

Zero-Cost Abstractions in C++20

Vincent Reverdý

[vincent.reverdý@lapp.in2p3.fr]

Researcher in Computer Science and Numerical Cosmology
CNRS - French National Centre for Scientific Research
LAPP - Laboratoire d'Annecy de Physique des Particules

July 11th, 2023



Table of contents

1 Code complexity

2 Software Architecture

3 Concept-based programming

4 Standardization

5 Advanced C++

6 Conclusions

Code complexity

- 1 Code complexity
- 2 Software Architecture
- 3 Concept-based programming
- 4 Standardization
- 5 Advanced C++
- 6 Conclusions

Why do we want better and better supercomputers?

Keep the same physics

- Better resolution
- Better accuracy
- Better statistics

Keep the same resolution

- Improve physical modeling
- Multiphysics

Usual code limitations

Runtime performance

- Limitations on execution time or energy consumption

Memory

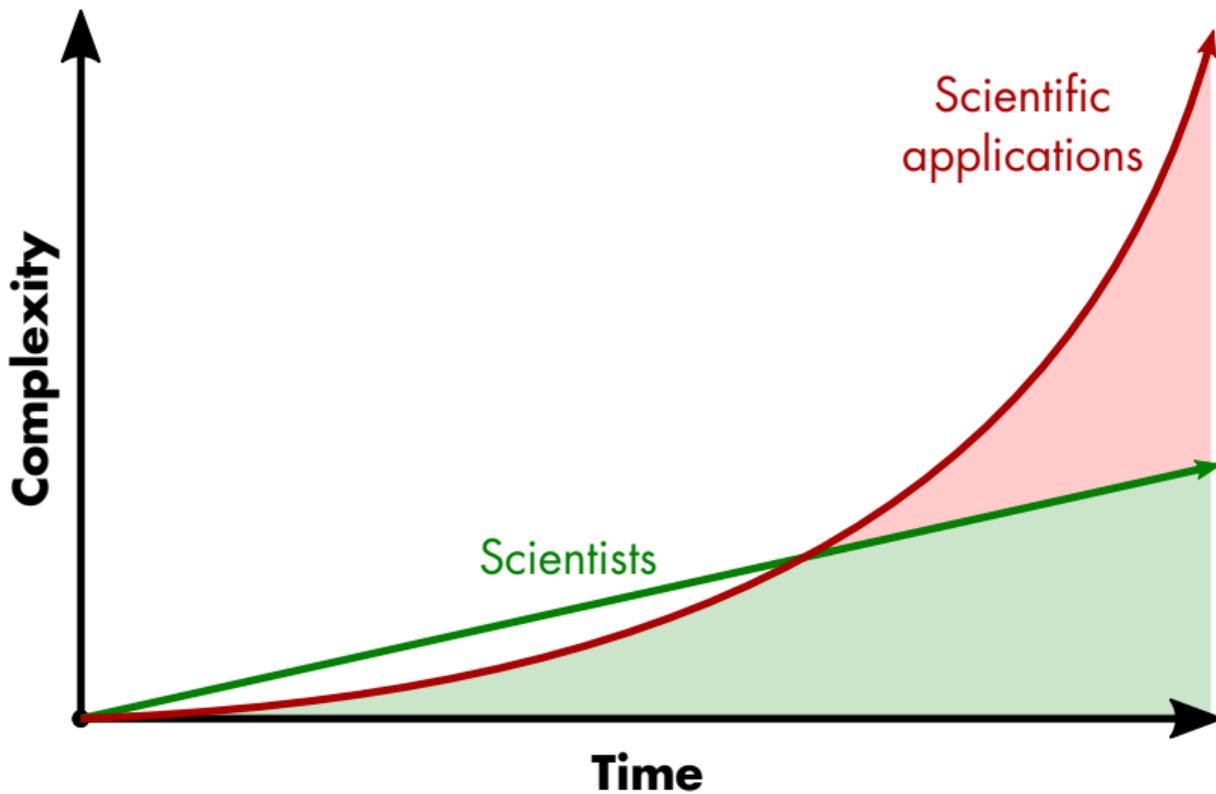
- Limitations on memory usage

The most problematic code limitation

Structural code complexity

- Codes only exist if humans can write them in the first place

The wall of software complexity



What was special about this game?



