2 == 0) ? u : v };
6 }
7 auto main(int argc, char* argv[]) -> int {
8     auto iter = sr::min_element(get_vec(argc));
9     // iter is valid, even if its parent span
10    // has expired.
11    std::cout << "Minimum " << *iter << "\n";
12 }
```

```
template <class T>
concept borrowed_range = range<T> &&
    ( is_lvalue_reference_v<T> || enable_borrowed_range<remove_cvref_t<T>> )
```

- Sometimes, an iterator can point to a valid element even when the “container” (imposter) has been destructed. `span`, `string_view` etc. do not own the elements in their range.
- No harm in returning real iterators of these objects, even if they are R-values. Even in this case, there is no danger of dangling.
- A `borrowed_range` is a range so that its iterators can be returned from a function without the danger of dangling, i.e.,
it is an L-value reference or
has been explicitly certified to be a borrowed range

VIEW ADAPTORS

```
1 namespace sv = std::views;  
2 std::vector v{1,2,3,4,5};  
3 auto v3 = sv::take(v, 3);  
4 // v3 is some sort of object so  
5 // that it represents the first  
6 // 3 elements of v. It does not  
7 // own anything, and has constant  
8 // time copy/move etc. It's a view.  
9  
10 // sv::take() is a view adaptor
```

- A `view` is a range with constant time copy, move etc. Think `string_view`
- A view adaptor is a function object, which takes a “viewable” range as an input and constructs a view out of it. `viewable` is defined as “either a `borrowed_range` or already a view.”
- View adaptors in the `<ranges>` library have very interesting properties, and make some new ways of coding possible.

VIEW ADAPTORS

```
Adaptor(Viewable) -> View  
Viewable | Adaptor -> View  
V | A1 | A2 | A3 ... -> View
```

```
Adaptor(Viewable, Args...) -> View  
Adaptor(Args...)(Viewable) -> View  
Viewable | Adaptor(Args...) -> View
```

- A **view** itself is trivially viewable.
- Since a view adaptor produces a view, successive applications of such adaptors makes sense.
- If an adaptor takes only one argument, it can be called using the pipe operator as shown. These adaptors can then be chained to produce more complex adaptors.
- For adaptors taking multiple arguments, there exists an equivalent adaptor taking only the viewable range.

VIEW ADAPTORS

```
Adaptor(Viewable) -> View  
Viewable | Adaptor -> View  
V | A1 | A2 | A3 ... -> View
```

```
Adaptor(Viewable, Args...) -> View  
Adaptor(Args...) (Viewable) -> View  
Viewable | Adaptor(Args...) -> View
```

- A `view` itself is trivially viewable.
- Since a view adaptor produces a view, successive applications of such adaptors makes sense.
- If an adaptor takes only one argument, it can be called using the pipe operator as shown. These adaptors can then be chained to produce more complex adaptors.
- For adaptors taking multiple arguments, there exists an equivalent adaptor taking only the viewable range.

VIEW ADAPTORS

```
Adaptor(Viewable) -> View  
Viewable | Adaptor -> View  
V | A1 | A2 | A3 ... -> View
```

```
Adaptor(Viewable, Args...) -> View  
Adaptor(Args...) (Viewable) -> View  
Viewable | Adaptor(Args...) -> View
```

- A `view` itself is trivially viewable.
- Since a view adaptor produces a view, successive applications of such adaptors makes sense.
- If an adaptor takes only one argument, it can be called using the pipe operator as shown. These adaptors can then be chained to produce more complex adaptors.
- For adaptors taking multiple arguments, there exists an equivalent adaptor taking only the viewable range.

VIEW ADAPTORS

```
Adaptor(Viewable) -> View  
Viewable | Adaptor -> View  
V | A1 | A2 | A3 ... -> View
```

```
Adaptor(Viewable, Args...) -> View  
Adaptor(Args...) (Viewable) -> View  
Viewable | Adaptor(Args...) -> View
```

- A `view` itself is trivially viewable.
- Since a view adaptor produces a view, successive applications of such adaptors makes sense.
- If an adaptor takes only one argument, it can be called using the pipe operator as shown. These adaptors can then be chained to produce more complex adaptors.
- For adaptors taking multiple arguments, there exists an equivalent adaptor taking only the viewable range.

VIEW ADAPTORS

```
Adaptor(Viewable) -> View  
Viewable | Adaptor -> View  
V | A1 | A2 | A3 ... -> View
```

