

Doc No: SC22/WG21/N2103
J16/06-0173
Date: 2006-10-31
Project: JTC1.22.32
Reply to: Herb Sutter
hsutter@microsoft.com

A Modest Proposal: Fixing ADL (revision 2)

ADL is widely known to be an essential feature. On the other hand, it is also unarguably a source of many subtle problems.

This paper first reviews the motivation behind ADL. It then summarizes characteristic examples of the real problems that arise in practice with ADL, shows their common root cause in that ADL is overeager, and proposes a minor refinement to ADL that I believe solves all of the reported problem cases, and does so with little or no change to the standard library (even though it's still all in one namespace).

The proposed change can cause some existing code to fail to compile:

- Usually all that is necessary is to add a qualification (or `using`) at the call site.
- To preserve the ability to use some library code organization styles, this revision adds a facility to explicitly name the current associated namespaces of a type or object in special cases (e.g., a closed source base) when lookup using the C++98 ADL rules is desired.

The differences between this proposal and revision 1 are:

- *Clarified the question:* Discussion in Portland made it clear that the basic question is whether ADL should find unconstrained templates or not.
- *Removed the proposal to restrict the set of associated namespaces.* The previous version recommended two changes: (a) restricting the entities that are looked up; and (b) optionally also restricting the set of associated namespaces. In my view, only (a) is necessary and sufficient to “fix ADL.” This draft recommends only (a), and drops (b) because it is not necessary and may be controversial.
- *Added `using func;`* As directed by EWG in Mont Tremblant and Berlin, added this backward compatibility feature to explicitly get the C++98 ADL behavior.

1. Background

1.1. Why ADL

ADL is a result of having a language that supports both namespaces and nonmember functions, especially when some functions that clearly are part of a type's public interface (e.g., `operator<<`) are required to be nonmembers.

Specifically, ADL exists to treat nonmember and member name lookup more uniformly. The principle behind ADL is that if you call member functions they naturally “come along” with an object for name lookup:

```
o.f();           // look into the scope of o's type, no "using" needed
```

and nonmember functions can be equally part of the interface and should be treated the same way:

```
f(o);           // look into the namespace containing o's type, no "using" needed
```

because some functions on a type must be nonmembers (notably operators `<<` and `>>`) and clearly ought to be first-class members of the type's interface. Languages like C# and Java don't have this problem only because there are no nonmember methods. ADL exists because people realized that the nonmember function parts of a type's interface ought to “come along for the ride” and be naturally usable just as the member functions already come along for free (because of the `.mf()` and `->mf()` syntax which implies the scope of the object's type).

So ADL is all about ensuring that, just as members “come along” without special namespace qualification (because you allude to the scope of the `this` argument with the `o.f()` syntax), so nonmembers that are equally part of the delivered interface of a type should come along for free without special namespace qualification (because you allude to the scope of the argument via ADL in the `f(o)` case).

There is a clear need for nonmember function names to be looked up in the scope of an argument's type's declaration. The canonical example is that otherwise the usual “Hello, world” wouldn't compile, because the required `operator<<` is declared in namespace `std`:

```
// Example 1
//
#include <iostream>

int main() {
    std::cout << "Hello, world" << std::endl;    // error, without ADL
}
```

Absent ADL, the choices to fix Example 1 would be: a) to qualify the `<<` call which defeats the purpose of operator overloading to get natural infix notation; b) to write a using-declaration for `std::operator<<` which would pull in all the other `<<` functions as well (e.g., the ones that work on iterators); or c) to write a using-directive for `std` (some people still really detest the latter, and although I don't agree with that I do think it would be embarrassing if that was the only reasonable way to make this code work).

For more discussion on why ADL is desirable, see [Sutter 2000] Items 31-34. In particular, the key guiding principle in those Items is the following:

“The Interface Principle

For a class X, all functions, including free functions, that both

(a) “mention” X, and

(b) are “supplied with” X [i.e., are in the same namespace as X]

are logically part of X, because they form part of the interface of X.”