The role of abstraction

Complexity reduction

$$\mathcal{C} = \mathcal{O} \left(\prod_i \alpha_i \right) \quad \Rightarrow \quad \mathcal{C} = \mathcal{O} \left(\sum_i \alpha_i \right)$$

- \mathcal{C} : structural complexity
- i : concept
- α_i : number of instances of that concept

The critical role of software architecture



Software Architecture

- 1 Code complexity
- 2 Software Architecture**
- 3 Concept-based programming
- 4 Standardization
- 5 Advanced C++
- 6 Conclusions

Starting from an example

A navigation code used to actually fly airplanes

```
1
2 void xY_Brg_Rng(double X_1, double Y_1, double X_2, double Y_2, double *Bearing, double *Range);
3
4 void DistanceBearing(double lat1, double lon1,
5                     double lat2, double lon2,
6                     double *Distance, double *Bearing);
7
8 double DoubleDistance(double lat1, double lon1,
9                      double lat2, double lon2,
10                     double lat3, double lon3);
11
12 void FindLatitudeLongitude(double Lat, double Lon,
13                           double Bearing, double Distance,
14                           double *lat_out, double *lon_out);
15
16 double CrossTrackError(double lon1, double lat1,
17                       double lon2, double lat2,
18                       double lon3, double lat3,
19                       double *lon4, double *lat4);
20
21 double ProjectedDistance(double lon1, double lat1,
22                         double lon2, double lat2,
23                         double lon3, double lat3,
24                         double *xtd, double *crs);
25
26 void LatLon2Flat(double lon, double lat, int *scx, int *scy);
```

A few guiding principles

Small functions

- Write small functions when possible (less than 30 lines)

Few parameters

- Try to minimize the number of parameters (less than 4 most of the time)

Keep the same pattern

- Keep the same pattern of parameters for similar functions

Type everything!

- Encode as much information as possible in types

Bikeshedding

- Finding good names for things is hard, but critical

Core idea

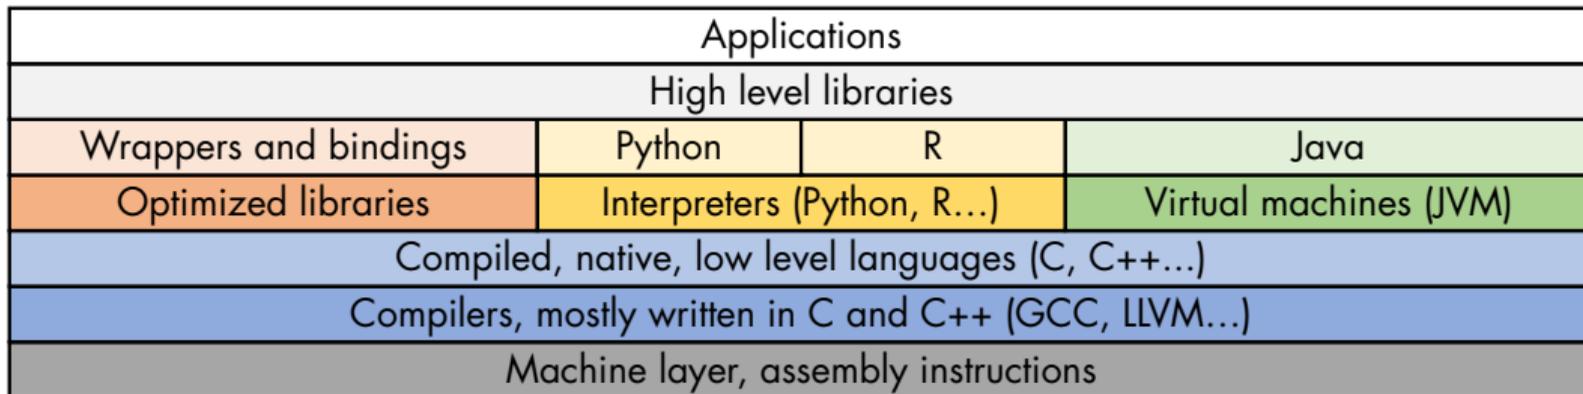
Structure matters

- Encode application domain information in the structure of the program itself

Aside note

- Programs can be seen as mathematical structures on which mathematical metrics can be computed (eg: the abstract shape or the topology of a program)

Software stack



The GPE Principle

The Holy Triad

- Genericity
- Performance
- Expressivity

Genericity: Optimize for the library's author lines of code

- How many special cases can I cover with my code?

