



MODULES ARE COMING

Bryce Adelstein Lelbach

CUDA C++ Core Libraries Lead



@blelbach

ISO C++ Library Evolution Incubator Chair, ISO C++ Tooling Study Group Chair

#include <C++>



MODULES ARE COMING

Bryce Adelstein Lelbach

CUDA C++ Core Libraries Lead



@blelbach

ISO C++ Library Evolution Incubator Chair, ISO C++ Tooling Study Group Chair

Modules are:

- A new compilation model for C++.
- A new way to organize C++ projects.

Textual Inclusion

math.hpp

```
#pragma once  
  
int square(int a);
```

math.cpp

```
#include "math.hpp"  
  
int square(int a) { return a * a; }
```

main.cpp

```
#include "math.hpp"  
  
int main() { return square(42); }
```

Textual Inclusion

math.hpp

```
#pragma once  
  
int square(int a);
```

math.cpp

```
#include "math.hpp"  
  
int square(int a) { return a * a; }
```

main.cpp

```
#include "math.hpp"  
  
int main() { return square(42); }
```

Modular Import

math.cppm

```
export module math;  
  
export int square(int a);
```

math.cpp

```
module math;  
  
int square(int a) { return a * a; }
```

main.cpp

```
import math;  
  
int main() { return square(42); }
```

Modules will have a greater impact than any other feature added post C++-98.

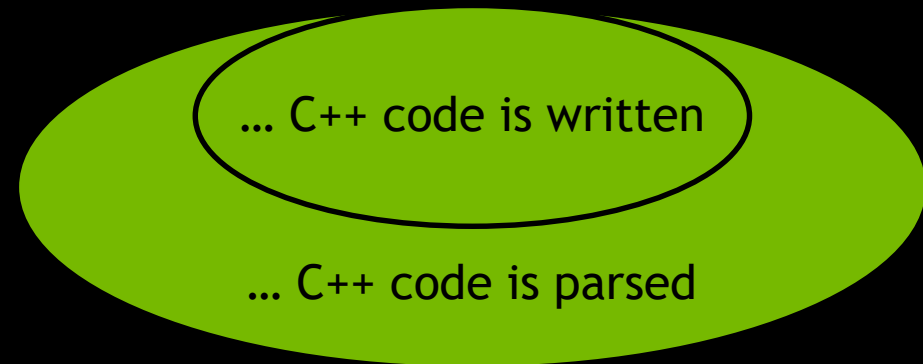
C++11's <thread> changes how...

... C++ code is written

C++11's lambdas change how...

... C++ code is written

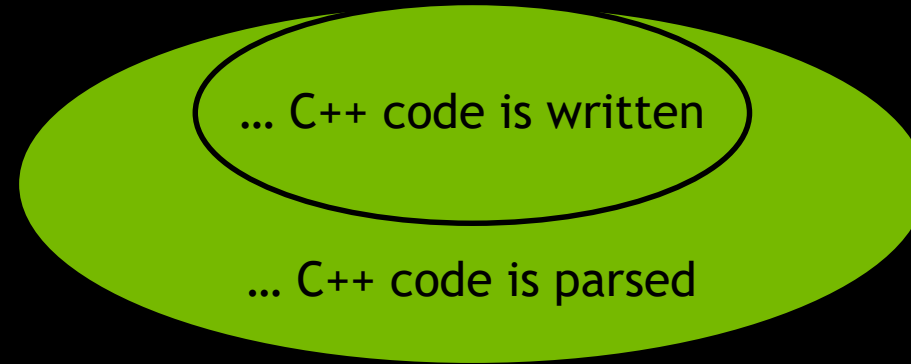
C++11's lambdas change how...




C++20's modules change how...

... C++ code is written

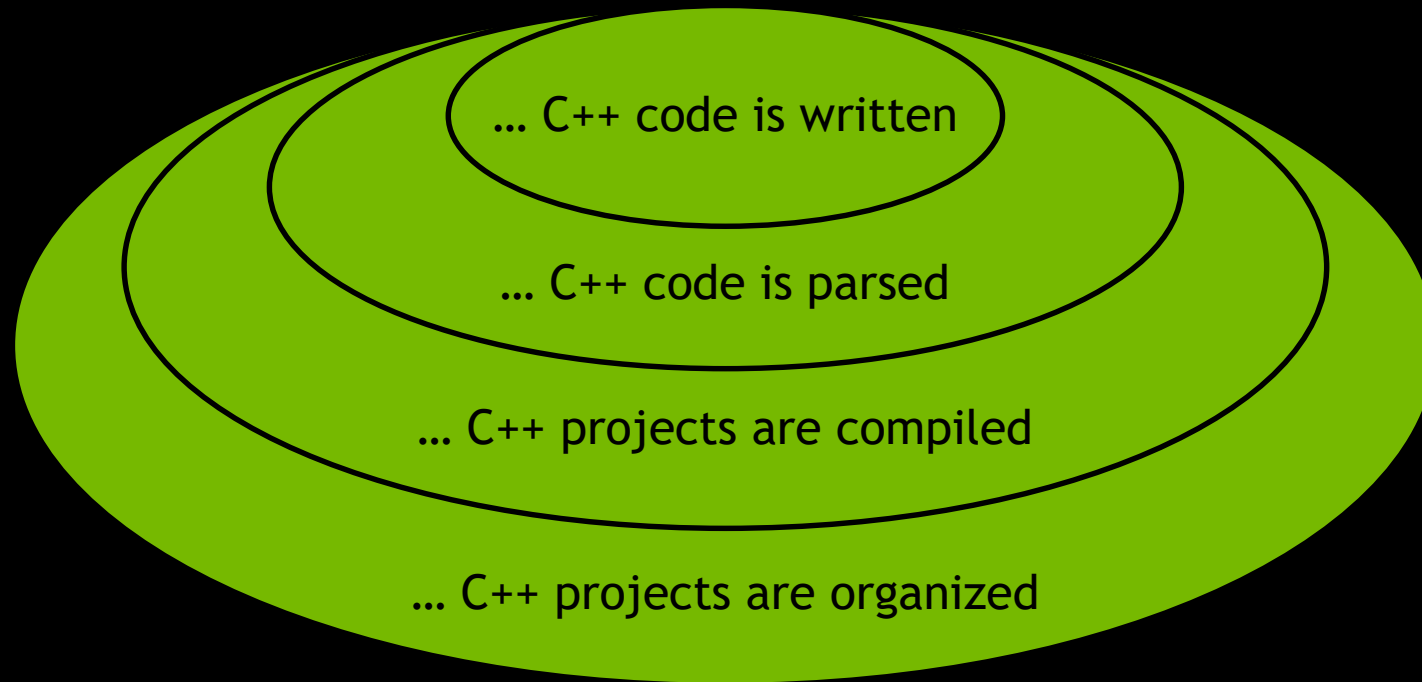
C++20's modules change how...



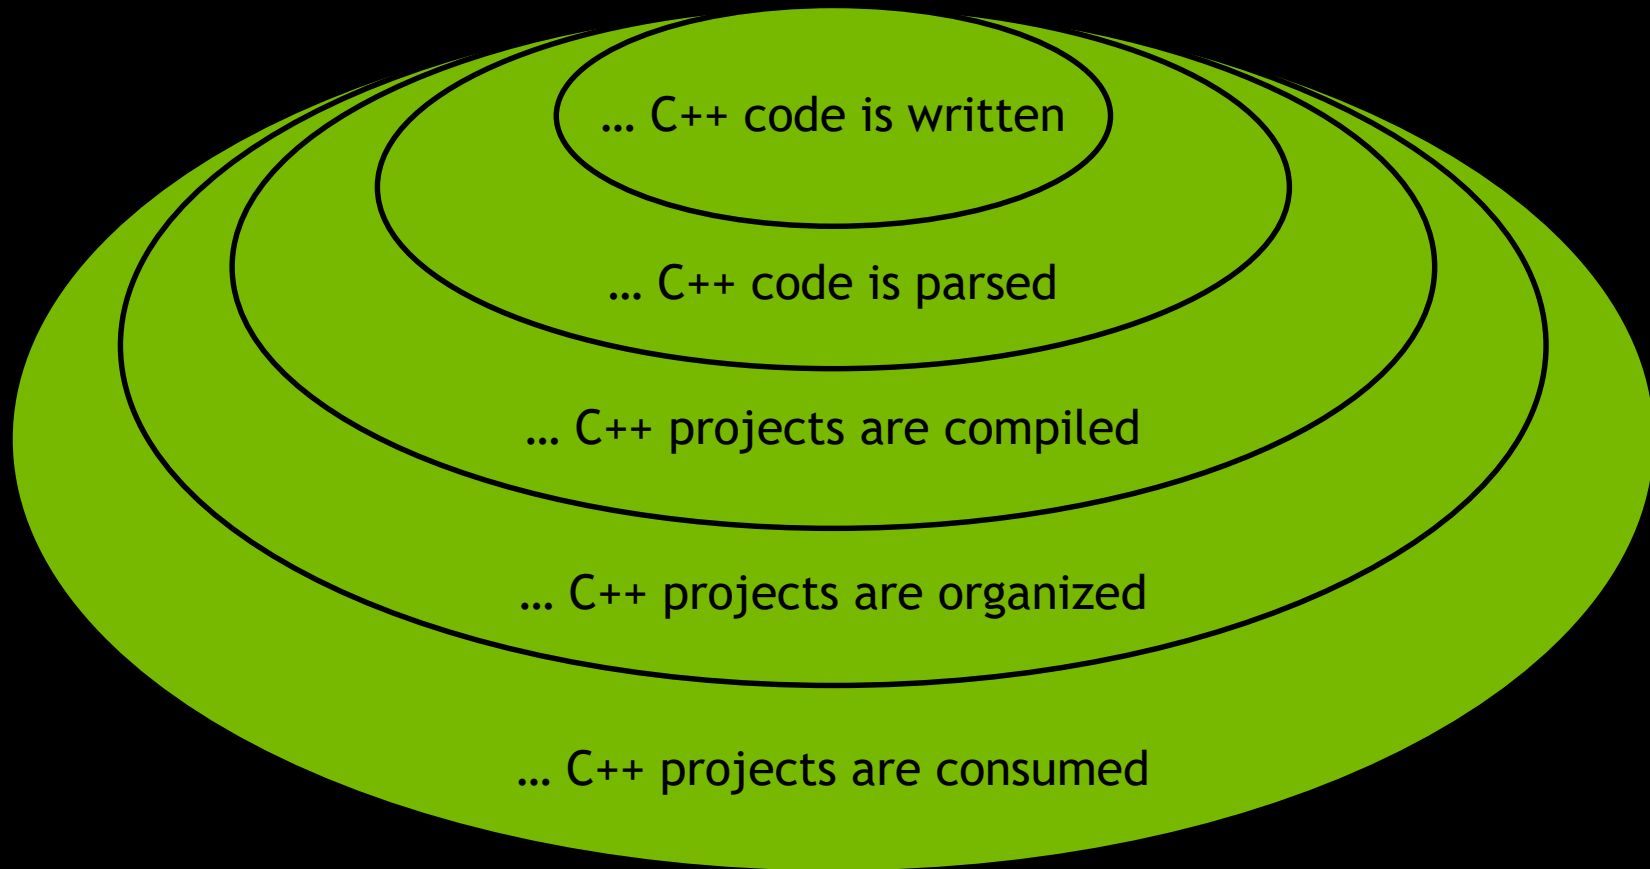
C++20's modules change how...