For example, clearly the `std::operator<<` template that takes `basic_string` parameters, which cannot be a member function, is nevertheless just as much a part of the supplied interface of `basic_string` as are `basic_string`'s member functions `operator+`, `find`, etc. They are all part of the algebra of operations that “come in the box” of things that you can do with `basic_string` objects. Equally clearly, not every nonmember function that happens to mention `basic_string` is part of the interface; for example, a random function `void f(std::string)` in a client's namespace is not part of the supplied interface of `basic_string`. (See [Sutter00] for further justification of the IP.)

It is because ADL is valuable, and because of the principle underlying it, that [Sutter/Alexandrescu 2005] includes the Item:

“57. Keep a type and its nonmember function interface in the same namespace.”

1.2. Why Not ADL

But it is because ADL is overeager that [Sutter/Alexandrescu 2005] also includes the Item:

“58. Keep types and functions in separate namespaces unless they're specifically intended to work together.”

This is one of the things programmers have to learn today to steer clear of ADL problems. This paper will point out that Item 58 above exists to avoid the problem that this paper would resolve.

The error in ADL is that it pulls in all functions in the namespace of the type, rather than only the functions that were part of the interface of the type (i.e., those that specifically mention the type). When ADL looks up too much, unintended and unwanted names are found, causing us to select the wrong function (in the case when the function found via ADL is a better match, which usually happens with unconstrained templates) or to give deeply mysterious error messages.

2. Problem Examples

2.1. Unspecified behavior for simple code: N1691, Example 1

This example is based on [Abrahams 2004]:

```

// Example 2.1
//
#include <vector>

namespace user {
  class Customer { };

  typedef std::vector<Customer*> index;

  // copy an index, selecting deep or shallow copy for Customers
  void copy( index const&, index&, int deep ) { }

  void g( index& x ) {
    copy( x, x, false );           // legal? not specified by the standard
  }
}

```

Because the argument `x` has a type declared in namespace `std`, ADL reaches over and considers `std::copy` — and `std::copy` happens to be a better match. Additionally, whether ADL will find `std::copy` or not depends on whether `<vector>` includes `<algorithm>`, which is left up to the implementation. I view it as embarrassing that Example 2.1 is nonportable code, and that the standard doesn't say what it does.

Under this paper's proposal, the Example 2.1 code would be well-specified and call `user::copy`.

Note: This example could also be fixed by using `enable_if` to constrain `std::copy` to take only iterators. The following example cannot be fixed this way.

2.2. N1691, Example 2

[Abrahams 2004] also contained this example, which also unintentionally looks up an unconstrained template via ADL, and again finds it to be a better match:

```

// Example 2.2
//
#include <shared_ptr>
#include <tuple>

namespace communications {
  class Channel;
}

namespace user {
  typedef std::shared_ptr<communications::Channel> pipe;

  // Connect p2 to the output of p1
  void tie( pipe const& p1, pipe& p2 );

  void g( pipe p1, pipe p2 ) {
    tie(p1, p2);           // oops, calls std::tie
  }
}

```

(Note that here the constness of the reference is sufficient to make `user::tie` a worse match.)

Note: We are not aware of a way to fix this example by using `enable_if` to constrain `std::tie`, because `tuple` template parameters (which `tie` is meant to match) must be unconstrained.

Under this paper's proposal, the Example 2.2 code would be well-specified and call `user::tie`.

2.3. From the Real World

This example from `comp.lang.c++.moderated` is similar to Example 2.4, but more subtle.

```
// Example 2.3
//
#include <vector>

namespace N {
    struct X { };

    template<typename T>
    int* operator+( T , unsigned )
        { static int i ; return &i ; /* just to stub in the function body */ }
}

int main() {
    std::vector< N::X > v(5);
    v[ 0 ];
}
```

Question: Will this program compile? Answer: Maybe yes, maybe no, and when it fails it fails for different and unrelated reasons on different combinations of compilers and standard library implementations.