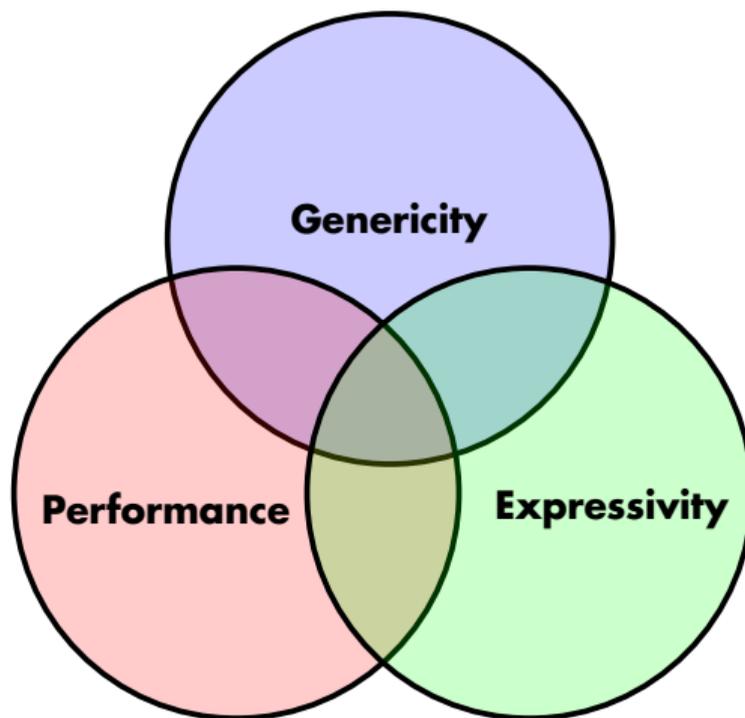
Performance: Optimize for runtime

- How fast my code is?

Expressivity: Optimize for the library's user lines of code

- How much can I express in a single line of code?

The GPE Principle



Concept-based programming

- 1 Code complexity
- 2 Software Architecture
- 3 Concept-based programming**
- 4 Standardization
- 5 Advanced C++
- 6 Conclusions

Software Architecture in C++20

Concept and constraints

- Concepts as constrained generic programming
- A way to formalize and specify abstractions and software architecture in C++

Concepts and constraints in C++20

Concept

- Named set of requirements
- Must appear at namespace scope

```
1 template < /*template-parameter-list*/ >  
2 concept /*concept-name*/ = /*constraint-expression*/;
```

Constraints

- Sequence of logical operations and operands
- Requirements on template arguments
- 3 types: conjunctions / disjunctions / atomic constraints

Example of concepts

A simple arithmetic concept

```
1 // Concept definition
2 template <class T>
3 concept arithmetic = std::is_arithmetic_v<T>;
4
5 // Constrained function v1
6 template <arithmetic T>
7 void print_v1(T x) {
8     std::cout << x << std::endl;
9 }
10
11 // Constrained function v2
12 template <class T>
13 requires arithmetic<T>
14 void print_v2(T x) {
15     std::cout << x << std::endl;
16 }
17
18 // Constrained function v3
19 void print_v3(arithmetic auto x) {
20     std::cout << x << std::endl;
21 }
```

Example of constraints

Constraints on addability

```
1  template <class T>
2  concept addable = requires {std::declval<T>() + std::declval<T>()};
3
4  template <class T>
5  concept addable = requires (T x) {x + x;};
6
7  template <class T1, class T2>
8  concept addable2 = requires (T1 x, T2 y) {x + y;};
9
10 template <class T>
11 concept addable_and_multiplicable = addable<T> && requires (T x) {x * x;};
12
13 template <class T>
14 requires requires (T x) {x + x;}
15 T add(T x, T y) {
16     return x + y;
17 }
```