C++20's modules change how...



C++20's modules change how...



Modules will have a greater impact than any other feature added post C++-98.



Today's Compilation Model

What is C++'s compilation model today?

How do we organize C++ projects today?

“The text of the **program** is kept in units called **source files** in this International Standard.”

[lex.separate] p1 s1

“A **source file** together with all the **headers** and **source files** included via the preprocessing directive `#include`, less any source lines skipped by any of the conditional inclusion preprocessing directives, is called a **translation unit**.”

[\[lex.separate\] p1 s2](#)

Kinds of Translation Units

	Example	Extension	Artifact	Notes
Non-Modular Unit	<code>#include "..."</code> ...	<code>.cpp</code>	<code>.o</code>	

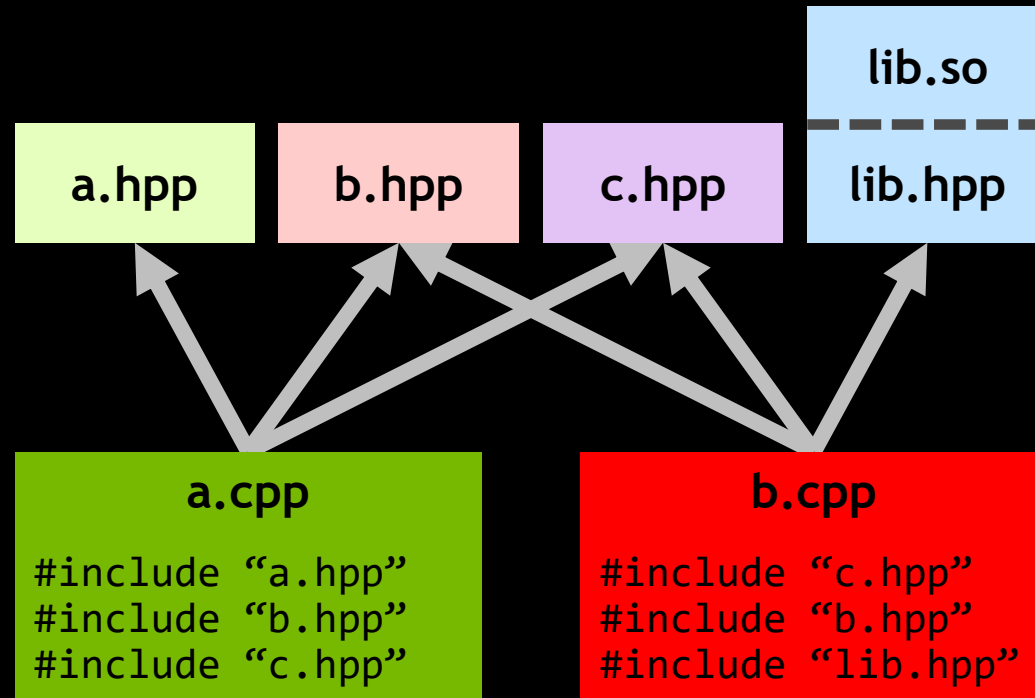
“A program consists of one or more *translation units* linked together.”

[basic.link] p1 s1

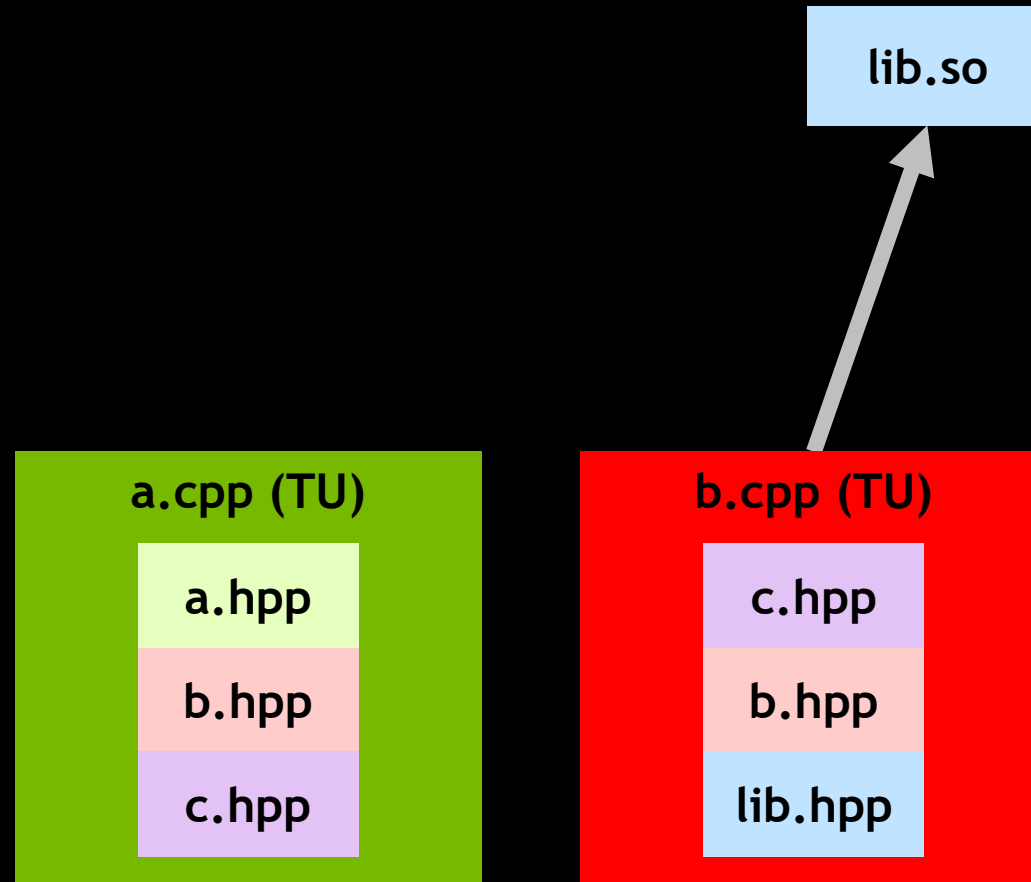
“Previously translated *translation units* and instantiation units can be preserved individually or in libraries.”