First, just to acknowledge it and get it out of the way: Is this the programmer's fault (pilot error)? After all, it's true that he did write an unconstrained template. But, well, so does namespace `std`... therefore either he's in good company, or it's the pot calling the kettle black. (He did put the operator template into his own namespace, after all, so he's clearly trying to be a good citizen and play nice.) So no, we can't just blame the programmer here.

I'll demonstrate the problems by extracting code and showing analysis for both `libstdc++ 3.2` and `Dinkumware 3.13`, and why they fail (always or sometimes) and why. Let's first look at an edited extract from `libstdc++ 3.2`:

```
// Example 2.3(b): An extract from libstdc++ 3.2, distilled and edited
//
namespace __gnu_cxx {
    template<typename T> class __normal_iterator {
    public:
        __normal_iterator operator+(const long& __n) const {...}
    };
}

namespace std {
    template <class T> class vector {
```

```

public:
    typedef __gnu_cxx::__normal_iterator<T> iterator;
    iterator operator[](unsigned __n) { return iterator() + __n; }
};
}

```

In this case, we might argue that this is a problem with ADL. Under the current rules `N` is an associated namespace of `__normal_iterator<N::X>`, so `N::operator+<N::X>(unsigned)` should be considered, and it happens to be a better match. Yes, the library implementer could have written the code more defensively. But why should he have to? (This goes to the heart of Abrahams' argument in [Abrahams 2004].)

But what about a library that *is* written to avoid ADL? From Dinkumware 3.13:

```

// Example 2.3(c): An extract from Dinkumware 3.13, distilled and edited
//
namespace std {
    template<class T> struct _Ranit { };
    template<class T> class vector {
    public:
        class iterator : public _Ranit<T> {
        public:
            iterator operator+(int) const { return iterator(); }
        };
        iterator operator[](unsigned _Pos) { return iterator()+_Pos; }
    };
}

```

Compiling this code under various compilers produces various results. This is a case of “Compiler Bugs Bite,” because `N` is *not* an associated namespace of `std::vector<N::X>::iterator` (the `_Ranit<T>` base doesn't matter). But all shipping versions I tried of Gnu g++ (up to 3.4.3) and Microsoft VC++ (up to 7.1) do incorrectly find it. Comeau 4.3.3 and the current beta of Microsoft VC++ 8 are two compilers that correctly do not find it.

An even simpler variant:

```

namespace std {
    template<class T> class vector {
    public:
        class iterator { // removed inheritance from _Ranit<T>
        public:
            iterator operator+(int) const { return iterator(); }
        };
        iterator operator[](unsigned _Pos) { return iterator()+_Pos; }
    };
}

```

This workaround prevents incorrect ADL on some of the incorrect compilers, but not others which even here still consider `N` an associated namespace. So there are two different compiler bugs at work

here: Incorrectly looking into the namespace of a base class template argument, and incorrectly looking into the namespace of an enclosing class template argument.

Interestingly, Example 2.3 compiled fine on other compilers and standard library implementations, including overlaps with the above-mentioned compilers and libraries. Why did it work there? Here's a sampling:

- **STLport, libcomo, Borland's Rogue Wave version:** In these implementations, when I tried them, `vector<T>::iterator` is just `T*`, so the operation is `T*+val`, not `iter+val`. Why does that make everything okay? Because ADL still looks into namespace `N`, but `N::operator+<T*>(T*, unsigned)` is not found because it doesn't exist – you can't replace the builtin `operator+` for pointers, so a general `operator+` template won't generate it! (Note: Always using raw pointer iterators isn't a "solution"... there are important benefits to having iterators of class type, e.g., "safe STL" iterator checks.) But surely "because you don't replace built-in operators" is an *extremely* subtle reason to rely on for why these implementations didn't break on this example.
- **Metrowerks:** Metrowerks's `vector<T>::iterator` can be configured to be many things, including `T*` and `__debug_iterator<vector<T>, T*>`. But `operator[]` always does "`data() + n`" (not `iterator() + n`) – so again that's always `T*+val`, not `iter+val`.

