

Semantic Tension: The Edge Cases of D's Introspective Power

DConf London '25

Timon Gehr



Compile-time Introspection in D

Enabled by (roughly):

- ▶ Code Generation
 - ▶ `template` (hygienic)
 - ▶ `mixin` (non-hygienic)
- ▶ Queries and Unification
 - ▶ `--traits`
 - ▶ `is` expressions
 - ▶ IFTI
- ▶ Conditional Compilation
 - ▶ `static if`
 - ▶ `static foreach`

--traits

TraitsExpression:

_traits (TraitsKeyword , TraitsArguments)

TraitsKeyword:

isAbstractClass
isArithmetic
isAssociativeArray
isFinalClass
isPOD
isNested
isFuture
isDeprecated
isFloating
isIntegral
isScalar
isStaticArray
isUnsigned
isDisabled
isVirtualFunction
isVirtualMethod
isAbstractFunction
isFinalFunction
isStaticFunction
isOverrideFunction
isTemplate
isRef
isOut
isLazy
isNone

__traits(compiles, ...)

26.5.29 *compiles*

1. Returns a bool **true** if all of the arguments compile (are semantically correct). The arguments can be symbols, types, or expressions that are syntactically correct. The arguments cannot be statements or declarations - instead these can be wrapped in a [function literal](#) expression.
2. If there are no arguments, the result is **false**.

```
static assert(!__traits(compiles));
static assert(__traits(compiles, 1 + 1)); // expression
static assert(__traits(compiles, typeid(1))); // type
static assert(__traits(compiles, object)); // symbol
static assert(__traits(compiles, 1, 2, 3, int, long));
static assert(!__traits(compiles, 3[1])); // semantic error
static assert(!__traits(compiles, 1, 2, 3, int, long, 3[1]));

enum n = 3;
// wrap a declaration/statement in a function literal
static assert(__traits(compiles, { int[n] arr; }));
static assert(!__traits(compiles, { foreach (e; n) {} }));

struct S
{
    static int s1;
    int s2;
}

static assert(__traits(compiles, S.s1 = 0));
static assert(!__traits(compiles, S.s2 = 0));
static assert(!__traits(compiles, S.s3));
```

```
__traits(allMembers, ...)
```

26.5.26 *allMembers*

1. Takes a single argument, which must evaluate to either a module, a struct, a union, a class, an interface, an enum, or a template instantiation.
2. A sequence of string literals is returned, each of which is the name of a member of that argument combined with all of the members of its base classes (if the argument is a class). No name is repeated. Built-in properties are not included.

```
import std.stdio;

class D
{
    this() {}
    ~this() {}
    void foo() {}
    int foo(int) { return 0; }
}

void main()
{
    auto b = [ __traits(allMembers, D) ];
    writeln(b);
    // ["__ctor", "__dtor", "foo", "toString", "toHash", "opCmp", "opEquals",
    // "Monitor", "factory"]
}
```

static if

25. Conditional Compilation

1. *Conditional compilation* is the process of selecting which code to compile and which code to not compile.

```
ConditionalDeclaration:  
  Condition DeclarationBlock  
  Condition DeclarationBlock else DeclarationBlock  
  Condition : DeclDefsopt  
  Condition DeclarationBlock else : DeclDefsopt
```

```
ConditionalStatement:  
  Condition NoScopeNonEmptyStatement  
  Condition NoScopeNonEmptyStatement else  
  NoScopeNonEmptyStatement
```

2. If the *Condition* is satisfied, then the following *DeclarationBlock* or *Statement* is compiled in. If it is not satisfied, the *DeclarationBlock* or *Statement* after the optional *else* is compiled in.
3. Any *DeclarationBlock* or *Statement* that is not compiled in still must be syntactically correct.
4. No new scope is introduced, even if the *DeclarationBlock* or *Statement* is enclosed by { }.
5. *ConditionalDeclarations* and *ConditionalStatements* can be nested.
6. The *StaticAssert* can be used to issue errors at compilation time for branches of the conditional compilation that are errors.