Subsumption of concepts (1/3)

Subsumption as partial ordering

```
1  template <class T>
2  concept addable = requires (T x) {x + x;};
3
4  template <class T>
5  concept shiftable = requires (T x) {x << 1;};
6
7  template <class T>
8  concept addable_and_shiftable = addable<T> && shiftable<T>;
9
10 template <addable T>
11 void f(T x) {std::cout << "only addable" << std::endl;}
12
13 template <addable_and_shiftable T>
14 void f(T x) {std::cout << "addable and shiftable" << std::endl;}
15
16 f(5.1); // only addable
17 f(5); // addable and shiftable
```

Subsumption of concepts (2/3)

Works only with concepts

```
1  template <class T, class = void>
2  struct is_addable: std::false_type {};
3  template <class T>
4  struct is_addable<T, std::void_t<decltype(
5      std::declval<T>() + std::declval<T>())
6  >>: std::true_type {};
7  template <class T>
8  inline constexpr bool is_addable_v = is_addable<T>::value;
9
10 template <class T, class = void>
11 struct is_shiftable: std::false_type {};
12 template <class T>
13 struct is_shiftable<T, std::void_t<decltype(
14     std::declval<T>() << std::declval<std::size_t>())
15 >>: std::true_type {};
16 template <class T>
17 inline constexpr bool is_shiftable_v = is_shiftable<T>::value;
```

Subsumption of concepts (3/3)

Works only with concepts

```
1  template <class T>
2  concept addable = is_addable_v<T>;
3
4  template <class T>
5  concept shiftable = is_shiftable_v<T>;
6
7  template <class T>
8  concept addable_and_shiftable = is_addable_v<T> && is_shiftable_v<T>;
9
10 template <addable T>
11 void f(T x) {std::cout << "only addable" << std::endl;}
12
13 template <addable_and_shiftable T>
14 void f(T x) {std::cout << "addable and shiftable" << std::endl;}
15
16 f(5.1); // ambiguous
17 f(5); // ambiguous
```

C++ concepts vs Rust traits

C++ concepts

- Structural typing
- A type may accidentally satisfy a concept
- No coupling between concepts (architecture) and types (implementation)
- Concepts are optional
- Constraints work on allowed expression for the whole language
- Subsumption and logical expressions of constraints (`&&`, `||`, `!`)

Rust traits

- Nominal typing
- `impl Trait for Type` explicitly indicates that a type satisfy a trait
- Coupling between traits (architecture) and types (implementation)
- Traits are mandatory
- Traits can only check for a subset of the language

C++ concepts can do nominal typing

Nominal typing implementation

```
1  template <class Trait>
2  struct implements_trait {};
3
4  template <class T, class Trait, class = void>
5  struct is_implementing: std::false_type {};
6
7  template <class T, class Trait>
8  struct is_implementing<T, Trait, std::enable_if_t<
9      std::is_base_of_v<implements_trait<Trait>, T>
10 >>: std::true_type {};
11
12 template <class T, class Trait>
13 concept implements = is_implementing<T, Trait>::value;
14
15 struct mytrait {};
16
17 struct mytype: implements_trait<mytrait> {};
18
19 template <implements<mytrait> T>
20 void f(T x) {}
21
22 f(mytype{}); // OK
23 f(3); // ERROR
```

Combining concepts and constraints with if constexpr

With an external concept

```
1 template <class T>
2 concept shiftable = requires {std::declval<T>() << std::declval<int>()};
3
4 template <class T>
5 void is_shiftable() {
6     if constexpr (shiftable<T>) {std::cout << "shiftable" << std::endl;}
7     else {std::cout << "not shiftable" << std::endl;}
8 }
9
10 // is_shiftable<int>() -> "shiftable"
11 // is_shiftable<double>() -> "not shiftable"
```

With an inline requires clause

```
1 template <class T>
2 void is_shiftable() {
3     if constexpr (requires {std::declval<T>() << std::declval<int>()};) {std::cout << "shiftable" << std::endl;}
4     else {std::cout << "not shiftable" << std::endl;}
5 }
```