Clearly, neither the standard nor the implementations have been in great shape when it comes to ADL. Those that fail, fail for various unrelated and inconsistent reasons. Those that work, work for exceedingly subtle reasons and only narrowly escape falling over the brink. This isn't a great place to be.

The good news is that this is easy to fix in the standard, and that the fix in the standard will make it easier (not harder) for implementations to conform correctly.

Under this paper's proposal, the Example 2.3 code would be well-specified and would not consider the user-supplied template in namespace `N`.

2.4. Pointers as iterators

Most people agree that ADL looks up too much, as already shown.

Some people, however, feel that ADL doesn't look up *enough*. Specifically, pointers (which can be used as iterators) have no associated type and therefore can't trigger ADL like a full iterator type can do. This is an inconsistency between pointers, which are supposed to be iterators over arrays, and other iterators. For example:

```
// Example 2.4 (a): Today's situation
//
// code in some other namespace:
int a[10], b[1000];
```

```

int *astart = &a[0], *aend = &a[6];
std::copy( astart, aend, &b[0] );           // qualification IS required

std::istream_iterator<int> istart(std::cin), iend;
copy( istart, iend, &b[0] );             // qualification IS NOT required

```

With pointers, `std::copy` is not looked up by ADL because a pointer has no associated namespaces. With `std::istream_iterators`, `std::copy` is looked up by ADL.

I agree this is inconsistent behavior between pointers and UDT iterators, but I believe that the problem is not that pointers don't trigger ADL, but that UDT iterators *do* trigger ADL; that is, ADL lets you get away with not qualifying `std::copy` — it would be consistent to require that a call to `std::copy` be qualified (`copy` is after all a general-purpose function template, and not merely a part of the interface of a specific iterator type).

For one thing, in the same example with *container* iterators, the programmer must qualify the call to have portable code:

```

std::deque<int> d;
std::deque<int>::iterator dstart = d.begin(), dend = d.end();
std::copy( dstart, dend, &b[0] );           // qualification IS required

```

With pointers, `std::copy` is not looked up by ADL because a pointer has no associated namespaces. With `std::deque<T>::iterators`, `std::copy` is looked up by ADL *only if* `deque<T>::iterator` is a type defined in namespace `std` (it is not always; see for example Example 2.3(b) about `libstdc++`); otherwise, qualification is still required. This means that, today, portable code must qualify the above example.

Alternatively, if the previous example instead used a `vector`, then the developer doesn't know whether he has to qualify it (because `vector::iterator` might be either a `std::` class type or just a pointer), and this can change just from going from a debug to a release build:

```

std::vector<int> v;
std::vector<int>::iterator vstart = v.begin(), vend = v.end();
???:copy( vstart, vend, &b[0] );           // qualification MIGHT BE required

```

Also, notice that qualification is already required not only for pointers but also for every user-defined iterator type (unless the container is instantiated with a type from `std::`):

```

My::Container<int> m;
My::Container<int>::iterator mstart = m.begin(), mend = m.end();
std::copy( mstart, mend, &b[0] );           // qualification IS required

```

So this is already unpredictably inconsistent to the user, and the problem is that ADL still pulls in too much, giving unwarranted “convenience” to the standard library's own iterator types in namespace `std` (and to any iterators a vendor might also add to `std`), and only sometimes.

This paper's position is that the desirable outcome would be:

```

// Example 2.4 (b): Under this proposal...
//
// code in some other namespace:

```



```

int a[10];
int *astart = &a[0], *aend = &a[6];
std::copy( astart, aend, &b[0] );           // qualification required

deque<int> d;
deque<int>::iterator dstart = d.begin(), dend = d.end();
std::copy( dstart, dend, &b[0] );         // qualification required

My::Container<int> m;
My::Container<int>::iterator mstart = m.begin(), mend = m.end();
std::copy( mstart, mend, &b[0] );        // qualification required

```

Under the proposed change, the qualification would be uniformly required in all the above cases.