[lex.separate] p2 s1

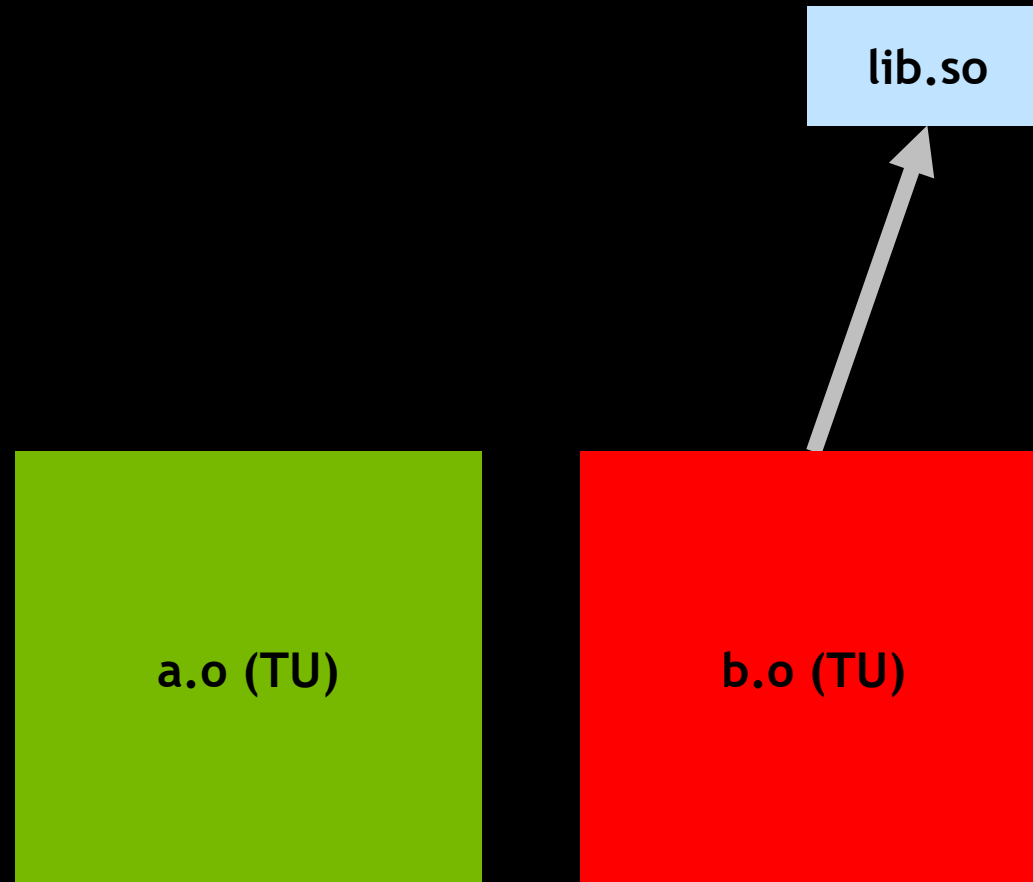
Textual Inclusion: Preprocess



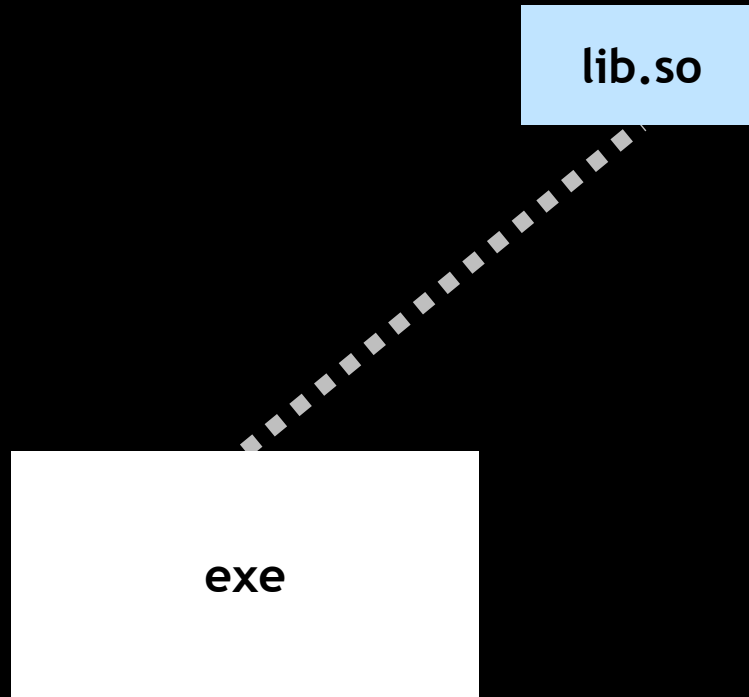
Textual Inclusion: Compile



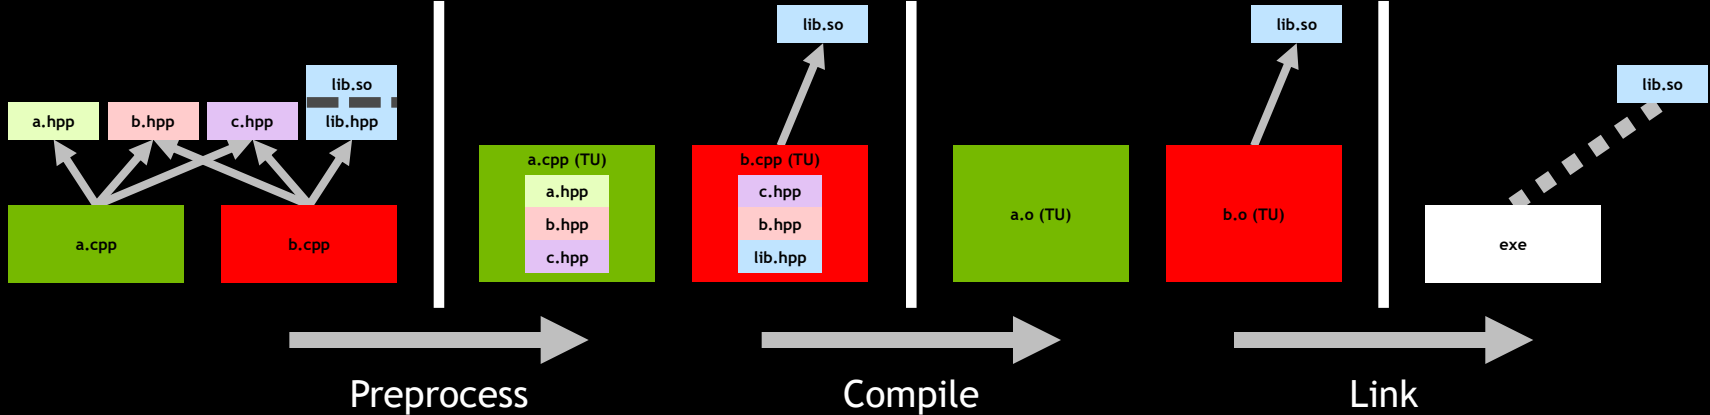
Textual Inclusion: Link



Textual Inclusion



Textual Inclusion




#include is problematic:

- Slow to compile.
- ODR violations.
- Lack of encapsulation.
- Cyclic dependencies.
- Order dependent.


#include is problematic:

- **Slow to compile.**
- ODR violations.
- Lack of encapsulation.
- Cyclic dependencies.
- Order dependent.

Textual Inclusion: Compile



Textual Inclusion



Pro: Embarrassingly parallel.

Con: x.hpp, y.hpp, and z.hpp are compiled 7 times.

#include is problematic:

- Slow to compile.
- **ODR violations.**
- Lack of encapsulation.
- Cyclic dependencies.
- Order dependent.

“There can be more than one definition of [an entity] in a program provided that no prior definition is necessarily reachable at the point where a definition appears, and provided the definitions satisfy the following requirements ...”

[\[basic.def.odr\] p12](#)

Ill formed, no diagnostic required
(IFNDR)

tree_node.hpp

```
#pragma once

template <typename T>
struct tree_node {
    T value;
    std::vector<tree_node*> children;
#ifdef DEBUG
    tree_node* parent;
#endif
};
```

a.cpp

```
#define DEBUG
#include "tree_node.hpp"

// ...
```

b.cpp

```
#include "tree_node.hpp"

// ...
```


#include is problematic:

- Slow to compile.
- ODR violations.
- Lack of encapsulation.
- Cyclic dependencies.
- Order dependent.

a.hpp

```
#pragma once

namespace c {
    struct A {
        private:
            template <typename T>
            void impl();
    };

    namespace detail::unsupported {
        template <typename T>
        void __please_dont_use();
    }
}
```

#include is problematic:

- Slow to compile.
- ODR violations.
- Lack of encapsulation.
- **Cyclic dependencies.**
- Order dependent.

a.hpp

```
#pragma once
#include "b.hpp"

struct Y;
struct X { Y* y; };
```

b.hpp

```
#pragma once
#include "a.hpp"

struct X;
struct Y { X* x; };
```

#include is problematic:

- Slow to compile.
- ODR violations.
- Lack of encapsulation.
- Cyclic dependencies.
- Order dependent.

a.hpp

```
#pragma once  
  
struct S { /* ... */ };
```

b.hpp

```
#pragma once  
  
void foo(S s);
```

c.cpp

```
#include "a.hpp"  
#include "b.hpp"
```

a.hpp

```
#pragma once  
  
struct S { /* ... */ };
```

b.hpp

```
#pragma once  
  
void foo(S s);
```

c.cpp

```
#include "a.hpp"  
#include "b.hpp"
```

d.cpp

```
#include "b.hpp"  
#include "a.hpp"
```

#include is problematic:

- Slow to compile.
- ODR violations.
- Lack of encapsulation.
- Cyclic dependencies.
- Order dependent.