```
Adaptor(Viewable, Args...) -> View  
Adaptor(Args...) (Viewable) -> View  
Viewable | Adaptor(Args...) -> View
```

- A `view` itself is trivially viewable.
- Since a view adaptor produces a view, successive applications of such adaptors makes sense.
- If an adaptor takes only one argument, it can be called using the pipe operator as shown. These adaptors can then be chained to produce more complex adaptors.
- For adaptors taking multiple arguments, there exists an equivalent adaptor taking only the viewable range.

So what are we going to do with this ?

VIEW ADAPTORS

Pretend that you want to verify that $\sin^2(x) + \cos^2(x) = 1$

VIEW ADAPTORS

Pretend that you want to verify that $\sin^2(x) + \cos^2(x) = 1$

- Start with a range of integers from 0 to $N = 10000$.
- $R_0 = \{0, 1, 2, 3 \dots\}$

VIEW ADAPTORS

Pretend that you want to verify that $\sin^2(x) + \cos^2(x) = 1$

- Start with a range of integers from 0 to $N = 10000$.
- Map the integer range to real numbers in the range $[0, 2\pi)$
- $R_0 = \{0, 1, 2, 3 \dots\}$
- $R_1 = T_{10}R_0 = T(n \mapsto \frac{n\pi}{N})R_0$

VIEW ADAPTORS

Pretend that you want to verify that $\sin^2(x) + \cos^2(x) = 1$

- Start with a range of integers from 0 to $N = 10000$.
 - Map the integer range to real numbers in the range $[0, 2\pi)$
 - Evaluate $\sin^2(x) + \cos^2(x) - 1$ over the resulting range
- $R_0 = \{0, 1, 2, 3 \dots\}$
 - $R_1 = T_{10}R_0 = T(n \mapsto \frac{n\pi}{N})R_0$
 - $R_2 = T_{21}R_1 = T(x \mapsto (\sin^2(x) + \cos^2(x) - 1))R_1$

$$\begin{aligned} R_2 &= T_{21}R_1 = T_{21}T_{10}R_0 \\ &= R_0|T_{10}|T_{21} \\ &= R_0|(T_{10}|T_{21}) \end{aligned}$$

VIEW ADAPTORS

Pretend that you want to verify that $\sin^2(x) + \cos^2(x) = 1$

- Start with a range of integers from 0 to $N = 10000$.
 - Map the integer range to real numbers in the range $[0, 2\pi)$
 - Evaluate $\sin^2(x) + \cos^2(x) - 1$ over the resulting range
 - If absolute value of any of the values in the result exceeds ϵ , we have found a counter example
- $R_0 = \{0, 1, 2, 3 \dots\}$
 - $R_1 = T_{10}R_0 = T(n \mapsto \frac{n\pi}{N})R_0$
 - $R_2 = T_{21}R_1 = T(x \mapsto (\sin^2(x) + \cos^2(x) - 1))R_1$

$$\begin{aligned} R_2 &= T_{21}R_1 = T_{21}T_{10}R_0 \\ &= R_0|T_{10}|T_{21} \\ &= R_0|(T_{10}|T_{21}) \end{aligned}$$

VIEW ADAPTORS

Pretend that you want to verify that $\sin^2(x) + \cos^2(x) = 1$

- Start with a range of integers from 0 to $N = 10000$.
 - Map the integer range to real numbers in the range $[0, 2\pi)$
 - Evaluate $\sin^2(x) + \cos^2(x) - 1$ over the resulting range
 - If absolute value of any of the values in the result exceeds ϵ , we have found a counter example
 - Intuitive left-to-right readability
- $R_0 = \{0, 1, 2, 3 \dots\}$
 - $R_1 = T_{10}R_0 = T(n \mapsto \frac{n\pi}{N})R_0$
 - $R_2 = T_{21}R_1 = T(x \mapsto (\sin^2(x) + \cos^2(x) - 1))R_1$

$$\begin{aligned} R_2 &= T_{21}R_1 = T_{21}T_{10}R_0 \\ &= R_0|T_{10}|T_{21} \\ &= R_0|(T_{10}|T_{21}) \end{aligned}$$

VIEW ADAPTORS

Pretend that you want to verify that $\sin^2(x) + \cos^2(x) = 1$

- Start with a range of integers from 0 to $N = 10000$.
- Map the integer range to real numbers in the range $[0, 2\pi)$
- Evaluate $\sin^2(x) + \cos^2(x) - 1$ over the resulting range
- If absolute value of any of the values in the result exceeds ϵ , we have found a counter example
- Intuitive left-to-right readability

- $R_0 = \{0, 1, 2, 3 \dots\}$
- $R_1 = T_{10}R_0 = T(n \mapsto \frac{n\pi}{N})R_0$
- $R_2 = T_{21}R_1 = T(x \mapsto (\sin^2(x) + \cos^2(x) - 1))R_1$

$$\begin{aligned} R_2 &= T_{21}R_1 = T_{21}T_{10}R_0 \\ &= R_0 | T_{10} | T_{21} \\ &= R_0 | (T_{10} | T_{21}) \end{aligned}$$

```
find . -name "*.cc" | xargs grep "if" | grep -v "constexpr" | less
```

VIEW ADAPTORS

```
find . -name "*.cc" | xargs grep "if" | grep -v "constexpr" | less
```

- Command line of Linux, Mac OS ...

VIEW ADAPTORS

```
find . -name "*.cc" | xargs grep "if" | grep -v "constexpr" | less
```

- Command line of Linux, Mac OS ...
- Small utilities. Each program does one thing, and does it well.

VIEW ADAPTORS

```
find . -name "*.cc" | xargs grep "if" | grep -v "constexpr" | less
```

- Command line of Linux, Mac OS ...
- Small utilities. Each program does one thing, and does it well.
- There is a way to chain them together with the pipe

VIEW ADAPTORS

```
find . -name "*.cc" | xargs grep "if" | grep -v "constexpr" | less
```

- Command line of Linux, Mac OS ...
- Small utilities. Each program does one thing, and does it well.
- There is a way to chain them together with the pipe
- Overall usefulness of the tool set is amplified exponentially!

VIEW ADAPTORS

```
find . -name "*.cc" | xargs grep "if" | grep -v "constexpr" | less
```

- Command line of Linux, Mac OS ...
- Small utilities. Each program does one thing, and does it well.
- There is a way to chain them together with the pipe
- Overall usefulness of the tool set is amplified exponentially!
- What about writing something similar in C++ ?

VIEW ADAPTORS

Pretend that you want to verify that $\sin^2(x) + \cos^2(x) = 1$

VIEW ADAPTORS

Pretend that you want to verify that $\sin^2(x) + \cos^2(x) = 1$

- Start with a range of integers from 0 to $N = 10000$. `R0 = iota(0, N)`

VIEW ADAPTORS

Pretend that you want to verify that $\sin^2(x) + \cos^2(x) = 1$

- Start with a range of integers from 0 to $N = 10000$. `R0 = iota(0, N)`
- Map the integer range to real numbers in the range $[0, 2\pi)$, i.e., perform the transformation $n \mapsto \frac{2\pi n}{N}$ over the range: `R1 = R0 | transform([](int n) -> double { return 2*pi*n/N; })`

VIEW ADAPTORS

Pretend that you want to verify that $\sin^2(x) + \cos^2(x) = 1$

- Start with a range of integers from 0 to $N = 10000$. `R0 = iota(0, N)`
- Map the integer range to real numbers in the range $[0, 2\pi)$, i.e., perform the transformation $n \mapsto \frac{2\pi n}{N}$ over the range: `R1 = R0 | transform([](int n) -> double { return 2*pi*n/N; })`
- `R2 = R1 | transform([](double x) -> double { return sin(x)*sin(x)+cos(x)*cos(x); });`

VIEW ADAPTORS

Pretend that you want to verify that $\sin^2(x) + \cos^2(x) = 1$

- Start with a range of integers from 0 to $N = 10000$. `R0 = iota(0, N)`
- Map the integer range to real numbers in the range $[0, 2\pi)$, i.e., perform the transformation $n \mapsto \frac{2\pi n}{N}$ over the range: `R1 = R0 | transform([](int n) -> double { return 2*pi*n/N; })`
- `R2 = R1 | transform([](double x) -> double { return sin(x)*sin(x)+cos(x)*cos(x); })`;
- If absolute value of any of the values in the result exceeds ϵ , we have found a counter example
`if (any_of(R2, [](auto x){return fabs(x) > eps;})) ...`

VIEW ADAPTORS