With an inline requires clause with a parameter list

```
1 template <class T>
2 void is_shiftable() {
3     if constexpr (requires (T x, int y) {x << y;}) {std::cout << "shiftable" << std::endl;}
4     else {std::cout << "not shiftable" << std::endl;}
5 }
```

Checking if a function exists

The traditional way: the preprocessor

```
1 #ifdef __SUPPORTS_THEFUNCTION
2 /* Doing something here */
3 #endif
```

The metaprogramming way

```
1 //void thefunction(int x);
2
3 template <class T, class = decltype(thefunction(std::declval<T>()))>
4 std::true_type supports_thefunction_for(T);
5 template <class T, class... X>
6 std::false_type supports_thefunction_for(T, X...);
7
8 inline constexpr bool supports_thefunction
9 = decltype(supports_thefunction_for(std::declval<int>()))::value;
10 // true if thefunction(int x) is active
11 // false if thefunction(int x) is commented out
```

The concept-based way

```
1 //void thefunction(int x);
2
3 template <class T = int>
4 concept supports_thefunction_for
5 = requires (T x) {thefunction(x);};
```

Checking if a function exists: forcing template dependency

Leveraging alias templates

```
1 //void thefunction(int x);
2
3 // The concept checks for a particular type provided by the user
4 template <class T>
5 concept supports_thefunction_for
6 = requires (T x) {thefunction(x);};
7
8 // Alias template keeping only the first type
9 template <class T, class...>
10 using first_type = T;
11
12 // The concept ignores its template parameter and tests only the relevant type
13 template <class... Dummy>
14 concept supports_thefunction
15 = requires {thefunction(std::declval<first_type<int, Dummy...>>());};
```

Constexpr if and requires clauses

The problem of undefined symbols

```
1 /*
2 void thefunction(int) {
3     std::cout << "thefunction" << std::endl;
4 }*/
5
6 template <class T>
7 void check(T x) {
8     if constexpr (requires (T y) {thefunction(y);} ) {
9         thefunction(x); // OK
10        thefunction(3); // ERROR
11        []<class U>(U x){thefunction(x);} (3); // OK
12    } else {
13        std::cout << "not thefunction" << std::endl;
14    }
15 }
16
17 check(1); //
```

Coming back to the problem of printing

Better than `std::enable_if`

```
1 // For numbers
2 template <printable T>
3 void print(const T& x) {
4     std::cout << x << std::endl;
5 }
6
7 // For container of numbers
8 template <range R>
9 requires printable<decltype(*std::begin(std::declval<R>()))>
10 void print(const R& range) {
11     for (auto it = std::begin(range); it != std::end(range); ++it) {
12         std::cout << *it << " ";
13     } std::cout << std::endl;
14 }
```

Software architecture with concepts

Solves the problems of metaprogramming-based approaches

- Easy to read
- Easy to implement
- Nice error messages

Contrast with Object Oriented Programming

- Types are not stuck in a fixed hierarchy
- Types come first, abstractions second
- No runtime overhead, pure compile-time check

Important notes

- A way to guide the compiler in the compilation process
- Bottom-up approach
- Designing concepts can be crazy hard