Using Modules


Textual Inclusion

```
#include "foo.hpp"
```


Modular Import

```
import foo;
```

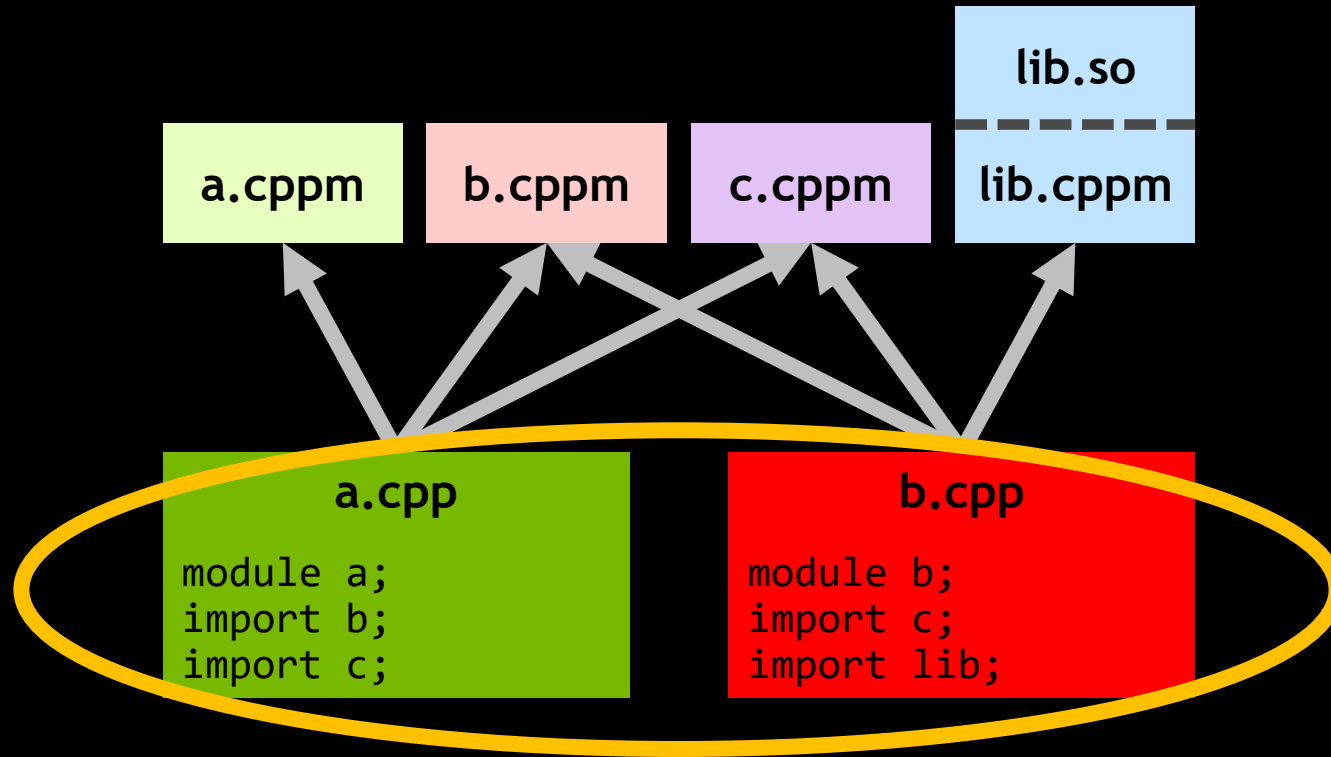
Modular Import



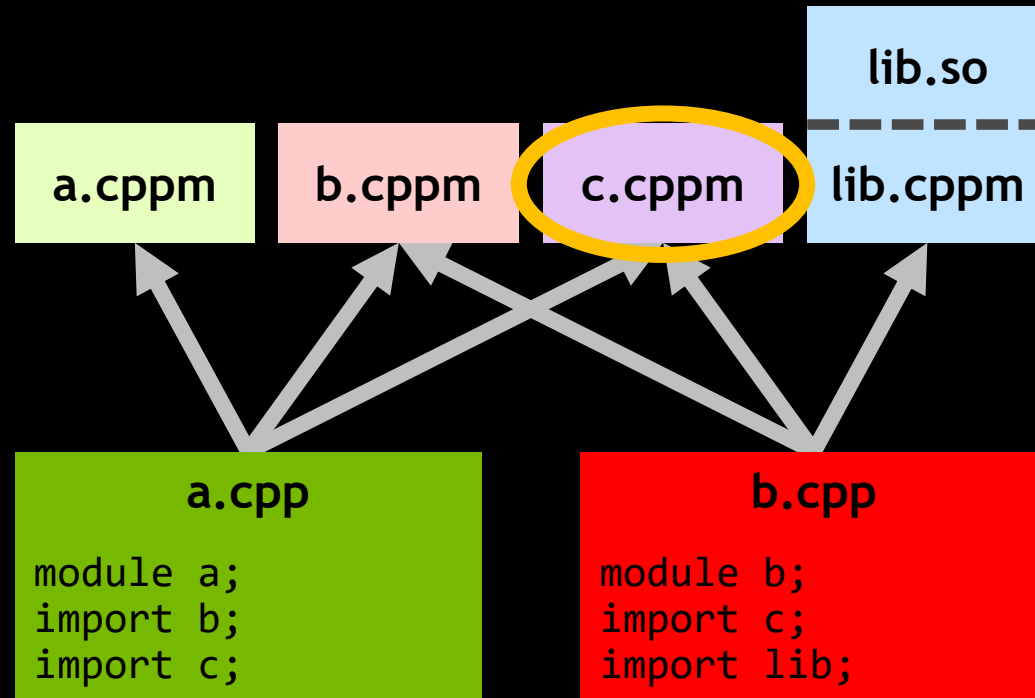
Modular Import




Modular Import




Modular Import




Modular Import: Precompile



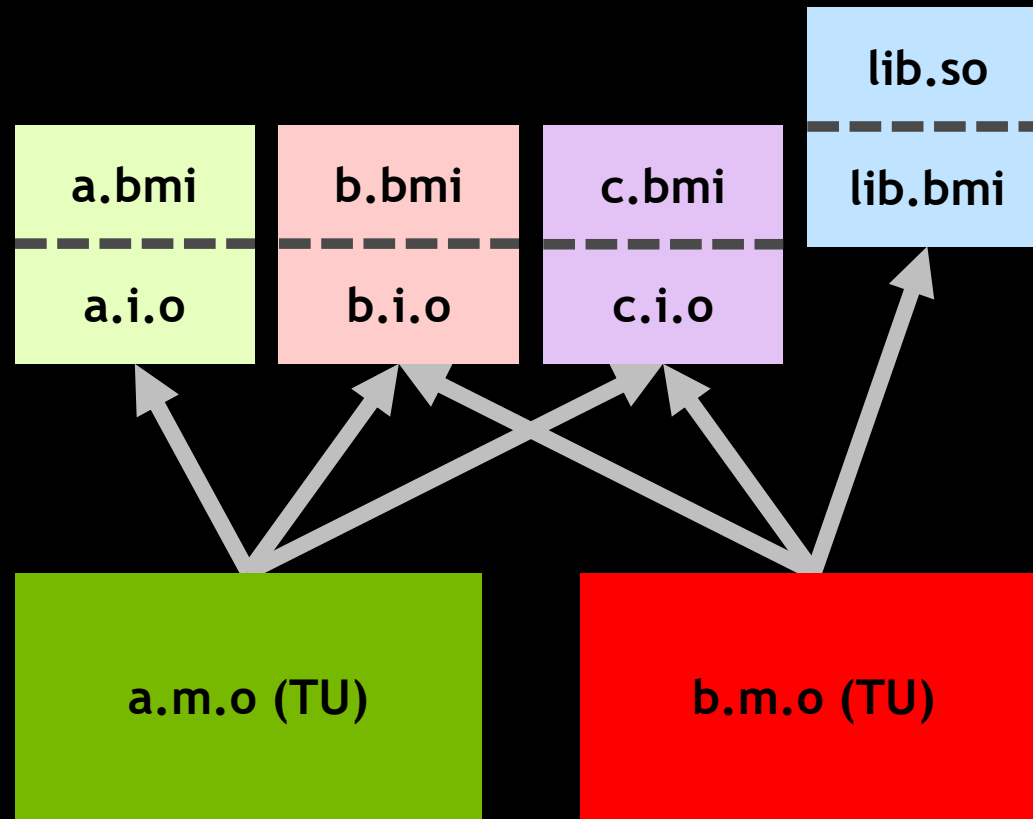
Modular Import: Preprocess



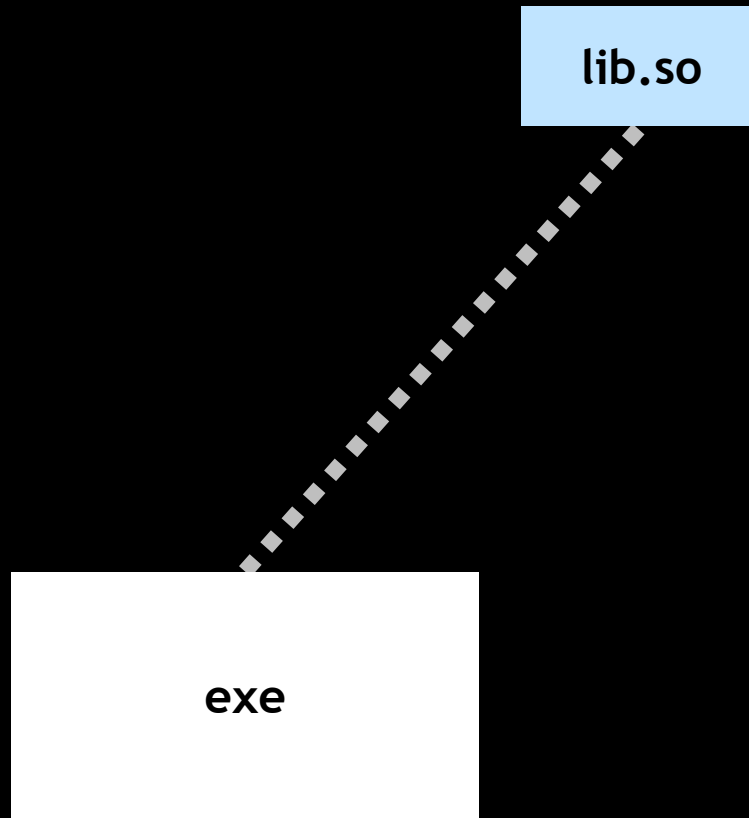
Modular Import: Compile



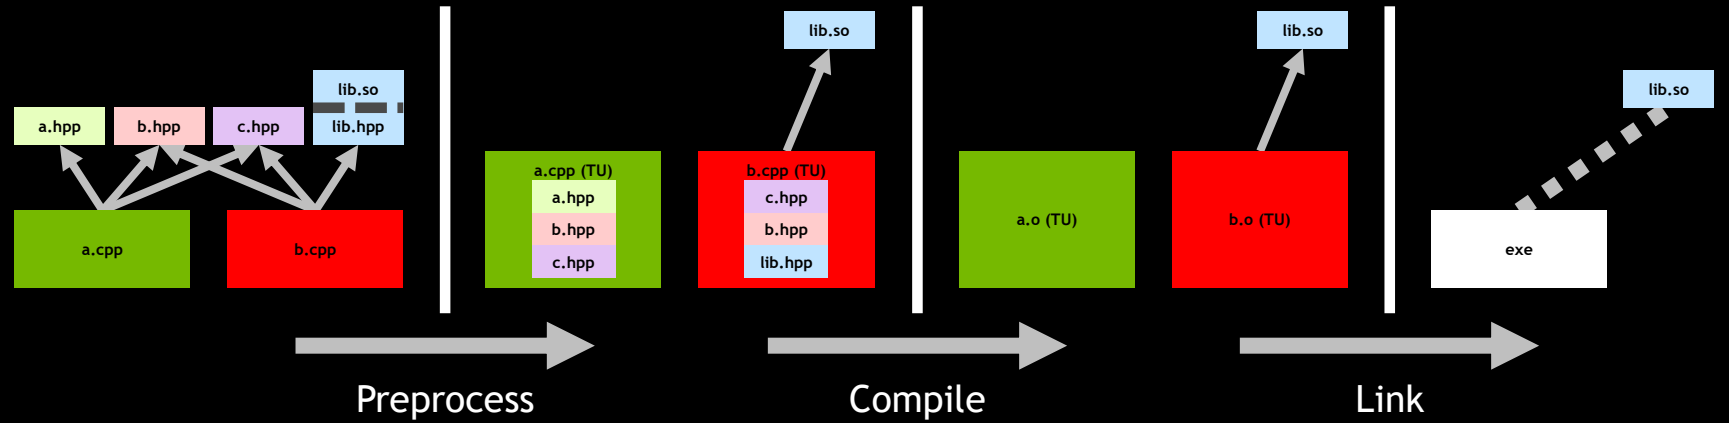
Modular Import: Link



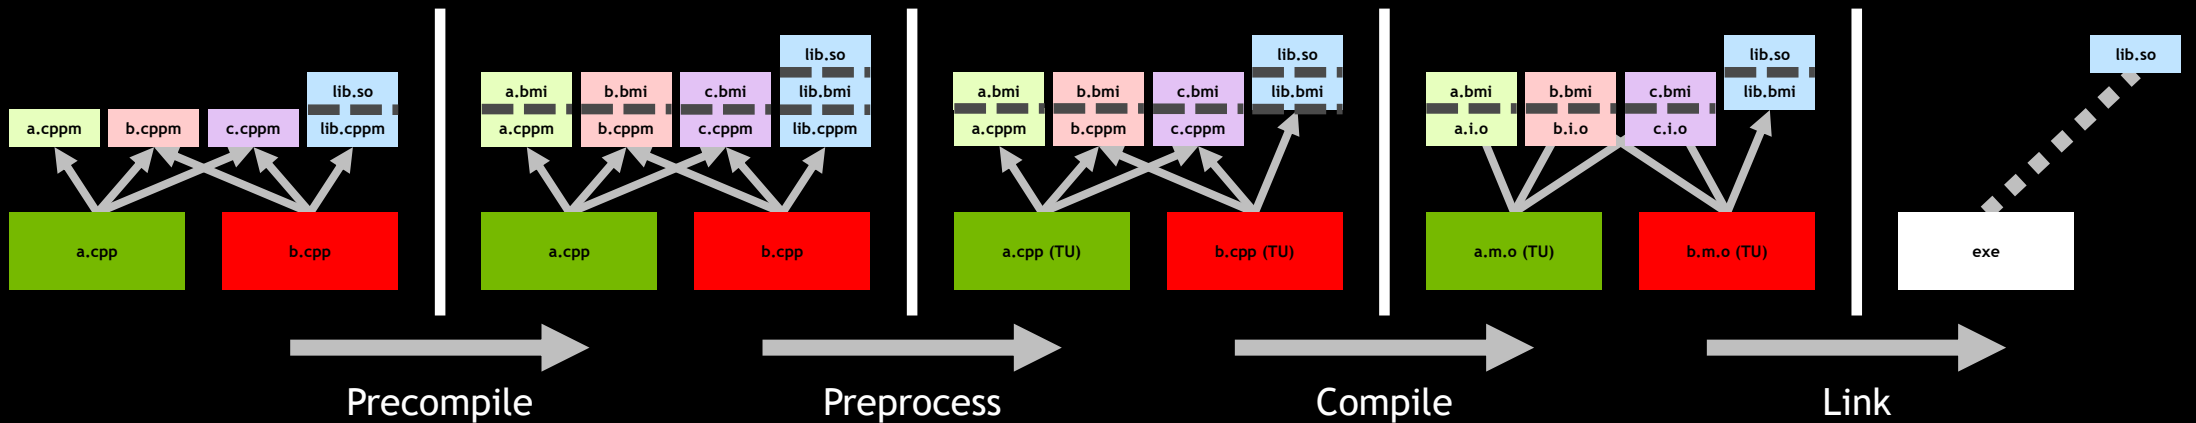
Modular Import




Textual Inclusion



Modular Import



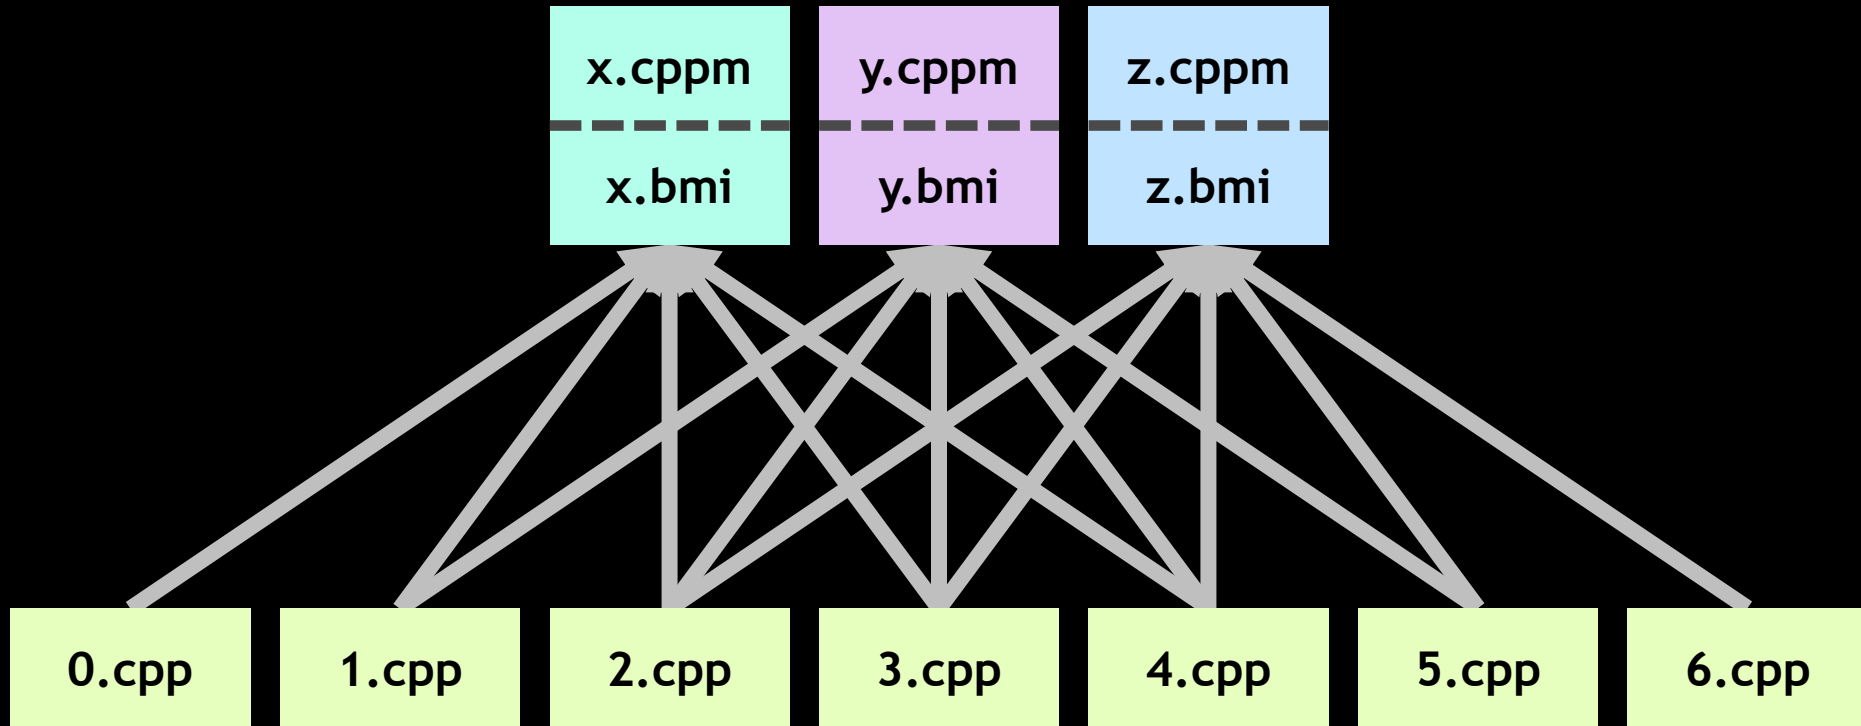
Textual Inclusion



Pro: Embarrassingly parallel.

Con: x.hpp, y.hpp, and z.hpp are compiled 7 times.

Modular Inclusion



Pro: x.cppm, y.cppm, and z.cppm are precompiled once.

Con: Not embarrassingly parallel.

Kinds of Translation Units

	Example	Extension	Artifact	Notes
Non-Modular Unit	<code>#include "..."</code> ...	<code>.cpp</code>	<code>.o</code>	
Module Interface Unit	<code>export module ...;</code> ...	<code>.cppm</code>	<code>.bmi</code> <code>.o</code>	Exactly one per module.
Module Implementation Unit	<code>module ...;</code> ...	<code>.cpp</code>	<code>.o</code>	

Textual Inclusion

```
#include <foo.hpp>  
#include "foo.hpp"
```

Modular Import

```
import foo;  
import <foo.hpp>;  
import "foo.hpp";
```


Importable headers:

- Most C++ standard library headers*.
- Some system headers.
- Headers you proclaim importable†.

* C standard library headers (<cfoo>, <foo.h>) are not required to be importable.

† Today, the mechanism for indicating importability is implementation defined (see [P1905](#)).

assert.h

```
#ifdef NDEBUG
#define assert(condition) ((void)0)
#else
#define assert(condition) /* ... */
#endif
```

cstddef

```
namespace std { /* ... */ }  
#define NULL /*see definition*/  
#define offsetof(P, D) /*see definition*/
```

```
import <foo>;
```

Ill-formed if foo is not importable.

Your implementation defines the set of importable headers.

Your Code

```
#include <vector>
#include <iostream>

// ...
```

Your Code

```
#include <vector>
#include <iostream>

// ...
```

Compiler Interpretation

```
import <vector>;
import <iostream>;

// ...
```

Include Translation: Implementations are free to interpret *#includes* of importable headers as imports.

Kinds of Translation Units

	Example	Extension	Artifact	Notes
Non-Modular Unit	<pre>#include "..." ...</pre>	.cpp	.o	
Header Unit	<pre>// Created by: import <...>;</pre>	.hpp	.bmi .o	
Module Interface Unit	<pre>export module ...; ...</pre>	.cppm	.bmi .o	Exactly one per module.
Module Implementation Unit	<pre>module ...; ...</pre>	.cpp	.o	

Textual Inclusion

Interfaces (headers) are not translation units; they are expanded into translation units.

Modular Import

Interfaces are translation units.