```
1  auto main() -> int {
2      namespace sr = std::ranges;
3      namespace sv = std::views;
4      const auto pi = std::acos(-1);
5      constexpr auto npoints = 10'000'00UL;
6      constexpr auto eps = 100 * std::numeric_limits<double>::epsilon();
7      auto to_0_2pi = [=](size_t idx) -> double {
8          return std::lerp(0., 2*pi, idx * 1.0 / npoints);
9      };
10     auto x_to_fx = [ ](double x) -> double {
11         return sin(x) * sin(x) + cos(x) * cos(x) - 1.0;
12     };
13     auto is_bad = [=](double x){ return std::fabs(x) > eps; };
14
15     auto res = sv::iota(0UL, npoints) | sv::transform(to_0_2pi)
16         | sv::transform(x_to_fx);
17     if (sr::any_of(res, is_bad) ) {
18         std::cerr << "The relation does not hold.\n";
19     } else {
20         std::cout << "The relation holds for all inputs\n";
21     }
22 }
```

VIEW ADAPTORS

- The job of each small transform in the previous example was small, simple, easily verified for correctness.

VIEW ADAPTORS

- The job of each small transform in the previous example was small, simple, easily verified for correctness.
- The view adaptors allow us to chain them to produce a resulting range

VIEW ADAPTORS

- The job of each small transform in the previous example was small, simple, easily verified for correctness.
- The view adaptors allow us to chain them to produce a resulting range
- Algorithms like `std::range::any_of` work on ranges, so they can work on the views resulting from chained view adaptors.

VIEW ADAPTORS

- The job of each small transform in the previous example was small, simple, easily verified for correctness.
- The view adaptors allow us to chain them to produce a resulting range
- Algorithms like `std::range::any_of` work on ranges, so they can work on the views resulting from chained view adaptors.
- No operation is done on any range when we create the variable `res` above.

VIEW ADAPTORS

- The job of each small transform in the previous example was small, simple, easily verified for correctness.
- The view adaptors allow us to chain them to produce a resulting range
- Algorithms like `std::range::any_of` work on ranges, so they can work on the views resulting from chained view adaptors.
- No operation is done on any range when we create the variable `res` above.
- When we try to access an element of the range in the `any_of` algorithm, one element is taken on the fly out of the starting range, fed through the pipeline and catered to `any_of`

VIEW ADAPTORS

- The job of each small transform in the previous example was small, simple, easily verified for correctness.
- The view adaptors allow us to chain them to produce a resulting range
- Algorithms like `std::range::any_of` work on ranges, so they can work on the views resulting from chained view adaptors.
- No operation is done on any range when we create the variable `res` above.
- When we try to access an element of the range in the `any_of` algorithm, one element is taken on the fly out of the starting range, fed through the pipeline and catered to `any_of`
- `any_of` does not process the range beyond what is necessary to establish its truth value. The remaining elements in the result array are never calculated.

Exercise 1.4:

The code used for the demonstration of view adaptors is `examples/trig_views.cc`. Build this code with GCC and Clang. If the version of your compiler does not have a usable `<ranges>` header, we can use a redirecting header `<cxx20ranges>` examples. When the compiler implements the ranges library, it includes `<ranges>`. Otherwise, it tries to include equivalent headers from the `rangev3` library. It also defines alias namespaces `sr` and `sv` for `std::ranges` and `std::std::views`. To compile, you would need to have the location of this redirecting header in your include path:

```
g++ -std=c++20 -I course_home/local/include trig_views.cc  
./a.out
```

```
clang++ -std=c++20 -stdlib=libc++ -I course_home/local/include trig_views.cc  
./a.out
```


Exercise 1.5:

The trigonometric relation we used is true, so not all possibilities are explored. In `examples/trig_views2.cc` there is another program trying to verify the bogus claim $\sin^2(x) < 0.99$. It's mostly true, but sometimes it isn't, so that our `if` and `else` branches both have work to do. The lambdas in this program have been rigged to print messages before returning. Convince yourself of the following:

- The output from the lambdas come out staggered, which means that the program does not process the entire range for the first transform and then again for the second ...
- Processing stops at the first instance where `any_of` gets a `true` answer.

VIEW ADAPTORS

```
1 // examples/gerund.cc
2 using itertype = std::istream_iterator<std::string>;
3 std::ifstream fin { argv[1] };
4 auto gerund = [] (std::string_view w) { return w.ends_with("ing"); };
5 auto in = sr::istream_view<std::string>(fin);
6 std::cout << (in | sv::filter(gerund)) << "\n";
7
```

- `sr::istream_view<T>` creates an (input) iterable range from an input stream. Each element of this range is of the type `T`.
- `sv::filter` is a view adaptor, which when applied to a range, produces another containing only the elements satisfying a given condition
- In the above, `std::cout` is shown writing out a range. This works via a separate header file included in `gerund.cc` called `range_output.hh`, which is provided to you with the course material. Ranges in C++20 are not automatically streamable to the standard output.

VIEW ADAPTORS

A program to print the alphabetically first and last word entered on the command line, excluding the program name.