25.5 Static If Condition

```
StaticIfCondition:  
  static if ( AssignExpression )
```

1. *AssignExpression* is implicitly converted to a boolean type, and is evaluated at compile time. The condition is satisfied if it evaluates to **true**. It is not satisfied if it evaluates to **false**.

Contents [hide]

1. Version Condition
2. Version Specification
 1. Predefined Versions
3. Debug Condition
 1. Debug Statement
4. Debug Specification
5. Static If Condition
6. Static Foreach
 1. break and continue
7. Static Assert

Forward References

```
1 static if(__traits(compiles, z)) {
2     pragma(msg, "z exists");
3 }
4
5 pragma(msg, x); // 1
6
7 static if(__traits(compiles, x)) {
8     enum y = 2;
9 }
10 enum x = 1;
11
12 pragma(msg, x, "", y); // 1 2
```

```
$ dmd -o- test.d
1
1 2
```

Multiple Solutions

The following code is appended to all examples:

```
1 static if(__traits(compiles, x)) {
2     pragma(msg, "x exists");
3 }
4 static if(__traits(compiles, y)) {
5     pragma(msg, "y exists");
6 }
7 static if(__traits(compiles, z)) {
8     pragma(msg, "z exists");
9 }
```

Multiple Solutions

```
1 enum x = 1;
2
3 static if(__traits(compiles, x)) {
4     enum y = 2;
5 }
```

```
$ dmd -o- test.d
x exists
y exists
```

Multiple Solutions

```
1 static if(__traits(compiles, y)) {
2     enum x = 1;
3 }
4
5 enum y = 2;
```

```
$ dmd -o- test.d
x exists
y exists
```

Multiple Solutions

```
1 static if(__traits(compiles, y)) {
2     enum x = 1;
3 }
4 static if(__traits(compiles, x)) {
5     enum y = 2;
6 }
```

```
$ dmd -o- test.d
$
```

Multiple Solutions (Cont.)

```
1 enum x = 1;
2
3 static if(!__traits(compiles, x)) {
4     enum y = 2;
5 }
```

```
$ dmd -o- test.d
x exists
$
```

Multiple Solutions (Cont.)

```
1 static if(!__traits(compiles, y)) {
2     enum x = 1;
3 }
4
5 enum y = 2;
```

```
$ dmd -o- test.d
y exists
$
```

Multiple Solutions (Cont.)

```
1 static if(!__traits(compiles, y)) {
2     enum x = 1;
3 }
4 static if(!__traits(compiles, x)) {
5     enum y = 2;
6 }
```

```
$ dmd -o- test.d
x exists
$
```

Multiple Solutions (Cont.)

```
1 static if(__traits(compiles, z)) {
2     static if(!__traits(compiles, y)) {
3         enum x = 1;
4     }
5 }
6 static if(!__traits(compiles, x)) {
7     enum y = 2;
8 }
9
10 enum z = 3;
```

```
$ dmd -o- test.d
x exists
z exists
$
```

Multiple Solutions (Cont.)

```
1 static if(__traits(compiles, z)) {
2     static if(!__traits(compiles, y)) {
3         enum x = 1;
4     }
5 }
6 static if(!__traits(compiles, x)) {
7     enum y = 2;
8 }
9 static if(__traits(compiles, y)) {
10     enum z = 3;
11 }
```

```
$ dmd -o- test.d
y exists
z exists
$
```

Module order

```
1 module a;    import b;
2 static if(!__traits(compiles, y)) {
3     enum x = 1;
4 }
```

```
1 module b;    import a;
2 static if(!__traits(compiles, x)) {
3     enum y = 2;
4 }
```

```
$ dmd -o- a.d b.d
```

```
y exists
```

```
$
```

```
$ dmd -o- b.d a.d
```

```
x exists
```

```
$
```

Self-contradictory code

```
1 static if(__traits(compiles, x)) {
2     pragma(msg, "x defined above");
3 }
4 static if(!__traits(compiles, x)) {
5     enum x = 2;
6 }
```

```
$ dmd -o- test.d
x exists
$
```

The same condition in the same scope evaluates to different results.