Modules are not translation units.



Writing Modules

```
import boost.spirit;  
import ctre;  
import blas.level1;
```

Module names are dot-separated identifiers.

```
import boost.spirit;  
import ctre;  
import blas.level1;
```

Module names are dot-separated identifiers.
Dots in module names have no semantic meaning.

```
module a;  
// ...  
module b;  
// ...
```

Only one module declaration per translation unit.

```
export module a;  
// ...  
export module b;  
// ...
```

Only one module declaration per translation unit.

Module Unit Structure

```
export module ...;  
import ...;  
...
```

```
export module a;  
// ...
```

Module Interface Unit

```
module a;  
// ...
```

Module Implementation Unit


```
export declaration
```

```
export {  
  declaration ...  
}
```

```
export void f();  
  
export int i{0};  
  
export struct A;  
  
export using B = A;  
  
export using std::any;
```

```
export {  
    void f();  
    struct A;  
    int i{0};  
}
```

```
export template <typename T>  
T square(T t) { return t * t; }
```

```
export template <typename T>  
struct is_const : false_type {};
```

```
export template <typename T>  
struct is_const<T const> : true_type {};
```

```
export namespace foo { struct A; }
```

```
namespace foo { struct B; }
```

```
export namespace foo { struct A; }  
namespace foo { struct B; }
```

Only `foo::A` is exported.

```
export import a;
```


square.cppm

```
export module math.square;  
  
export template <typename T>  
T square(T a) { return a * a; }
```

add.cppm

```
export module math.add;  
  
export template <typename T>  
T add(T a, T b) { return a + b; }
```

square.cppm

```
export module math.square;  
  
export template <typename T>  
T square(T a) { return a * a; }
```

add.cppm

```
export module math.add;  
  
export template <typename T>  
T add(T a, T b) { return a + b; }
```

math.cppm

```
export module math;  
  
export import math.square;  
export import math.add;
```

a.cppm

```
export module a;  
  
struct S { int m; };  
export S foo();
```

main.cpp

```
import a;  
  
int main() {  
    auto s0 = foo();  
    s0.m = 42;  
}
```

a.cppm

```
export module a;  
  
struct S { int m; };  
export S foo();
```

main.cpp

```
import a;  
  
int main() {  
    auto s0 = foo();  
    s0.m = 42;  
  
    S s1{};  
}
```

a.cppm

```
export module a;  
  
struct S { int m; };  
export S foo();
```

main.cpp

```
import a;  
  
int main() {  
    auto s0 = foo();  
    s0.m = 42;  
  
    decltype(foo()) s1{};  
}
```

Visible: In scope, can be named.

Reachable: In scope, not necessarily namable.

a.cppm

```
export module a;  
  
struct S { int m; };  
export S foo();
```

main.cpp

```
import a;
```

In main.cpp:

- S is reachable.
- foo is reachable and visible.

Modules enable true encapsulation.