2.5. A caution about namespaces and helpers

Note that a library implementation (including an implementation of the standard library) might choose to define helper functions, possibly templates, inside its own namespace. Because today these can easily be picked up via ADL whenever users of the library use a library type from the same namespace, today the only way for a library to have “really private” helpers is to put them in their own nested namespace, and specifically one that contains no types. With the proposed change, most such helpers would automatically be made immune from inadvertent lookup, as the library author intended.

3. Non-Solutions

3.1. “Operator << should be global”

In response to Example 1, one claim I’ve heard is that we should have designed the standard library differently, specifically that:

“If operator<< were defined as global this would cease to be a problem.” —[Anon]

There are several serious problems with that, but pointing out one is sufficient: As illustrated in real-world code in sections 2.1-3, a global `operator<<` will be hidden by any other `operator<<` in a namespace, with the result that code in that namespace (or in further nested scopes such as classes inside that namespace) cannot use the global operator without qualification. That is unacceptable because the whole point of operator overloading is to facilitate the use of the operator using natural operator notation.

For example, this option fails to work in the simplest case:

```

// assume the standard library put its operator<< functions
// here in global scope, instead of in namespace std
...

```

```

namespace N1 {
    // provide a class X that happens to be streamable
    class X { ... };
    std::ostream& operator<<( const X&, std::ostream& );
    // now write a function...
    void f() {
        std::cout << "Hello, world" << std::endl;    // error, global operator is hidden
    }                                                // the programmer probably gets a weird error message that
}                                                    // there's no conversion from const char[13] to N::X... huh?!

```

Besides, putting more names in the global namespace is pollution, and if this were to be our recommendation then namespaces really would have failed (which I don't believe is true at all).

4. Proposal

4.1. Restricting ADL

The key proposed change is to not look up unconstrained templates in an associated namespace. ADL should look up only those functions, including template instantiations, that could exactly match the argument that triggered the ADL in the first place (i.e., that have a parameter type that is the same as, or an accessible base of, the argument type, in the same parameter position).

For example, the expression `std::cout << std::endl` would look into namespace `std` for a function named `operator<<` but consider only those free functions that have a first parameter of type `std::ostream` (or bases thereof, and possibly array-, reference-, or cv-qualified).

4.2. Extending using

4.2.1. Motivation

Consider this example that rely on ADL to find a function in an associated namespace even though the function does not mention the types that triggered the lookup:

```

// Example 4.2.1(a)
namespace N {
    struct S {};
    template<typename T> void forward( T, int ) {}
}

void test() {
    N::S s;
    forward( s, 2 );    // ok in today's ADL, error under this proposal
}

```

Today's ADL looks up `N::forward`, the proposed ADL would not. To make this possible under the proposed restricted ADL, we can extend `using` to allow the same effect to be restored by writing:

```
void test() {
    N::S s;
    using forward; // find any function/template in associated namespaces
    forward( s, 2 );
}
```

The second example is:

```
// Example 4.2.1(b)
namespace N {
    struct S {};
    template<typename T> void forward( T, int ) { }
}

template<typename U>
void test() {
    iterator_helper<U> i;
    forward( i, 2);
}

test<N::S>(); // ok in today's ADL, error under this proposal
```

Today's ADL looks up `N::forward`, the proposed ADL would not. To make this possible under the proposed restricted ADL, we can extend `using` to allow the same effect to be restored by writing:

```
template<typename U>
void test() {
    iterator_helper<U> i;
    using forward;
    forward( i, 2);
}

test<N::S>(); // ok
```

4.2.2. Proposal

For both using-declarations and using-directives, permit `using func;` to find any `func` in the associated namespaces of the arguments to `func`.