```
1 // examples/views_and_span.cc
2 auto main(int argc, char* argv[]) -> int
3 {
4     if (argc < 2) return 1;
5     namespace sr = std::ranges;
6     namespace sv = std::views;
7
8     std::span args(argv, argc);
9     auto str = [] (auto cstr) -> std::string_view { return cstr; };
10    auto [mn, mx] = sr::minmax(args | sv::drop(1) | sv::transform(str));
11
12    std::cout << "Alphabetically first = " << mn << " last = " << mx << "\n";
13 }
```

FORMATTED OUTPUT

```
1  for (auto i = 0UL; i < 100UL; ++i) {  
2      std::cout << "i = " << i  
3          << ", E_1 = " << cos(i * wn)  
4          << ", E_2 = " << sin(i * wn)  
5          << "\n";  
6  }
```

```
i = 5, E_1 = 0.55557, E_2 = 0.83147  
i = 6, E_1 = 0.382683, E_2 = 0.92388  
i = 7, E_1 = 0.19509, E_2 = 0.980785  
i = 8, E_1 = 6.12323e-17, E_2 = 1  
i = 9, E_1 = -0.19509, E_2 = 0.980785  
i = 10, E_1 = -0.382683, E_2 = 0.92388  
i = 11, E_1 = -0.55557, E_2 = 0.83147
```

- While convenient and type safe and extensible, the interface of `ostream` objects like `std::cout` isn't by itself conducive to regular well-formatted output

FORMATTED OUTPUT

```
1  for (auto i = 0UL; i < 100UL; ++i) {  
2      std::cout << "i = " << i  
3          << ", E_1 = " << cos(i * wn)  
4          << ", E_2 = " << sin(i * wn)  
5          << "\n";  
6  }
```

```
i = 5, E_1 = 0.55557, E_2 = 0.83147  
i = 6, E_1 = 0.382683, E_2 = 0.92388  
i = 7, E_1 = 0.19509, E_2 = 0.980785  
i = 8, E_1 = 6.12323e-17, E_2 = 1  
i = 9, E_1 = -0.19509, E_2 = 0.980785  
i = 10, E_1 = -0.382683, E_2 = 0.92388  
i = 11, E_1 = -0.55557, E_2 = 0.83147
```

- While convenient and type safe and extensible, the interface of `ostream` objects like `std::cout` isn't by itself conducive to regular well-formatted output
- C `printf` often has a simpler path towards visually uniform columnar output, although it is neither type safe nor extensible
- C++ `<iomanip>` header allows formatting with a great deal of control, but has a verbose and inconsistent syntax

FORMATTED OUTPUT

```
1  for (auto i = 0UL; i < 100UL; ++i) {  
2      std::cout << fmt::format(  
3          "i = {:>4d}, E_1 = {:< 12.8f}, "  
4          "E_2 = {:< 12.8f}\n",  
5          i, cos(i * wn), sin(i * wn));  
6  }
```

```
i =      5, E_1 =  0.55557023 , E_2 =  0.83146961  
i =      6, E_1 =  0.38268343 , E_2 =  0.92387953  
i =      7, E_1 =  0.19509032 , E_2 =  0.98078528  
i =      8, E_1 =  0.00000000 , E_2 =  1.00000000  
i =      9, E_1 = -0.19509032 , E_2 =  0.98078528  
i =     10, E_1 = -0.38268343 , E_2 =  0.92387953  
i =     11, E_1 = -0.55557023 , E_2 =  0.83146961
```

- While convenient and type safe and extensible, the interface of `ostream` objects like `std::cout` isn't by itself conducive to regular well-formatted output
- C `printf` often has a simpler path towards visually uniform columnar output, although it is neither type safe nor extensible
- C++ `<iomanip>` header allows formatting with a great deal of control, but has a verbose and inconsistent syntax
- C++20 introduced the `<format>` header, which introduces Python like string formatting
- Based on the open source `fmt` library.

FORMATTED OUTPUT

```
1  for (auto i = 0UL; i < 100UL; ++i) {  
2      std::cout << fmt::format(  
3          "i = {:>4d}, E_1 = {:< 12.8f}, "  
4          "E_2 = {:< 12.8f}\n",  
5          i, cos(i * wn), sin(i * wn));  
6  }
```

```
i =      5, E_1 =  0.55557023 , E_2 =  0.83146961  
i =      6, E_1 =  0.38268343 , E_2 =  0.92387953  
i =      7, E_1 =  0.19509032 , E_2 =  0.98078528  
i =      8, E_1 =  0.00000000 , E_2 =  1.00000000  
i =      9, E_1 = -0.19509032 , E_2 =  0.98078528  
i =     10, E_1 = -0.38268343 , E_2 =  0.92387953  
i =     11, E_1 = -0.55557023 , E_2 =  0.83146961
```

Perfectly aligned, as all numeric output should be.

- While convenient and type safe and extensible, the interface of `ostream` objects like `std::cout` isn't by itself conducive to regular well-formatted output
- C `printf` often has a simpler path towards visually uniform columnar output, although it is neither type safe nor extensible
- C++ `<iomanip>` header allows formatting with a great deal of control, but has a verbose and inconsistent syntax
- C++20 introduced the `<format>` header, which introduces Python like string formatting
- Based on the open source `fmt` library.
- Elegant. Safe. Fast. Extensible.

FORMATTED OUTPUT

```
1 // Example of a redirecting header
2 #include <version>
3 #ifdef __cpp_lib_format
4     #include <format>
5     using std::format;
6 #elif __has_include(<fmt/format.h>)
7     #define FMT_HEADER_ONLY
8     #include <fmt/core.h>
9     using fmt::format;
10 #else
11     #error No suitable header for C++20 format!
12 #endif
```

- GCC 13 has an implementation. Our redirecting header can help us work with clang as well.
- We can use a redirecting header to use the `fmt` library when the compiler does not have the library feature
- Code simplification and compilation (and runtime) speed \implies useful to learn it. Eventually all compilers will have it.

FORMATTED OUTPUT

- `std::format("format string {}", {} etc.", args...)` takes a compile time constant format string and a parameter pack to produce a formatted output string
- `std::vformat` can be used if the format string is not known at compilation time
- If instead of receiving output as a newly created string, output into a container or string is desired, `std::format_to` or `std::format_to_n` are available
- The string contains python style placeholder braces to be filled with formatted values from the argument list
- The braces can optionally contain `id : spec` descriptors. `id` is a 0 based index to choose an argument from `args...` for that slot. `spec` controls how the value is to be written: width, precision, alignment, padding, base of numerals etc. Details of the format specifiers can be found [here](#)!

Exercise 1.6:

A simple example demonstrating the text formatting library of C++20 is in `examples/format1.cc`. When this C++20 header is not available in the standard library implementation, we use headers from the `fmt` library giving us approximately the same functionality. Although `fmt` is usually compiled to a static or shared library to link, we define the macro `FMT_HEADER_ONLY` to pretend that we got everything from the standard library.

Exercise 1.7:

The program `examples/word_count.cc` is an improved version of the word counter program from day 4. Here we clear any trailing non-alphabetic characters from strings read as words, e.g., treat "instance," as "instance". We use the ranges algorithms to clean up the string. We then use the formatting library to write the histogram.

REGULAR EXPRESSIONS USING C++20