Textual Inclusion

a.hpp

```
#pragma once  
  
struct foo { /* ... */ };
```

b.hpp

```
#pragma once  
  
void bar(foo f);
```

main.cpp

```
#include "a.hpp"  
#include "b.hpp"
```

Textual Inclusion

a.hpp

```
#pragma once  
  
struct foo { /* ... */ };
```

b.hpp

```
#pragma once  
  
void bar(foo f);
```

main.cpp

```
#include "a.hpp"  
#include "b.hpp"
```

Modular Import

a.cppm

```
export module a;  
  
export struct foo { /* ... */ };
```

b.cppm

```
export module b;  
  
export void bar(foo f);
```

main.cpp

```
import a;  
import b;
```

Textual Inclusion

a.hpp

```
#pragma once  
  
struct foo { /* ... */ };
```

b.hpp

```
#pragma once  
  
void bar(foo f);
```

main.cpp

```
#include "a.hpp"  
#include "b.hpp"
```

Modular Import

a.cppm

```
export module a;  
  
export struct foo { /* ... */ };
```

b.cppm

```
export module b;  
  
export void bar(foo f);
```

main.cpp

```
import a;  
import b;
```

Textual Inclusion

a.hpp

```
#pragma once

// Implementation detail/not part of the API.
template <typename T>
T id(T t) { return t; }
```

b.hpp

```
#pragma once

// Implementation detail/not part of the API.
template <typename T>
struct id { using type = T; };
```

main.cpp

```
#include "a.hpp"
#include "b.hpp"
```

Textual Inclusion

a.hpp

```
#pragma once

// Implementation detail/not part of the API.
template <typename T>
T id(T t) { return t; }
```

b.hpp

```
#pragma once

// Implementation detail/not part of the API.
template <typename T>
struct id { using type = T; };
```

main.cpp

```
#include "a.hpp"
#include "b.hpp"
```

Textual Inclusion

a.hpp

```
#pragma once

namespace a {
    template <typename T>
    T id(T t) { return t; }
}
```

b.hpp

```
#pragma once

namespace b {
    template <typename T>
    struct id { using type = T; };
}
```

main.cpp

```
#include "a.hpp"
#include "b.hpp"
```

Textual Inclusion

a.hpp

```
#pragma once

namespace a::detail {
    template <typename T>
    T id(T t) { return t; }
}
```

b.hpp

```
#pragma once

namespace b::detail {
    template <typename T>
    struct id { using type = T; };
}
```

main.cpp

```
#include "a.hpp"
#include "b.hpp"
```

Textual Inclusion

a.hpp

```
#pragma once

namespace a::detail {
    template <typename T>
    T __dont_use_id(T t) { return t; }
}
```

b.hpp

```
#pragma once

namespace b::detail {
    template <typename T>
    struct __dont_use_id { using type = T; };
}
```

main.cpp

```
#include "a.hpp"
#include "b.hpp"
```


Textual Inclusion

a.hpp

```
#pragma once

namespace a::detail {
    template <typename T>
    T __dont_use_id(T t) { return t; }
}
```

b.hpp

```
#pragma once

namespace b::detail {
    template <typename T>
    struct __dont_use_id { using type = T; };
}
```

main.cpp

```
#include "a.hpp"
#include "b.hpp"
```

Modular Import

a.cppm

```
export module a;

template <typename T>
T id(T t) { return t; }
```

b.cppm

```
export module b;

template <typename T>
struct id { using type = T; };
```

main.cpp

```
import a;
import b;
```

No more `detail/impl` namespaces.

No more uglifying identifiers.

Kinds of Linkage

	Example	Visible From	Notes
External Linkage	<pre>extern void foo(); export void bar(); extern int i{}; export bool b{};</pre>	Other translation units.	
Module Linkage	<pre>struct S; int foo(); int i{};</pre>	This module.	In non-modular units, entities with module linkage have external linkage.
Internal Linkage	<pre>static void foo(); static int i{}; bool const b{}; namespace { /* ... */ }</pre>	This translation unit.	
No Linkage	<pre>int main() { int i{}; }</pre>	This scope.	

Non-header-unit never expose macros.

```
import foo;           // No macros brought in.  
import <foo.hpp>;    // Macros brought in.  
import "foo.hpp";    // Macros brought in.
```

a.cppm

```
export module a;
```

```
#define foo int
```

main.cpp

```
import a;
```

```
foo f{};
```

a.cppm

```
export module a;
```

```
#define foo int
```

main.cpp

```
import a;
```

```
foo f{};
```

The definition of `foo` isn't brought in by the `import`.

a.hpp

```
#pragma once  
  
#define foo int
```

main.cpp

```
import <a>;  
  
foo f{};
```

a.hpp

```
#pragma once  
  
#define foo int
```

main.cpp

```
import <a>;  
  
foo f{};
```

The definition of `foo` is brought in by the `import`.

Modules never see the definition of programmatic macros defined in other modules.

Programmatic macros == `#define F00`

Non-programmatic macros == `-DF00`

a.cppm

```
export module a;  
  
export struct foo{};
```

main.cpp

```
#define foo bar  
import a;  
#undef foo  
  
foo f{};
```

a.cppm

```
export module a;  
  
export struct foo{};
```

main.cpp

```
#define foo bar  
import a;  
#undef foo  
  
foo f{};
```

The definition of `foo` isn't seen by the imported module.

a.cppm

```
export module a;  
  
#if defined(DEBUG)  
    // ...  
#else  
    // ...  
#endif
```

main.cpp

```
#define DEBUG  
import a;
```

a.cppm

```
export module a;  
  
#if defined(DEBUG)  
    // ...  
#else  
    // ...  
#endif
```

main.cpp

```
#define DEBUG  
import a;
```

The definition of DEBUG isn't seen by the imported module.

a.hpp

```
#pragma once

#if defined(DEBUG)
    // ...
#else
    // ...
#endif
```

main.cpp

```
#define DEBUG
import <a>;
```

The definition of DEBUG isn't seen by the imported header-unit.

Programmatic macros are:

- Never seen by modules.
- Never exported from non-header-unit modules.

With modules we no longer have to defend names:

- No more `detail/impl` namespaces.
- No more uglifying identifiers.

Bug ID: 2547402
Date: 3/6/2019 2:56:55 PM
Synopsis: CUB: NAMD defines `WARP_ANY` as a macro, which conflicts with CUB
Module: CUDA - C++ Core Libraries - Thrust

```
281  /**
282   * Warp any
283   */
284  __device__ __forceinline__ int WARP_ANY(int predicate, unsigned int member_mask)
285  {
286  #ifdef CUB_USE_COOPERATIVE_GROUPS
287      return __any_sync(member_mask, predicate);
288  #else
289      return ::__any(predicate);
290  #endif
291  }
```

```
#define _LIBCPP_NO_EXCEPTIONS  
import <vector>;
```

```
#define _LIBCPP_NO_EXCEPTIONS
import <vector>;
```

The definition of `_LIBCPP_NO_EXCEPTIONS` isn't seen by `<vector>`.

```
#define _LIBCPP_NO_EXCEPTIONS  
#include <vector>
```

```
#define _LIBCPP_NO_EXCEPTIONS
#include <vector>
```

If include translation occurs here,
the definition of `_LIBCPP_NO_EXCEPTIONS` isn't seen by `<vector>`.

```
#define NDEBUG
#include <cassert>

// For `readlink`.
#define _XOPEN_SOURCE
#include <unistd.h>
```

How do we deal with non-modular headers?

Textual Inclusion

a.hpp

```
#pragma once
#define NDEBUG
#include <cassert>

// ...
```

Modular Import

a.cppm

```
export module a;
#define NDEBUG
import <cassert>;

// ...
```

Textual Inclusion

a.hpp

```
#pragma once
#define NDEBUG
#include <cassert>

// ...
```

Modular Import

a.cppm

```
export module a;
#define NDEBUG
#include <cassert>

// ...
```


Module Unit Structure

```
module;  
#pp-directive ...;  
export module ...;  
import ...;  
...
```

Textual Inclusion

a.hpp

```
#pragma once
#define NDEBUG
#include <cassert>

// ...
```

Modular Import

a.cppm

```
module;
#define NDEBUG
#include <cassert>
export module a;

// ...
```

```
module;  
#include <boost/circular_buffer>  
export module boost.circular_buffer;  
namespace boost {  
    export using ::boost::circular_buffer;  
}
```

Modules are order independent.

```
import a;  
import b;
```

==

```
import b;  
import a;
```

Modules cannot have cycles.