Standardization

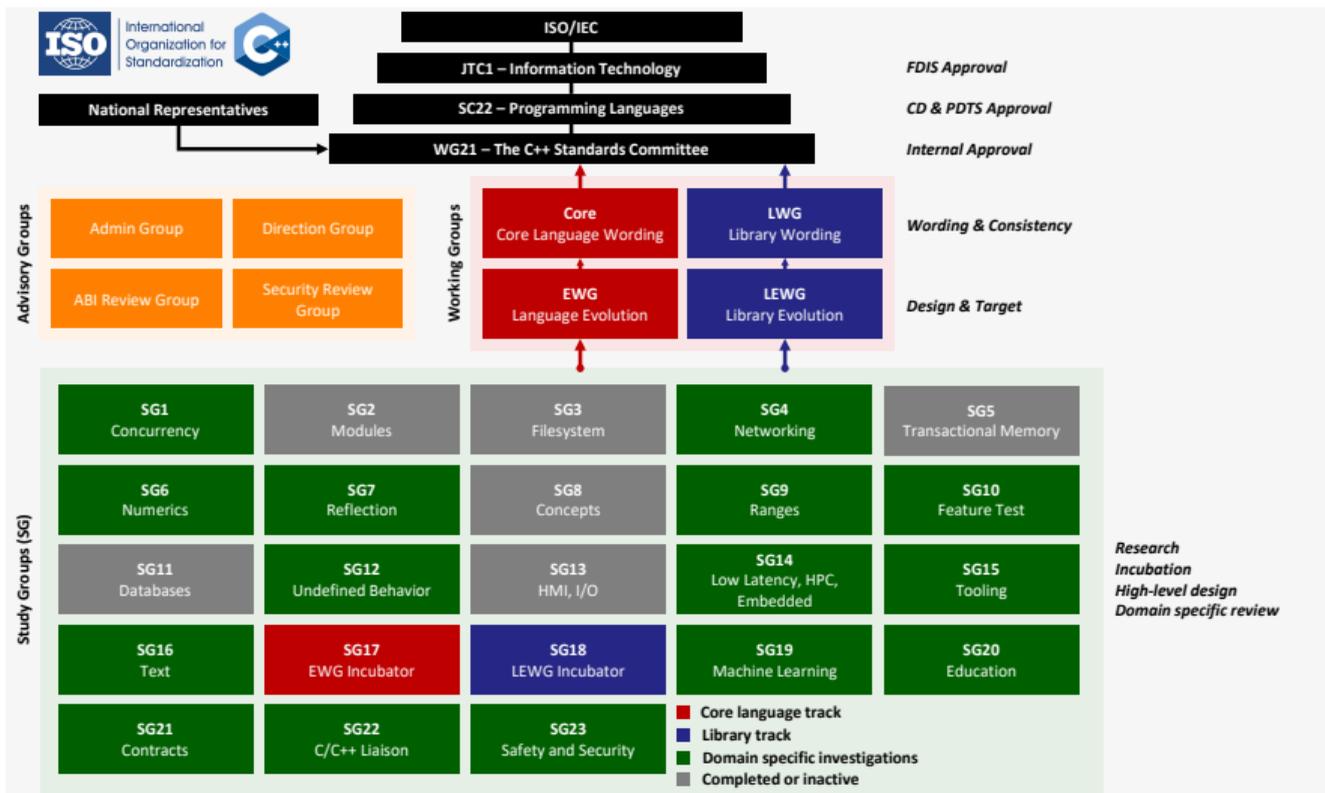
- 1 Code complexity
- 2 Software Architecture
- 3 Concept-based programming
- 4 Standardization**
- 5 Advanced C++
- 6 Conclusions

The C++ Standard

The standard

- [Link: N4950](#)
- Only a specification, not an implementation

The C++ Standards Committee



Going beyond

Why is it so complicated to standardize anything?

- Backward compatibility
- No ABI break
- Bad past decisions
- Insane levels of requirements
- Achieving Genericity, Performance, and Expressivity at the same time

Better than the standard library

- Still possible to do better than the standard and the standard library
- Implementers are not doing black magic and do not have infinite resources

An improved tuple: the overloaded log-tuple trick

Straightforward approach

`get<N>(tuple)` has to iterate over the first N types.

Advanced approach

There is a way to exploit overload resolution to have $\mathcal{O}(\log(N))$ compile-time access.

Indexing

```
1 // Index constant type
2 template <std::size_t I>
3 struct index_constant: std::integral_constant<std::size_t, I> {};
4
5 // Index constant variable template
6 template <std::size_t I>
7 inline constexpr index_constant<I> index = {};
```