```
1 constexpr ctll::fixed_string re{ R"xpr:^(https|http|www\.)\S*)xpr" };
2 auto urls_in_input = args | sv::drop(1)
3   | sv::transform( [=](auto inp) { return str(inp); })
4   | sv::filter([re](auto inp) { return ctre::search<re>(inp); });
5 if (auto m = ctre::match<trx>(diststr); m) {
6     auto numstr = m.get<1>().to_string();
7     // and so on...
8 }
```

- CTRE: "Compile time regular expressions", header only open source library
- Regular expressions parsed at compile time.
- Smaller binaries than `std::regex`
- Syntax makes excellent use of C++20 features for intuitive handling of regular expressions
- Compile time regex processing is possible, with great performance

REGULAR EXPRESSIONS USING CTRE

Exercise 1.8:

`examples/dist.cc` contains a rudimentary Distance class. Distances can be constructed by giving a value with a unit. Overloaded literal operators allow writing code like `auto d = 14.5_km;`. It is possible to write distances using `std::cout`, or read using `std::cin`. E.g.,

```
$ Enter distance: 13.99_cm
That is 0.1399_m

$ Enter distance: "23    km"
That is 23000_m
```

To read and interpret the input string in the correct units, we make use of regular expressions. Since these can be known at when writing the source code, we use the CTRE library to process our regular expressions. The example demonstrates many different topics explored during the course.

Modules

A PREVIEW OF C++20 MODULES

Traditionally, C++ projects are organised into header and source files. As an example, consider a simple

saxpy program ...

```
1  #ifndef SAXPY_HH
2  #define SAXPY_HH
3  #include <algorithm>
4  #include <span>
5  template <class T> concept Number = std::floating_point<T> or std::integral<T>;
6  template <class T> requires Number<T>
7  auto saxpy(T a, std::span<const T> x, std::span<const T> y, std::span<T> z){
8      std::transform(x.begin(), x.end(), y.begin(), z.begin(),
9          [a](T X, T Y) { return a * X + Y; });
10 }
11 #endif
```

```
1  #include "saxpy.hh"
2  auto main() -> int {
3      //declarations
4      saxpy(10., {inp1}, {inp2}, {outp});
5  }
```

PROBLEMS WITH HEADER FILES

- Headers contain declarations of functions, classes etc., and definitions of inline functions.
- Source files contain implementations of other functions, such as `main`.
- Since function templates and class templates have to be visible to the compiler at the point of instantiation, these have traditionally lived in headers.
- Standard library, TBB, Thrust, Eigen ... a lot of important C++ libraries consist of a lot of template code, and therefore in header files.
- The `#include <abc>` mechanism is essentially a copy-and-paste solution. The preprocessor inserts the entire source of the headers in each source file that includes it, creating large translation units.
- The same template code gets re-parsed over and over for every new translation unit.
- If the headers contain expression templates, CRTP, metaprogramming repeated processing of the templates is a waste of resources.

MODULES

- The `module` mechanism in C++20 offers a better organisation
- All code, including template code can now reside in source files
- Module source files will be processed once to produce “precompiled modules”, where the essential syntactic information has been parsed and saved.
- These compiled module interface (binary module interface) files are to be treated as caches generated during the compilation process. There are absolutely no guarantees of them remaining compatible between different versions of the same compiler, different machine configurations etc.
- Any source `import` ing the module immediately has access to the precompiled syntax tree in the precompiled module files. This leads to less overall work and faster compilation of individual translation units
- Since a source file may export a module to be imported by another source in the same project, sources must sometimes be compiled in a specific order. Automatically deducing this order is a difficult problem, and is one of the reasons tools like CMake have taken this long to support C++ modules

USING MODULES

```
1 // examples/hello_m.cc
2 import <iostream>;
3
4 auto main() -> int
5 {
6     std::cout << "Hello, world!\n";
7 }
```

- If a module is available, not much special needs to be done to use it. `import` the module instead of `#include` ing a header. Use the exported classes, functions and variables from the module.
- As of C++20, the standard library is not available as a module. But standard library headers can be imported as “header units”.

```
$ clang++ -std=c++20 -stdlib=libc++ -fmodules hello_m.cc
$ ./a.out
$ g++ -std=c++20 -fmodules-ts -xc++-system-header iostream
$ g++ -std=c++20 -fmodules-ts hello_m.cc
$ ./a.out
$
```

- GCC requires that the header units needed are first generated in a separate explicit step.
- If `iostream` had been the name of a module, we would have written `import iostream;` instead of `import <iostream>`

USING MODULES

Exercise 1.9:

Convert a few of the example programs you have seen during the course to use modules syntax instead. At the moment it means no more than replacing the `#include` lines with the corresponding `import` lines for the standard library headers. The point is to get used to the extra compilation options you need with modules at the moment. Use, for instance, the date time library functions like `feb.cc` and `advent.cc` from the day 4 examples.

CREATING A MODULE (EXAMPLE)

```
1 // saxpy.hh
2 #ifndef SAXPY_HH
3 #define SAXPY_HH
4 #include <algorithm>
5 #include <span>
6
7 template <class T>
8 concept Number = std::floating_point<T>
9                 or std::integral<T>;
10 template <Number T>
11 auto saxpy(T a, std::span<const T> x,
12            std::span<const T> y,
13            std::span<T> z)
14 {
15     std::transform(x.begin(), x.end(),
16                   y.begin(), z.begin(),
17                   [a](T X, T Y) {
18                       return a * X + Y;
19                   });
20 }
21 #endif
```

- A header file contains a function template `saxpy`, and a `concept` necessary to define that function
- A source file, `main.cc` which includes the header and uses the function

CREATING A MODULE (EXAMPLE)