Self-contradictory code (Cont.)

```
1 string foo() => q{static string foo() => q{int bar() => 2;};};
2
3 struct S{
4     mixin(foo());
5 }
6
7 pragma(msg, __traits(allMembers, S));
```

```
$ dmd -o- test.d
AliasSeq!("foo")
$
```

Self-contradictory code (Cont.)

```
1 string foo() => q{static string foo() => q{int bar() => 2;};};  
2  
3 struct S{  
4     mixin(foo());  
5     mixin(foo());  
6 }  
7  
8 pragma(msg, __traits(allMembers, S));
```

```
$ dmd -o- test.d  
AliasSeq!("foo","bar")  
$
```

Self-contradictory code (Cont.)

```
1 string foo() => q{static string foo() => q{int bar() => 2;};};  
2  
3 struct S{  
4     alias a = foo;  
5     mixin(foo());  
6     alias b = foo;  
7     static assert(__traits(isSame, a, b)); // fails  
8 }
```

Self-contradictory code (Cont.)

```
1 string foo() => q{static string foo() => q{int bar() => 2;};};  
2  
3 struct S{  
4     static if(__traits(compiles, b)){  
5         alias a = foo;  
6     }  
7     mixin(foo());  
8     alias b = foo;  
9     static assert(__traits(isSame, a, b)); // passes  
10 }
```

```
1 int foo(long z) => 1;
2 static if(foo(1) == 1) {
3     static assert(foo(1) == 1);
4
5     // ...
6 }
```

```
1 int foo(long z) => 1;
2 static if(foo(1) == 1) {
3     static assert(foo(1) == 1); // fails
4
5     int foo(int x) => 2;
6 }
```

Caching

```
1 int foo(long z) => 1;
2 enum Foo(alias x)=foo(x);
3
4 static assert(Foo!1 == foo(1)); // fails
5
6 alias a = Foo!1;
7 static if(foo(1) == 1) {
8     int foo(int x) => 2;
9 }
```

Invalid Solutions

```
1 static if(__traits(compiles, x)) {
2     enum y = true;
3 }
4 static if(!__traits(compiles, z)) {
5     enum x = true;
6 }
```

```
$ dmd -o- test.d
x exists
$
```

Invalid Solutions

```
1 static if(!__traits(compiles, z)) {
2     enum x = true;
3 }
4 static if(__traits(compiles, x)) {
5     enum y = true;
6 }
```

```
$ dmd -o- test.d
x exists
y exists
$
```

What else could D do?

- ▶ Use fixed evaluation order.
 - ▶ Breaks forward referencing, inconsistent.
 - ▶ Would need to define global order on modules. Brittle.

What else could D do?

- ▶ Actually solve as a constraint satisfaction problem.
 - ▶ Truly order-independent.
 - ▶ Infeasible in the general case.

N.B. Unique-1-in-3SAT

We can encode any unique-1-in-3SAT problem in conditional variable declarations, e.g.:

$$\varphi(\mathbf{x}) = (\neg x_1 \vee \neg x_5 \vee x_3) \wedge (\neg x_2 \vee \neg x_4 \vee x_1) \wedge \dots$$

```
1 static if(__traits(compiles, x1) && __traits(compiles, x5)) {
2     void x3(){} // using functions so we can generate overloads
3 }
4 static if(__traits(compiles, x2) && !__traits(compiles, x4)) {
5     void x1(){}
6 }
```

(Variables that only occur in negative clauses have to be manually pinned at true one-by-one to verify unsatisfiability.) This is NP-hard (under randomized reductions).

What else could D do?

- ▶ Explicit staging.
 - ▶ Can only reflect on properties determined at a lower stage.
 - ▶ Annotation overhead.
 - ▶ Suboptimal modularity.
 - ▶ May disallow useful code.