Log-tuple trick: elements

Element wrappers

```
1 // A basic element wrapper
2 template <class T>
3 struct tuple_element_wrapper {
4     constexpr tuple_element_wrapper(const T& x): value(x) {}
5     // Other constructors to be defined
6     T value;
7 };
8
9 // An indexed tuple element
10 template <std::size_t I, class T>
11 struct tuple_element: tuple_element_wrapper<T> {
12     constexpr tuple_element(const T& x): tuple_element_wrapper<T>(x) {}
13     constexpr T& operator[](index_constant<I>) {
14         return static_cast<wrapper<T>&>(*this).value;
15     }
16     constexpr const T& operator[](index_constant<I>) const {
17         return static_cast<const wrapper<T>&>(*this).value;
18     }
19 };
```

Log-tuple trick: tuple

Tuple

```
1 // Base class declaration
2 template <class Sequence, class... T>
3 struct tuple_base;
4
5 // Base class specialization for index sequence
6 template <std::size_t... I, class... T>
7 struct tuple_base<std::index_sequence<I...>, T...>
8 : tuple_element<I, T>... {
9     using index_sequence = std::index_sequence<I...>;
10    using tuple_element<I, T>::operator[]...;
11    constexpr tuple_base(const T&... x): tuple_element<I, T>(x)... {}
12    // Other constructors to be defined
13 };
14
15 // Actual tuple implementation
16 template <class... T>
17 struct tuple: tuple_base<std::index_sequence_for<T...>, T...> {
18     using base = tuple_base<std::index_sequence_for<T...>, T...>;
19     using base::base;
20     using base::operator[];
21 };
22 template <class... T>
23 tuple(const T&...) -> tuple<T...>;
```

Result

mytuple[index<3>] leverages overload resolution to access the element at compile-time.

Taking C++ to another level of genericity

Genericity in C++

- C++ is type-generic but NOT kind-generic

C++ is type-generic

```
1 template <class T>
2 struct wrapper {};
3
4 template <>
5 struct wrapper<int> {};
6
7 template <>
8 struct wrapper<double> {};
```

C++ is NOT kind-generic

```
1 template <class T>
2 struct wrapper1 {};
3
4 template <auto X>
5 struct wrapper2 {};
6
7 template <template <class...> class F>
8 struct wrapper3 {};
```

Problem with higher-order metafunctions

The problem

```
1 // Metafunction hierarchy
2 template <class T>
3 struct metafunction_wrapper_0 {};
4 template <template <class...> class F>
5 struct metafunction_wrapper_1 {};
6 template <template <template <class...> class...> class F>
7 struct metafunction_wrapper_2 {};
8 template <template <template <template <class...> class...> class...> class F>
9 struct metafunction_wrapper_3 {};
10 template <template <template <template <template <class...> class...>>>> class...> class...> class F>
11 struct metafunction_wrapper_4 {};
12
13 // Use cases
14 metafunction_wrapper_1<metafunction_wrapper_0> x1; // OK
15 metafunction_wrapper_2<metafunction_wrapper_1> x2; // OK
16 metafunction_wrapper_3<metafunction_wrapper_2> x3; // OK
17 metafunction_wrapper_4<metafunction_wrapper_3> x4; // OK
```

Proposal

- Currently no way of collapsing the hierarchy
- Introducing a new mechanism to make C++ kind-generic

Advanced C++

- 1 Code complexity
- 2 Software Architecture
- 3 Concept-based programming
- 4 Standardization
- 5 Advanced C++**
- 6 Conclusions

Symbolic calculus in C++

Unique identifiers for symbols

```
1 template <class T>
2 struct symbol_id {
3     static constexpr auto singleton = []{};
4 };
```

Symbol definition

```
1 template <class T = void, auto Id = symbol_id<decltype([]{})>{}>
2 struct symbol {
3     static constexpr auto symbol_id = Id;
4 };
```

Unique symbols

Symbol definition

```
1 template <class T = void, auto Id = symbol_id<decltype([]{})>{}>
2 struct symbol {
3     static constexpr auto symbol_id = Id;
4 };
```

In practice

```
1 int main() {
2     symbol x;
3     symbol y;
4     std::cout << std::is_same_v<decltype(x), decltype(y)> << std::endl; // 0
5 }
```