```
1 // usesaxpy.cc
2 #include <iostream>
3 #include <array>
4 #include <vector>
5 #include <span>
6 #include "saxpy.hh"
7
8 auto main() -> int
9 {
10     using namespace std;
11     const array inp1 { 1., 2., 3., 4., 5. };
12     const array inp2 { 9., 8., 7., 6., 5. };
13     vector outp(inp1.size(), 0.);
14
15     saxpy(10., {inp1}, {inp2}, {outp});
16     for (auto x : outp) cout << x << "\n";
17     cout << "::::::::::::::::::::::::::\n";
18 }
```

- A header file contains a function template `saxpy`, and a `concept` necessary to define that function
- A source file, `main.cc` which includes the header and uses the function

CREATING A MODULE (EXAMPLE)

Make a module interface unit

```
1  // saxpy.hh -> saxpy.ixx
2  #ifndef SAXPY_HH
3  #define SAXPY_HH
4  #include <algorithm>
5  #include <span>
6
7  template <class T>
8  concept Number = std::floating_point<T>
9                  or std::integral<T>;
10 template <Number T>
11 auto saxpy(T a, std::span<const T> x,
12            std::span<const T> y,
13            std::span<T> z)
14 {
15     std::transform(x.begin(), x.end(),
16                   y.begin(), z.begin(),
17                   [a](T X, T Y) {
18                       return a * X + Y;
19                   });
20 }
21 #endif
```

CREATING A MODULE (EXAMPLE)

Make a module interface unit

- Include guards are no longer required, since importing a module does not transitively import things used inside the module

```
1 // saxpy.hh -> saxpy.ixx
2 #ifndef SAXPY_HH
3 #define SAXPY_HH
4 #include <algorithm>
5 #include <span>
6
7 template <class T>
8 concept Number = std::floating_point<T>
9                 or std::integral<T>;
10 template <Number T>
11 auto saxpy(T a, std::span<const T> x,
12            std::span<const T> y,
13            std::span<T> z)
14 {
15     std::transform(x.begin(), x.end(),
16                   y.begin(), z.begin(),
17                   [a](T X, T Y) {
18                       return a * X + Y;
19                   });
20 }
21 #endif
```

CREATING A MODULE (EXAMPLE)

Make a module interface unit

- Include guards are no longer required, since importing a module does not transitively import things used inside the module

```
1  // saxpy.hh -> saxpy.ixx
2
3
4  #include <algorithm>
5  #include <span>
6
7  template <class T>
8  concept Number = std::floating_point<T>
9                  or std::integral<T>;
10
11 template <Number T>
12 auto saxpy(T a, std::span<const T> x,
13            std::span<const T> y,
14            std::span<T> z)
15 {
16     std::transform(x.begin(), x.end(),
17                   y.begin(), z.begin(),
18                   [a](T X, T Y) {
19                       return a * X + Y;
20                   });
21 }
```

CREATING A MODULE (EXAMPLE)

```
1 // saxpy.hh -> saxpy.ixx
2 module;
3 #include <algorithm>
4 #include <span>
5 export module saxpy;
6
7 template <class T>
8 concept Number = std::floating_point<T>
9                 or std::integral<T>;
10 template <Number T>
11 auto saxpy(T a, std::span<const T> x,
12            std::span<const T> y,
13            std::span<T> z)
14 {
15     std::transform(x.begin(), x.end(),
16                   y.begin(), z.begin(),
17                   [a](T X, T Y) {
18                       return a * X + Y;
19                   });
20 }
```

Make a module interface unit

- Include guards are no longer required, since importing a module does not transitively import things used inside the module
- A module interface unit is a file which exports a module. The lines between the `module;` and `export module saxpy;` constitute the “global module fragment”, where traditional `#include` statements can be used.

CREATING A MODULE (EXAMPLE)

```
1 // saxpy.hh -> saxpy.ixx
2
3 export module saxpy;
4 import <algorithm>;
5 import <span>;
6
7 template <class T>
8 concept Number = std::floating_point<T>
9                 or std::integral<T>;
10 template <Number T>
11 auto saxpy(T a, std::span<const T> x,
12            std::span<const T> y,
13            std::span<T> z)
14 {
15     std::transform(x.begin(), x.end(),
16                   y.begin(), z.begin(),
17                   [a](T X, T Y) {
18                       return a * X + Y;
19                   });
20 }
```

Make a module interface unit

- Include guards are no longer required, since importing a module does not transitively import things used inside the module
- A module interface unit is a file which exports a module. The lines between the `module;` and `export module saxpy;` constitute the “global module fragment”, where traditional `#include` statements can be used.
- If you can get by with only `import` s, replace `#include` lines with corresponding `import` lines. Omit the `module;` line in this case.

CREATING A MODULE (EXAMPLE)

```
1 // saxpy.hh -> saxpy.ixx
2
3 export module saxpy;
4 import <algorithm>;
5 import <span>;
6
7 template <class T>
8 concept Number = std::floating_point<T>
9                 or std::integral<T>;
10 export template <Number T>
11 auto saxpy(T a, std::span<const T> x,
12            std::span<const T> y,
13            std::span<T> z)
14 {
15     std::transform(x.begin(), x.end(),
16                   y.begin(), z.begin(),
17                   [a](T X, T Y) {
18                       return a * X + Y;
19                   });
20 }
```

Make a module interface unit

- Include guards are no longer required, since importing a module does not transitively import things used inside the module
- A module interface unit is a file which exports a module. The lines between the `module;` and `export module saxpy;` constitute the “global module fragment”, where traditional `#include` statements can be used.
- If you can get by with only `import` s, replace `#include` lines with corresponding `import` lines. Omit the `module;` line in this case.
- Explicitly export any definitions (classes, functions...) you want for users of the module. Anything not exported by a module is automatically private to the module

CREATING A MODULE (EXAMPLE)

Use your module