Textual Inclusion

a.hpp

```
#pragma once
#include "b.hpp"

struct Y;
struct X { Y* y; };
```

b.hpp

```
#pragma once
#include "a.hpp"

struct X;
struct Y { X* x; };
```

Textual Inclusion

a.hpp

```
#pragma once
#include "b.hpp"

struct Y;
struct X { Y* y; };
```

b.hpp

```
#pragma once
#include "a.hpp"

struct X;
struct Y { X* x; };
```

Modular Import

a.cppm

```
export module a;
import b;

struct Y;
export struct X { Y* y; };
```

b.cppm

```
export module b;
import a;

struct X;
export struct Y { X* x; };
```


Textual Inclusion

a.hpp

```
#pragma once
#include "b.hpp"

struct Y;
struct X { Y* y; };
```

b.hpp

```
#pragma once
#include "a.hpp"

struct X;
struct Y { X* x; };
```

Modular Import

a.cppm

```
export module a;
import b;

struct Y;
export struct X { Y* y; };
```

b.cppm

```
export module b;
import a;

struct X;
export struct Y { X* x; };
```

Modular Import

a.cppm

```
export module a;  
  
struct Y;  
export struct X { Y* y; };
```

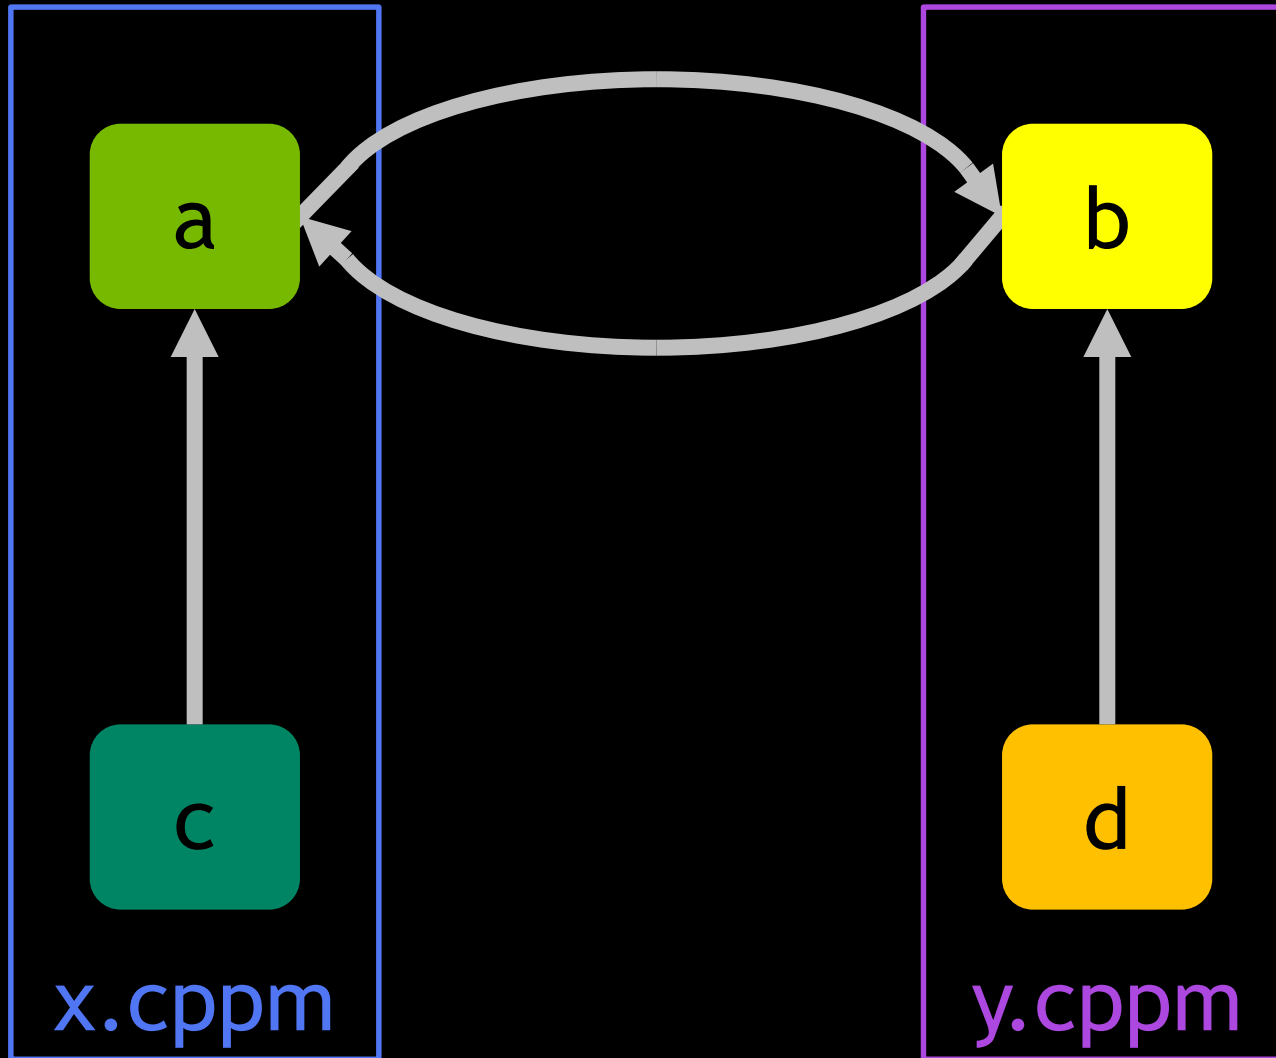
b.cppm

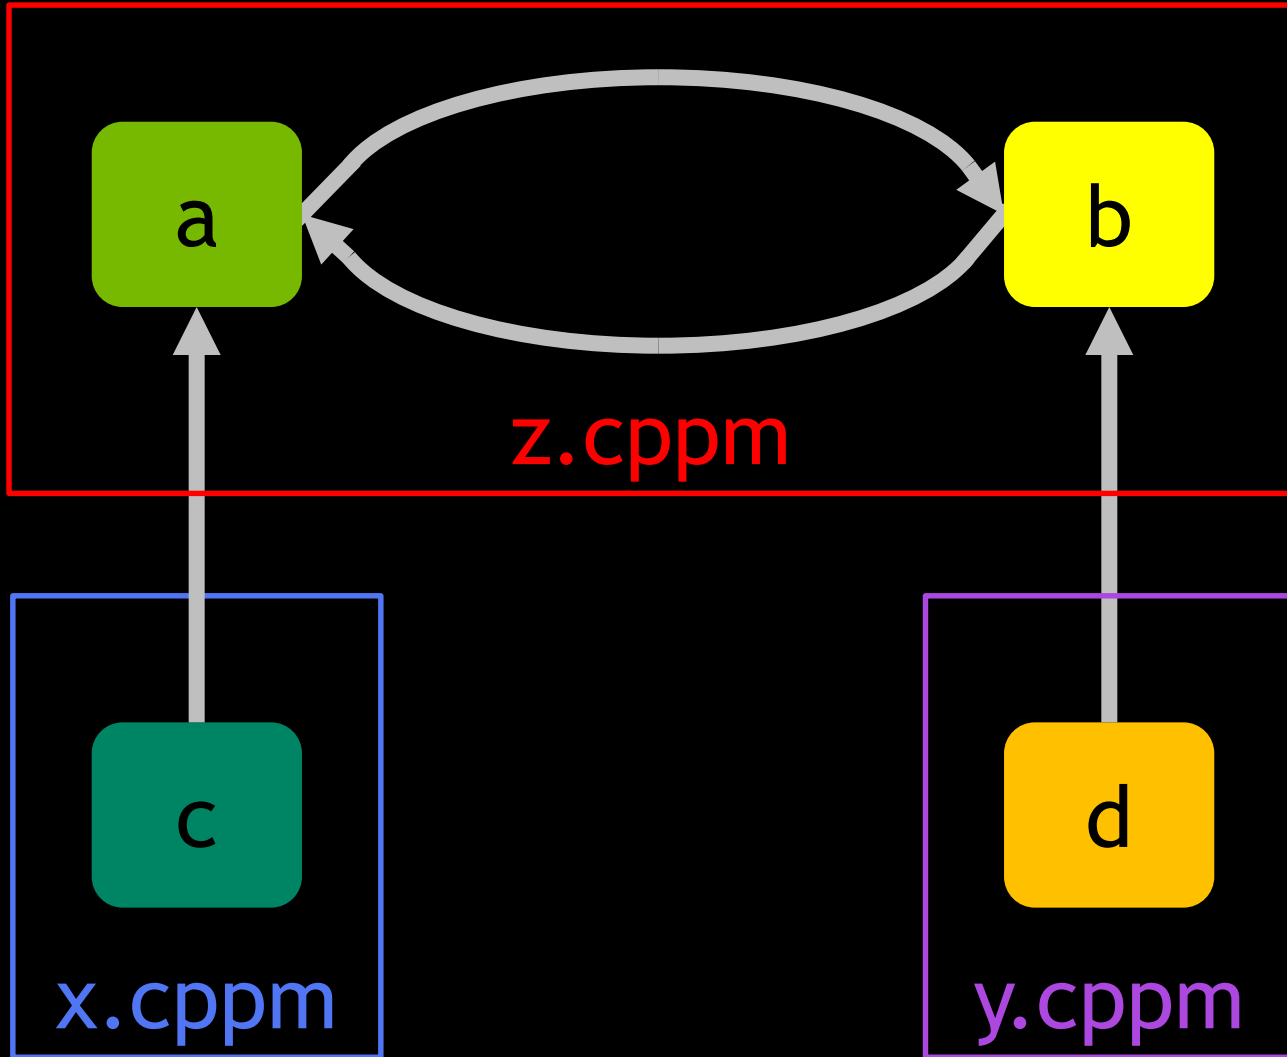
```
export module b;  
  
struct X;  
export struct Y { X* x; };
```

main.cpp

```
import a;  
import b;
```

How do we break cycles?





Breaking cyclic dependencies:

- Is difficult to automate.
- May require major structural changes.

Modules own their declarations.

Modular Import

a.cppm

```
export module a;  
  
struct Y;  
export struct X { Y* y; };
```

b.cppm

```
export module b;  
  
struct X;  
export struct Y { X* x; };
```

main.cpp

```
import a;  
import b;
```


You can not forward declare an entity
that lives in another module.

Textual Inclusion

a.hpp

```
#pragma once  
  
void foo();
```

a.cpp

```
#include "a.hpp"  
  
void foo() { /* ... */ }
```

b.cpp

```
#include "a.hpp"  
  
static void foo() { /* ... */ }
```

Modular Import

a.cppm

```
export module a;  
  
export void foo();
```

a.cpp

```
module a;  
  
void foo() { /* ... */ }
```

b.cpp

```
import a;  
  
static void foo() { /* ... */ }
```

Textual Inclusion

a.hpp

```
#pragma once  
  
void foo();
```

a.cpp

```
#include "a.hpp"  
  
void foo() { /* ... */ }
```

b.cpp

```
#include "a.hpp"  
  
static void foo() { /* ... */ }
```

Modular Import

a.cppm

```
export module a;  
  
export void foo();
```

a.cpp

```
module a;  
  
void foo() { /* ... */ }
```

b.cpp

```
import a;  
  
static void foo() { /* ... */ }
```

a.cppm

```
export module a;  
  
export void foo();
```

b.cppm

```
export module b;  
  
export void foo();
```

a.cppm

```
export module a;  
  
export void foo();
```

b.cppm

```
export module b;  
  
export void foo();
```

Ill formed, no diagnostic required (IFNDR).

Modules can be contained in one file.

Textual Inclusion

math.hpp

```
#pragma once

template <typename T>
T square(T a) { return a * a; }
```

main.cpp

```
import square;

int main() { return square(42); }
```

Modular Import

math.cppm

```
export module math;

export template <typename T>
T square(T a) { return a * a; }
```

main.cpp

```
import math;

int main() { return square(42); }
```

Textual Inclusion

a.hpp

```
#pragma once
```

```
struct pimpl;
```

a.cpp

```
#include "a.hpp"
```

```
struct pimpl { /* ... */ };
```

Modular Import

a.cppm

```
export module a;
```

```
export struct pimpl;
```

a.cpp

```
module a;
```

```
struct pimpl { /* ... */ };
```


Textual Inclusion

a.hpp

```
#pragma once
```

```
struct pimpl;
```

a.cpp

```
#include "a.hpp"
```

```
struct pimpl { /* ... */ };
```

Modular Import

a.cppm

```
export module a;
```

```
export struct pimpl;
```

```
module : private;
```

```
struct pimpl { /* ... */ };
```

Module Unit Structure

```
module;  
#pp-directive ...;  
export module ...;  
import ...;  
...  
module : private;  
...
```

Modules can be split across multiple files.