4.3. Compatibility

4.3.1. Fixed code

As shown in Examples 2.1 through 2.3, this change fixes code that people are trying to write today, but that is an error (or unspecified and nonportable) according to the current definition of ADL.

4.3.2. Calling standard algorithms

One example where the proposed change would break code that is currently correct is when a user is relying on ADL to get a function that is actually unrelated to (i.e., not part of the interface of) the types he's passing. To me this seems like something the user is just getting away with anyway, and the simple fix is to qualify the function name to explicitly reach over into the other namespace. For example:

```
#include <vector>

int main() {
    std::vector<int> v1, v2;
    some_algorithm(v1, v2);           // programmer would have to qualify
}                                     // with std:: to reach into std
```

We could optionally also preserve the meaning of such code, by supplying an overload (not specialization) of `some_algorithm` for `vectors` in namespace `std`. In most cases, exactly such an overload is already there; for example, there is already such an overload of `swap` for `vectors` which explicitly makes `swap` part of the interface of `vector`.

There is some evidence that requiring qualification would not be surprising to users. From `comp.lang.c++.moderated`, April 8, 2005, comes a note from a user who was surprised that the qualification wasn't required:

```
#include <vector>
std::vector<int> v;
transform(v.begin(), v.end(), v.begin(), Delete_ptr);
```

This code compiles even though I forgot #include <algorithm> and using std::transform;

4.4. Overlaps with other proposals

4.4.1. Unified call syntax proposal

There have been proposals to add a unified call syntax, whereby for example `f(a,b)` could potentially call a member `a.f(b)` or nonmember `f(a,b)` (yes, adding such lookup is problematic because it could break code, but it's worth considering and it keeps being suggested for reasonable reasons including to make it easier to write more flexible templates that aren't sensitive to whether a function or operator that could be either a member or a nonmember is one or the other).

If such a proposal is ever accepted, then it will be even more important to have consistent treatment of members and nonmembers. I say that because clearly when you say `a.f(b)` you want lookup in the namespace where `a`'s type is defined, and the meaning should be the same if you use the syntax `f(a,b)`. This is what ADL is all about, except that ADL picks up more than just functions matching the argument types, which is the problem that I'd like to fix.

4.5. Proposed wording

~~Strikeout~~ denotes removed text, and underline denotes inserted text.

In 3.4.2, modify paragraph 4:

- 4 When considering an associated namespace of a given argument type T in the function call, the lookup is the same as the lookup performed when the associated namespace is used as a qualifier (3.4.3.2) except that:
 - Any *using-directive* s in the associated namespace are ignored.
 - Any namespace-scope friend functions declared in associated classes are visible within their respective namespaces even if they are not visible during an ordinary lookup (11.4).
 - A function is considered if and only if it does not have an unconstrained template parameter in the parameter position of T.

In 7.33, 7.3.4, and A.6, extend the following productions:

using-declaration:

`using typenameopt ::opt nested-name-specifier unqualified-id ;`

`using :: unqualified-id ;`

`using unqualified-id ;`

In 7.3.3, append to paragraph 1:

The `using unqualified-id ;` form has the effect of a set of *using-declarations* for the *unqualified-id*, one for each associated namespace of the *identifier*.

5. Acknowledgments

Thanks to Dave Abrahams, Gabriel Dos Reis, Howard Hinnant, Thorsten Ottosen, John Spicer, Daveed Vandevoorde, and others for their very helpful comments on drafts of this material.

6. References

[Abrahams 2004] D. Abrahams. “Explicit Namespaces” (ISO/IEC JTC1/SC22/WG21/N1691 = ANSI/INCITS J16/04-131).

[Anon] Private correspondence.

[Sutter 2000] H. Sutter. *Exceptional C++* (Addison-Wesley, 2000).

[Sutter/Alexandrescu 2005] H. Sutter and A. Alexandrescu. *C++ Coding Standards* (Addison-Wesley, 2005).