What else could D do?

- ▶ Track dependencies, analyze until saturation, then force decisions.
 - ▶ Source-order-independent.
 - ▶ Still dependent on nesting order.
 - ▶ Only produces consistent solutions.
 - ▶ Sometimes disallows useful code.
 - ▶ May pick one of multiple possible solutions.
- ▶ Hybrid versions (not much explored)
 - ▶ E.g., iterate to saturation, then use tie-breakers based on source order to schedule decisions.

Analyzing until Saturation

```
1 static if(__traits(compiles, x)) {
2     enum y = 2;
3 }
4 static if(!__traits(compiles, z)) {
5     enum x = 1;
6     mixin(q{enum k = 0;});
7 }
8 static if(__traits(compiles, w)) {
9     enum z = 3;
10 }
```

```
$ dmd -o- test.d
x exists
$
```

Analyzing until Saturation

```
1 static if(__traits(compiles, x)) { // A: depends on B
2     enum y = 2;
3 }
4 static if(!__traits(compiles, z)) { // B: depends on C
5     enum x = 1;
6     mixin(q{enum k = 0;}); // treated as opaque
7 }
8 static if(__traits(compiles, w)) { // depends on B
9     enum z = 3;
10 }
```

A → B <-> C

Analyzing until Saturation

```
1 static if(__traits(compiles, x)) {  
2     enum y = 2;  
3 }  
4 // z does not exist  
5 enum x = 1;  
6 enum k = 0;  
7  
8 // w does not exist  
9  
10  
11 //
```

Analyzing until Saturation

```
1
2 enum y = 2; // y exists!
3
4 // z does not exist
5 enum x = 1;
6 enum k = 0;
7
8 // w does not exist
9
10
11 //
```

Analyzing Contradictions

```
1 static if(!__traits(compiles, x)) {
2     enum x = 1;
3 }
```

Analyzing Contradictions

```
1 // x does not exist
2 enum x = 1;
```

Analyzing Contradictions

```
1 // x does not exist
2 enum x = 1; // error: x introduced anyway
```

Experimental D frontend

<https://github.com/tgehr/d-compiler>

- ▶ Explores analyzing to saturation iteratively.
- ▶ Started out as a lexer and parser for D.
- ▶ My first side project using D.
- ▶ 2011-2017. (Manual version control 2011-2013, git later.)
- ▶ Very incomplete, but does have templates with IFTI, and CTFE.
- ▶ Does some things better than DMD, does not do some other things at all.

How is it implemented?

- ▶ Compile-time “visitor pattern”.
- ▶ Single *reentrant* method for semantic analysis.
- ▶ Conceptually, every AST node is a coroutine.
 - ▶ In practice, most analysis is handled via direct recursion.
 - ▶ Explicit *scheduler* detects and resolves cyclic dependencies.
 - ▶ String mixins for handling suspension.

Compile-time “visitor pattern”

```
1 // from visitors.d:  
2 mixin template Visitors(){  
3     mixin Semantic!(typeof(this)); // semantic analysis  
4     mixin Analyze; // ad-hoc visitors, often templated  
5     mixin CTFEInterpret!(typeof(this)); // ctfc bytecode  
       interpreter  
6     mixin DeepDup!(typeof(this)); // create a deep duplicate of  
       an AST  
7 }
```

Compile-time “visitor pattern” (Cont.)

Then in every node class:

```
1 class Node{
2     // ...
3     mixin Visitors
4 }
5 class Expression{
6     // ...
7     mixin Visitors;
8 }
9 class Statement{
10    // ...
11    mixin Visitors;
12 }
13 // ...
```

Compile-time “visitor pattern” (Cont.)