```
1 // usesaxpy.cc
2 #include <iostream>
3 #include <array>
4 #include <vector>
5 #include <span>
6 #include "saxpy.hh"
7
8 auto main() -> int
9 {
10     using namespace std;
11     const array inp1 { 1., 2., 3., 4., 5. };
12     const array inp2 { 9., 8., 7., 6., 5. };
13     vector outp(inp1.size(), 0.);
14
15     saxpy(10., {inp1}, {inp2}, {outp});
16     for (auto x : outp) cout << x << "\n";
17     cout << "::::::::::::::::::::::::::\n";
18 }
```

CREATING A MODULE (EXAMPLE)

```
1 // usesaxpy.cc
2 import <iostream>;
3 import <array>;
4 import <vector>;
5 import <span>;
6 #include "saxpy.hh"
7
8 auto main() -> int
9 {
10     using namespace std;
11     const array inp1 { 1., 2., 3., 4., 5. };
12     const array inp2 { 9., 8., 7., 6., 5. };
13     vector outp(inp1.size(), 0.);
14
15     saxpy(10., {inp1}, {inp2}, {outp});
16     for (auto x : outp) cout << x << "\n";
17     cout << "::::::::::::::::::::::::::\n";
18 }
```

Use your module

- Replace `#include` lines with corresponding `import` lines. Obs: `import` lines end with a semi-colon!

CREATING A MODULE (EXAMPLE)

```
1 // usesaxpy.cc
2 import <iostream>;
3 import <array>;
4 import <vector>;
5 import <span>;
6 import saxpy;
7
8 auto main() -> int
9 {
10     using namespace std;
11     const array inp1 { 1., 2., 3., 4., 5. };
12     const array inp2 { 9., 8., 7., 6., 5. };
13     vector outp(inp1.size(), 0.);
14
15     saxpy(10., {inp1}, {inp2}, {outp});
16     for (auto x : outp) cout << x << "\n";
17     cout << "::::::::::::::::::::\n";
18 }
```

Use your module

- Replace `#include` lines with corresponding `import` lines. Obs: `import` lines end with a semi-colon!
- When importing actual modules, rather than header units, use the `module name` without angle brackets or quotes

CREATING A MODULE (EXAMPLE)

```
1 // usesaxpy.cc
2 import <iostream>;
3 import <array>;
4 import <vector>;
5 import <span>;
6 import saxpy;
7
8 auto main() -> int
9 {
10     using namespace std;
11     const array inp1 { 1., 2., 3., 4., 5. };
12     const array inp2 { 9., 8., 7., 6., 5. };
13     vector outp(inp1.size(), 0.);
14
15     saxpy(10., {inp1}, {inp2}, {outp});
16     for (auto x : outp) cout << x << "\n";
17     cout << "::::::::::::::::::::\n";
18 }
```

Use your module

- Replace `#include` lines with corresponding `import` lines. Obs: `import` lines end with a semi-colon!
- When importing actual modules, rather than header units, use the `module name` without angle brackets or quotes
- Importing `saxpy` here, only grants us access to the explicitly exported function `saxpy`. Not other functions, classes, concepts etc. defined in the module `saxpy`, not any other module imported in the module interface unit.

COMPILATION OF PROJECTS WITH MODULES

- Different compilers require different (sets of) options
- GCC:
 - Auto-detects if a file is a module interface unit (exports a module), and generates the CMI as well as an object file together.
 - No special file extension required for module interface units(Just `.cc` , `.cpp` , ...like regular source files).
 - Requires that standard library header units needed by the project are explicitly generated
 - Does not recognise module interface file extensions used by other compilers (`.ixx` , `.ccm` etc.)
 - Still rather crashy in May 2023, especially if multiple standard library headers are in use.
- Clang:
 - Provides standard library header units.
 - Comparatively stable for module based code.
 - Lots of command line options required
 - Different options to translate module interfaces depending on file extensions!
 - `.ccm` or `.cppm` : `--precompile`
 - `.ixx` : `--precompile -xc++-module`
 - `.cc` or `.cpp` : `-Xclang -emit-module-interface`
 - Separate generation of object file required
 - Module partitions not implemented

Exercise 1.10:

Versions of the `saxpy` program written using header files and then modules can be found in the `examples/modules/saxpy/`. The recipe for building is described in the README.md files in the respective sub-folders. Familiarize yourself with the process of building applications with modules. Experiment by writing a new inline function in the module interface file without exporting it. Try to call that function from `main`. Check again after exporting it in the module.

Exercise 1.11:

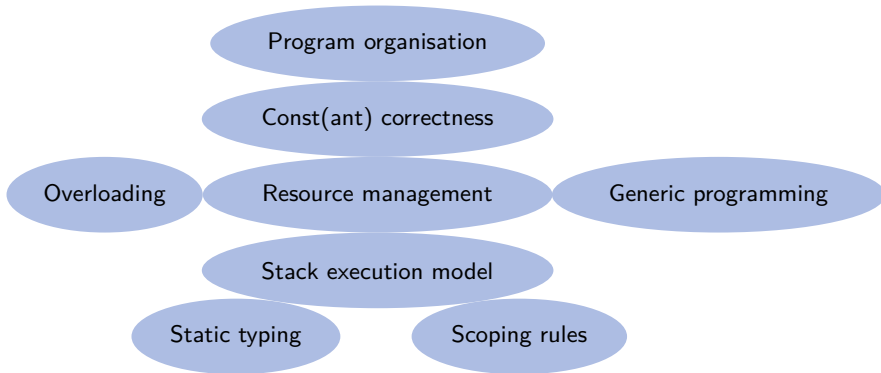
As a more complicated example, we have in `examples/modules/2_any` the second version of our container with polymorphic geometrical objects. The header and source files for each class `Point`, `Circle` etc have been rewritten for modules. Compare the two versions, build them, run them. The recipes for building are in the README.md files.

Exercise 1.12:

Each of the above exercise directories contains a `CMakeLists.txt` file. We demonstrate the use of the experimental support for modules in CMake using the combination Clang-16 and Ninja 1.11 to build these simple projects. Test it, and try using it for a different example from the course.

Closing remarks

C++ “GENES”



CLOSING REMARKS



- Most examples were simply demo code to show you how it works

CLOSING REMARKS



- Most examples were simply demo code to show you how it works
- To really internalise the ideas, you have to solve those or similar problems yourself

CLOSING REMARKS



- Most examples were simply demo code to show you how it works
- To really internalise the ideas, you have to solve those or similar problems yourself
- Information summarised for you, so that you hear it once, and then hopefully when you need it, you will remember having heard about a feature, and then look it up

CLOSING REMARKS



- Most examples were simply demo code to show you how it works
- To really internalise the ideas, you have to solve those or similar problems yourself
- Information summarised for you, so that you hear it once, and then hopefully when you need it, you will remember having heard about a feature, and then look it up
- Rapidly evolving language

CLOSING REMARKS



- Most examples were simply demo code to show you how it works
- To really internalise the ideas, you have to solve those or similar problems yourself
- Information summarised for you, so that you hear it once, and then hopefully when you need it, you will remember having heard about a feature, and then look it up
- Rapidly evolving language
- en.cppreference.com

CLOSING REMARKS



- Most examples were simply demo code to show you how it works
- To really internalise the ideas, you have to solve those or similar problems yourself
- Information summarised for you, so that you hear it once, and then hopefully when you need it, you will remember having heard about a feature, and then look it up
- Rapidly evolving language
- en.cppreference.com
- isocpp.org

CLOSING REMARKS



- Most examples were simply demo code to show you how it works
- To really internalise the ideas, you have to solve those or similar problems yourself
- Information summarised for you, so that you hear it once, and then hopefully when you need it, you will remember having heard about a feature, and then look it up
- Rapidly evolving language
- en.cppreference.com
- isocpp.org
- YouTube channel: Jason Turner's C++ weekly

CLOSING REMARKS



- Most examples were simply demo code to show you how it works
- To really internalise the ideas, you have to solve those or similar problems yourself
- Information summarised for you, so that you hear it once, and then hopefully when you need it, you will remember having heard about a feature, and then look it up
- Rapidly evolving language
- en.cppreference.com
- isocpp.org
- YouTube channel: Jason Turner's C++ weekly
- YouTube channel: CppCon conference talks

SOME USEFUL NON-STANDARD LIBRARIES

- `Lyra` command line processing
- `CTRE` Compile time regular expressions
- `xtensor` (the entire `xtensor-stack`) : C++ SIMD wrappers, numpy style multi-dimensional array operations...

XTENSOR: MULTI-DIMENSIONAL ARRAYS WITH LAZY EVALUATION