Need a total ordering on symbol identifiers

Modifying the definition of the symbol identifiers

```
1  template <class T>
2  struct symbol_id {
3      static constexpr auto singleton = []{};
4      static constexpr const void* address = static_cast<const void*>(&singleton);
5  };
6
7  template <class Lhs, class Rhs>
8  constexpr std::strong_ordering operator<=>(symbol_id<Lhs>, symbol_id<Rhs>) {
9      // Using the standard function object that defines a total order on pointers
10     return std::compare_three_way{}(
11         symbol_id<Lhs>::address,
12         symbol_id<Rhs>::address
13     );
14 }
15
16 int main() {
17     symbol x;
18     symbol y;
19     std::cout << (x.id < y.id) << std::endl;
20 }
```

Introducing some concepts

Specializable type traits

```
1 template <class>
2 struct is_symbolic: std::false_type {};
3
4 template <class T, auto Id>
5 struct is_symbolic<symbol<T, Id>>: std::true_type {};
6
7 template <class T>
8 inline constexpr bool is_symbolic_v = is_symbolic<T>::value;
```

Symbolic concept

```
1 template <class T>
2 concept symbolic = is_symbolic_v<T>;
```

Symbolic operators

Assignment

```
1 struct assignment_operator {
2     template <class Rhs, class Lhs>
3     constexpr decltype(std::declval<Rhs>() = std::declval<Lhs>())
4     operator()(Rhs&& rhs, Lhs&& lhs)
5     noexcept(noexcept(std::forward<Rhs>(rhs) = std::forward<Lhs>(lhs))) {
6         return std::forward<Rhs>(rhs) = std::forward<Lhs>(lhs);
7     }
8 };
```

Addition

```
1 struct addition_operator {
2     template <class Rhs, class Lhs>
3     constexpr decltype(std::declval<Rhs>() + std::declval<Lhs>())
4     operator()(Rhs&& rhs, Lhs&& lhs)
5     noexcept(noexcept(std::forward<Rhs>(rhs) + std::forward<Lhs>(lhs))) {
6         return std::forward<Rhs>(rhs) + std::forward<Lhs>(lhs);
7     }
8 };
```

Symbolic expressions

Expressions

```
1 template <class... Args>
2 struct symbolic_expression {
3 };
4
5 template <class... Args>
6 struct is_symbolic<symbolic_expression<Args...>>: std::true_type {};
7
8 template <symbolic Lhs, symbolic Rhs>
9 constexpr symbolic_expression<decltype(
10     []()-> std::tuple<operator_symbol<assignment_operator>, Lhs, Rhs>{return {};}
11 )>
12 operator+(Lhs, Rhs) noexcept {
13     return {};
14 }
```

Expression templates are not dead

Application

```
1 int main(int argc, char* argv[]) {  
2     symbol x;  
3     symbol y;  
4     symbol z;  
5     auto f = x + y + z; // Contains the AST  
6 }
```

Full symbolic language with AST manipulation

Basic application

```
1 int main(int argc, char* argv[]) {
2     // Real symbols
3     symbol<real> a;
4     symbol<real> b;
5     symbol<real> c;
6     symbol<real> d;
7
8     // Symbolic function
9     auto f = (a + b) * (c + d);
10
11    // Computation
12    f(a = 5., b = 13., c = 50., d = 12.)
13 }
```

For linear algebra

With matrices

```
1 int main(int argc, char* argv[]) {
2     // Real symbols
3     symbol<matrix<real>> a;
4     symbol<matrix<real>> b;
5     symbol<matrix<real>> c;
6
7     // Symbolic function
8     auto f = (a + b) * c;
9
10    // Computation
11    f(
12        a = std::mdspan(...),
13        b = std::mdspan(...),
14        c = std::mdspan(...)
15    );
16 }
```

Going beyond

A full symbolic language

- AST manipulation (simplification, ...)
- Solving equations
- Expressing parallelism
- Custom optimizer

Take-home lesson

Expressivity

Start from what users would like to write

Another example

```
std::ndarray<double, shape[4] () [3] [5]> myarray;
```

Conclusions

- 1 Code complexity
- 2 Software Architecture
- 3 Concept-based programming
- 4 Standardization
- 5 Advanced C++
- 6 **Conclusions**

Conclusion



Architecture is important. Please abstract things. Thanks.

Thank you for your attention

Any question?