In semantic.d:

```
1 template Semantic(T) if(is(T==Node)){
2     // visitors can add fields in addition to methods
3     Node rewrite; // for lazy substitution
4     SemState sstate = SemState.begin;
5     ubyte needRetry = false;
6
7     void semantic(Scope sc)in{assert(sstate>=SemState.begin);}body{
8         mixin(SemPrlg); // register with scheduler, return if completed, etc.
9         sc.error("feature not implemented",loc);
10        mixin(ErrEplg); // ‘‘terminate’’ node with error
11    }
12
13    // analysis is trapped because of circular await-relationship involving
14    // this node
15    void noHope(Scope sc){}
16
17 template Semantic(T) if(is(T==Expression)){ ... }
18 template Semantic(T) if(is(T==Statement)){ ... }
```

Tradeoffs

Pro:

- ▶ Good conceptual encapsulation.
 - ▶ Like visitor pattern, but nicer to use.
 - ▶ Can add multiple related fields and methods.
 - ▶ Everything related to one method is in one place.
 - ▶ Can add mixin templates in parent class as customization points.
- ▶ Only one virtual call instead of two.
- ▶ All information available at compile time for introspection.
- ▶ Can implement virtual methods for different templated classes at once.

Con:

- ▶ Terrible actual encapsulation.
 - ▶ Template mixins are unhygienic.
 - ▶ Therefore everything imports everything else.
 - ▶ Visibility modifier meanings don't correspond to code organization.

Coming clean

```
1 // actually from visitors.d:  
2 mixin template Visitors(){  
3     // workaround for DMD bug: Interpret goes first  
4     /*static if(is(typeof({mixin Semantic!(typeof(this));})))*/  
5     static if(is(typeof(this):Expression)&&!is(typeof(this):  
6         Type)) mixin Interpret!(typeof(this));// TODO: minimize  
7         and report bug  
8     static assert(is(TypeTuple==class));  
9     static if(!IsNonASTType!(typeof(this))) mixin Semantic!(  
10        typeof(this));  
11    // another workaround for DMD bug, other part is in  
12    expression.Node  
13    static if(!is(typeof(this)==Node)){  
14        static if(!is(typeof(this)==AggregateTy)){  
15            mixin Analyze; // wtf?  
16            mixin CTFEInterpret!(typeof(this));  
17        }  
18        static if(!is(typeof(this)==AggregateTy)) mixin DeepDup  
19        !(typeof(this));
```

Live code browsing

- ▶ Lexer
- ▶ Parser
- ▶ Semantic analysis
- ▶ Scheduler
- ▶ Bytecode Interpreter

Content warning

- ▶ I will proceed to show code with `is(typeof(...))`.
- ▶ Viewer discretion is advised.
- ▶ Don't use `is(typeof(...))`.
- ▶ `__traits(compiles, ...)` is strictly superior.
- ▶ Except my frontend does not implement `__traits...`
- ▶ `typeof` eagerly determines types for expressions that would not compile!

Demo

- ▶ Detection of contradictory/ambiguous code
- ▶ Template demo
- ▶ CTFE demo
- ▶ Improved IFTI
 - ▶ Matching template aliases
 - ▶ Improved unification features

Not explored: Attribute inference

I had also wanted the frontend to do better at attribute inference, but never got around to it.

- ▶ The frontend does not yet infer attributes.
- ▶ Attribute inference in DMD does not work for recursive functions.
- ▶ Interesting questions related to attribute inference and introspection.
- ▶ Can do something based on fixed points as there is a natural lattice structure.

`static foreach`

- ▶ `static foreach` first prototyped in my frontend
- ▶ Similar implementation worked in DMD.

Experimental Frontend Death

- ▶ With DMD 2.061, I was finally no longer able to work around the regressions:
semantic.d(7834): Error: function
type.BasicType.Visitors!().Semantic!(BasicType).implicitlyConvertsTo does
not override any function, did you mean
'expression.Expression.Visitors!().Semantic!(Expression).implicitlyConvertTo'
- ▶ Rainer Schütze made a patch that got the frontend running under 2.074,
but it was never pulled as it broke other tests.
- ▶ Some reduced issues were fixed but they invariably failed in other ways.

Thanks!

Questions?