```
1 np.linspace(0., 2., 10)
2 np.logspace(1., 10., 4)
3 np.zeros(10, 10)
4 A[1,2]
5 A.flat(4)
6 A[:,3]
7 A[:3, 3:]
8 np.vectorize(f)
9 A[A > 1.0]
10 A[[1,2], [0,1]]
11 np.random.rand(100,200)
12 np.random.shuffle(A)
13 np.where(a < 0, a , b)
14 np.loadtxt(file, delim)
15 np.linalg.svd(a)
16 np.linalg.eig(a)
```

- Syntax modelled after python numpy
- Sometimes more lazy evaluations

```
1 xt::linspace<double>(0., 2., 10UL);
2 xt::logspace<double>(2., 10., 4UL);
3 xt::zeros<double>({10UL, 10UL});
4 A(1,2);
5 A[4];
6 xt::col(A, 3) or xt::view(A, xt::all, 3);
7 xt::view(A, xt::range(_, 3), xt::range(3,_));
8 xt::vectorize(f);
9 xt::filter(A, A > 1.0);
10 xt::index_view(A, {{1,2}, {0,1}});
11 xt::random::rand<double>({100, 200});
12 xt::random::shuffle(A);
13 xt::where(A < 0, A, B);
14 xt::load_csv<double>(stream);
15 xt::linalg::svd(A);
16 xt::linalg::eig(A);
```

```
1  #include <xtensor/xtensor.hpp>
2  #include <xtensor/xarray.hpp>
3  #include <xtensor/xio.hpp>
4  #include <xtensor/xrandom.hpp>
5  #include <xtensor-blas/xlinalg.hpp>
6  #include <iostream>
7
8  auto main() -> int
9  {
10     auto R = xt::random::rand<double>({4, 4});
11     auto eigs = xt::linalg::eigvals(R);
12     std::cout << R << "\n\n";
13     std::cout << eigs << "\n";
14 }
```

Exercise 1.13:

The short program `examples/xt0.cc` demonstrates using `xtensor` with eigen value evaluation. The linear algebra functionality in `xtensor` is currently handled by an external project `xtensor-blas`, which offloads some of the work to a blas library. To build the program, set the include path to include headers from “xtensor-stack”, i.e., `xtl`, `xtensor`, `xsimd`, `xtensor-io` and `xtensor-blas`. They can be given a common installation prefix. For includes, use `$swhome/include` and for linking,

```
-L $swhome/lib64 -lopenblas -lpthread -lgfortran
```

LYRA: COMMAND LINE PROCESSING

```
1  auto cli = lyra::help(showhelp)
2      | lyra::opt(N, "N_samples")["-N"]["--number-of-samples"]
3          ("The number of samples you want to generate")
4      | lyra::opt(mean, "mu")["-m"]["--mean"]("The mean of the distribution")
5      | lyra::opt(stdv, "sigma")["-s"]["--standard-deviation"]("Standard deviation");
6
7  auto cli_good = cli.parse({ argc, argv });
8
9  if (not cli_good) {
10      std::cerr << "Error in command line: " << cli_good.errorMessage() << "\n";
11      return 1;
12  }
13
14  if (showhelp or argc == 1) {
15      std::cout << cli;
16      return 0;
17  }
18
```

CTRE: COMPILE TIME REGULAR EXPRESSIONS

```
1     constexpr ctll::fixed_string re{ R"xpr^(https:|http:|www\.)\S*xpr" };
2
3     auto urls_in_input = args | sv::drop(1)
4         | sv::transform( [=](auto inp) { return str(inp); })
5         | sv::filter([re](auto inp) { return ctre::search<re>(inp); });
6
```

- Regular expressions parsed at compile time.
- Smaller binaries than `std::regex`