square.cppm

```
export module math.square;  
  
export template <typename T>  
T square(T a) { return a * a; }
```

add.cppm

```
export module math.add;  
  
export template <typename T>  
T add(T a, T b) { return a + b; }
```

math.cppm

```
export module math;  
  
export import math.square;  
export import math.add;
```

square.cppm

```
export module math:square;  
  
export template <typename T>  
T square(T a) { return a * a; }
```

add.cppm

```
export module math:add;  
  
export template <typename T>  
T add(T a, T b) { return a + b; }
```

math.cppm

```
export module math;  
  
export import :add;  
export import :square;
```

math_vector.cppm

```
export module math.vector;  
export import :dot_product;  
float sqrt(float x);
```

math_vector.cppm

```
export module math.vector;  
export import :dot_product;  
float sqrt(float x);
```

math_vector_dot_product.cppm

```
export module math.vector:dot_product;  
import :sum_squares;  
import <span>;  
export float vector_norm(std::span<float> x) {  
    return sqrt(sum_squares(x));  
}
```

math_vector.cppm

```
export module math.vector;  
export import :dot_product;  
float sqrt(float x);
```

math_vector_dot_product.cppm

```
export module math.vector:dot_product;  
import :sum_squares;  
import <span>;  
export float vector_norm(std::span<float> x) {  
    return sqrt(sum_squares(x));  
}
```

math_vector_sum_squares.cpp

```
module math.vector:sum_squares;  
import <span>;  
import <algorithm>;  
export float sum_squares(std::span<float> x) {  
    return std::transform_reduce(x);  
}
```


math_vector.cppm

```
export module math.vector;  
export import :dot_product;  
float sqrt(float x);
```

math_vector_dot_product.cppm

```
export module math.vector:dot_product;  
import :sum_squares;  
import <span>;  
export float vector_norm(std::span<float> x) {  
    return sqrt(sum_squares(x));  
}
```

math_vector_sum_squares.cpp

```
module math.vector:sum_squares;  
import <span>;  
import <algorithm>;  
float sum_squares(std::span<float> x) {  
    return std::transform_reduce(x);  
}
```

math_vector.cpp

```
module math.vector;  
float sqrt(float x) { /* ... */ }
```

Kinds of Translation Units

	Example	Extension	Artifact	Notes
Non-Modular Unit	<pre>#include "..." ...</pre>	.cpp	.o	
Header Unit	<pre>// Created by: import <...>;</pre>	.hpp	.bmi .o	
Module Interface Unit	<pre>export module ...; ...</pre>	.cppm	.bmi .o	Exactly one per module.
Module Implementation Unit	<pre>module ...; ...</pre>	.cpp	.o	
Module Partition Interface Unit	<pre>export module ...:...; ...</pre>	.cppm	.bmi .o	
Module Partition Implementation Unit	<pre>module ...:...; ...</pre>	.cpp	.o	

Modules do not force a file layout on you.

Modules are not translation units.

#include is problematic:

- Slow to compile.
- ODR violations.
- Lack of encapsulation.
- Cyclic dependencies.
- Order dependent.

Modules are coming:

- Modules will bring substantial algorithmic build throughput improvements.
- Modules offer true encapsulation which will eliminate many structural challenges with writing C++ at scale.
- Transitioning to modules will not be free.



Ecosystem Impact

math.hpp

```
#pragma once  
  
int square(int a);
```

math.cpp

```
#include "math.hpp"  
  
int square(int a) { return a * a; }
```

How are headers found?

- Not specified by the standard.
- In practice, all implementations assume a mapping between file names and #includes.
- The file is searched for in a set of include paths.

COMPILER EXPLORER

Add... More

Share Other Policies

C++ source #1 x x86-64 gcc 9.1 (Editor #1, Compiler #1) C++ x

Save/Load Add new... CppInsights C++ x86-64 gcc 9.1 -std=c++2a -Os

```
1 #include <https://compile-time.re/dfa.hpp>
2
3 bool match(std::string_view subject) {
4     return ctre::fast_match<"aloha|[a-z]*">(subject);
5 }
```

A

11010 .LX0: lib.f: .text // \s+ Intel Demangle

Libraries + Add new... Add tool...

```
1 match(std::basic_string_view<
2     add    rdi, rsi
3 .L3:
4     cmp    rdi, rsi
5     je     .L4
6     movsx  eax, BYTE PTR
7     sub    eax, 97
8     cmp    eax, 25
9     ja     .L5
10    inc    rsi
11    jmp    .L3
12 .L4:
13    mov    al, 1
14    ret
15 .L5:
16    xor    eax, eax
17    ret
```

Output (0/0) x86-64 gcc 9.1 - cached (1569389B)

How are modules found?

- Not specified by the standard.
- Unlike headers, modules are programmatically named.
- A file name <-> module name mapping is not straightforward.
 - Modules have to be precompiled.
 - Partitions span multiple files.

```
bar.cppm
```

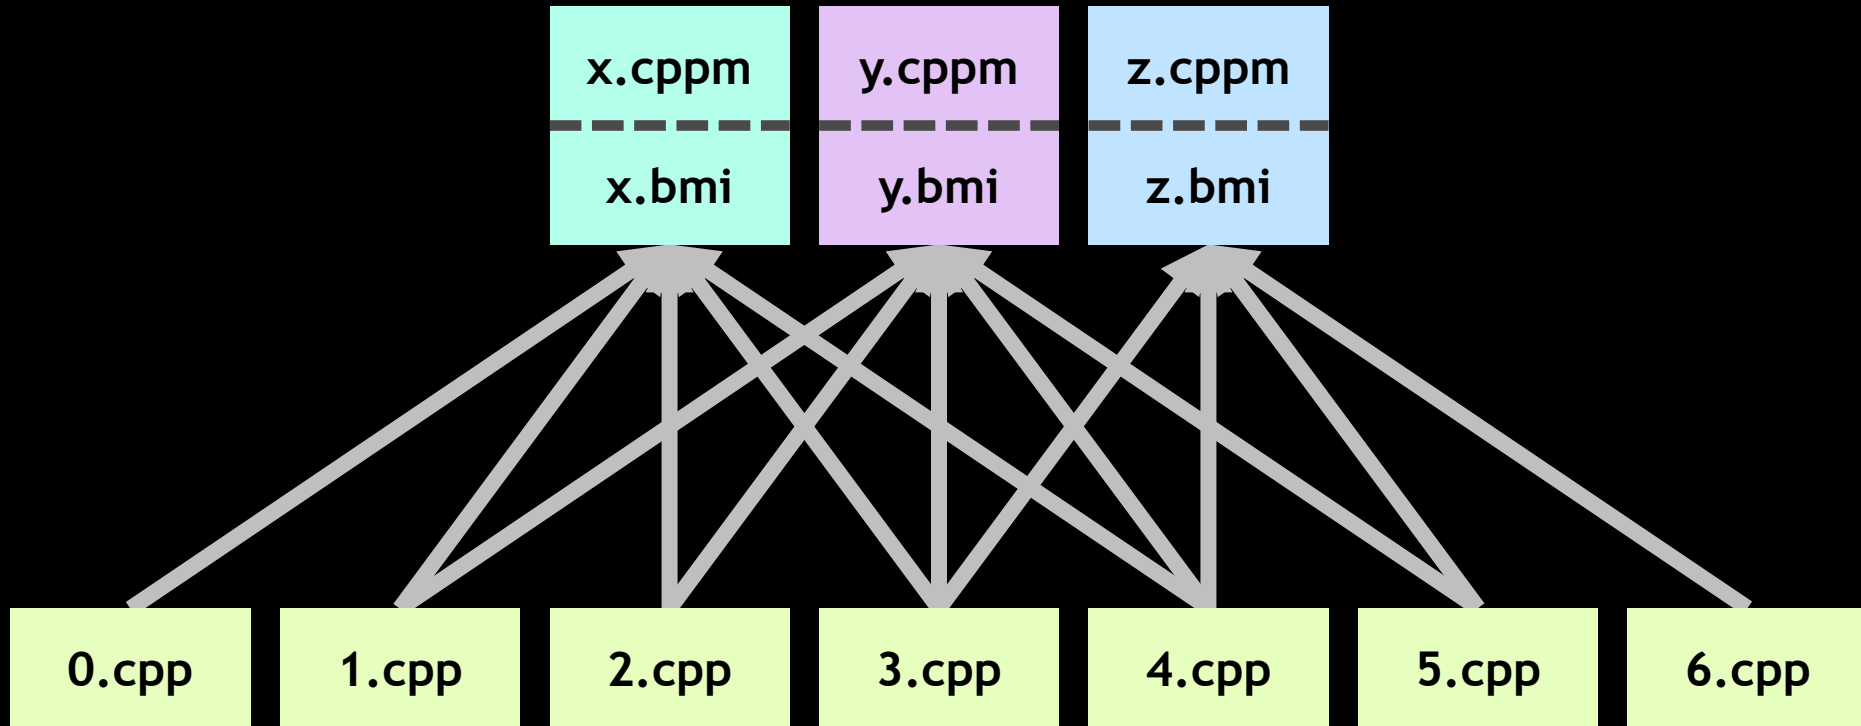
```
export module foo;
```

```
// ...
```

Strategies for Module Lookup

- (Fast) Dependency Scanner.
- Assume File Name == Module Name and Search.
- Explicit Mapping.
- Client/Server Mapper.

Implicit Precompilation is Problematic



If precompilation is implicit and builds are parallel, how do we decide who builds the BMs?

Implicit Precompilation is Problematic

Assuming the existence of fast dependency scanners, you can implicitly discover what BMIs need to be built.

This comes with a cost, however; you'll need to do this scanning as an additional pass prior to compilation.

Module Recipes

- BMIs are built with a certain set of compiler options and global macro definitions (the module recipe).
 - Ex: GCC 7.3, -O3, -DDEBUG
- The BMI may only be used when compiling with the same set of compiler options and global macro definitions.
- Tools need to either interface with the compiler and read BMIs, or be able to build BMIs of their own using the module recipe.
- Module lookup isn't just a matter of mapping a module name to a module interface unit (MIU) + a BMI. It's mapping a module name + module recipe to a MIU + a BMI.

Guidelines for BMI Redistribution

- BMIs are built with a certain set of compiler options and global macro definitions (the module recipe).
- The source for module interface units (MIUs) must always be distributed.
- BMIs may be distributed as an optimization alongside MIU sources.
- BMIs must be distributed with module recipes to allow tools to rebuild BMIs if needed.

Tools can no longer rely on simple lookup mechanism (include directories and header file names) to understand C++ projects.

Tools must be able to either consume BMIs or rebuild BMIs using module recipes.

Dependency scanning now requires a C++ parser, not just a C preprocessor.

C++ Ecosystem Technical Report

([P1688](#), [P1838](#))

C++ Ecosystem Technical Report:

- Packaging Modules ([P1767](#)).
- Module Mapping ([P1302](#), [P1484](#)).
- Module Naming ([P1634](#)).
- Module Transition Path.
- Module Recipes ([P1788](#)).
- Built Module Interface (BMI) Distribution ([P1788](#)).
- Dependency Scanning ([P1689](#), [P1857](#)).

C++ Ecosystem Technical Report Metadata Formats:

- Package Manifest Format ([P1767](#)).
- Dependency Format ([P1689](#)).
- Module Recipe Format ([P1788](#)).

Modules are coming:

- Modules will bring substantial algorithmic build throughput improvements.
- Modules offer true encapsulation which will eliminate many structural challenges with writing C++ at scale.
- Transitioning to modules will not be free.

Modules are a team effort.

Thanks to everyone involved!



@blelbach



nVIDIA®