

Doc No: WG21/N2100=J16/06-0170

Date: 2006-09-09

Reply to: Bjarne Stroustrup
bs@cs.tamu.edu

Initializer lists (Rev 2.)

Bjarne Stroustrup and Gabriel Dos Reis
Texas A&M University

Abstract

This paper presents a synthesis of initialization based on consistent use of initializer lists. The basic idea is to allow the use of initializer lists wherever initialization occurs and to have identical semantics for all such initialization.

For user-defined types, initialization by an arbitrarily long list of values of a specified type can be defined by a sequence constructor.

The discussion is based on the earlier papers and on discussions in the evolution working group. Much of this paper summarizes the discussions of alternatives.

In addition to the main discussion and proposal, we present two subsidiary proposals: (1) to allow the use of initializer lists as sub-expressions and (2) to disallow narrowing in initializations using initializer lists. The latter would put an end to many narrowing problems.

If this proposal is accepted, we will propose that a sequence constructor be added to each standard library container.

Suggested working paper text is an appendix (yet to be completed).

1 Previous work

The direct ancestor to this paper is N1919==05-179, which in turn was based on N1890=05-0150 “Initialization and initializers”. The other parts of “the initialization puzzle” presented in N1890 are presented in companion papers, such as Gabriel Dos Reis and Bjarne Stroustrup’s “Generalized constant expressions” (N1920=05-0180). Here is a list of problems and suggested improvements that has led to the current design:

- General use of initializer lists (Dos Reis & Stroustrup N1509, Gutson N1493, Meredith N1806, Meredith N1824, Glassborow N1701)

- There are four different syntaxes for initializations (Glassborow N1584, Glassborow N1701)
- C99 aggregate initialization (C99 standard)
- Type safe variable length argument lists (C++/CLI)
- Overloading “new style” casts
- Making **T(v)** construction rather than conversion (casting)
- Variadic templates (N1603 and N1704)

In each case, the person and paper referred to is just one example of a discussion, suggestion, or proposal. In many cases, there are already several suggested solutions. This is not even a complete list: initialization is one of the most fertile sources of ideas for minor improvements to C++. Quite likely, the potential impact on the programmer of sum of those suggestions is not minor. In addition to the listed sources, we are influenced by years of suggestions in email, newsgroups, etc. Thanks; apologies those who contributed, but are not explicitly mentioned here.

2 Summary

As the result of the detailed discussion presented in the following sections we propose:

- To allow an initializer list (e.g., **{1,2,3}** or **={1,2,3}**) wherever an initializer can appear (incl. as a return expression, an function argument, a base or member initializer, and an initializer for an object created using **new**). An initializer list appears to the programmer as an rvalue.
- To introduce type **std::initializer_list** for the programmer to use as an argument type when an initializer list is to be accepted as an argument. The name **initializer_list** is known to the compiler, but its use requires including the definition of **initializer_list** from namespace **std**.
- To distinguish sequence constructors (a single-argument constructor with a **initializer_list** argument type) in the overload resolution rules.
- To use a type name to indicate the intended type of an initializer list (e.g., **X{1,2,3}**). This construct is primarily for disambiguation.
- To allow an initializer list to be used as arguments to a constructor of a class when no sequence constructor can be used (e.g. **f({1,2})** can be interpreted as **f(X(1,2))** when **f()** unambiguously takes an **X** argument and **X** does not have a sequence constructor for **{1,2}**). This mirrors the traditional (back to K&R C) use of initializer lists for both arrays and **structs** and is needed for initializer lists to be used for all initializers, thus providing a single notation with a single semantics for all initialization.
- Initialization using an initializer list, for example **X x = { y };** is direct initialization, not copy initialization.
- A separate subsidiary proposal is to disallow narrowing conversions when using the initializer list notation. For example, **char c = { 1234 };** would become an error. See §7

- A separate subsidiary proposal is to allow an initializer list as a sub-expression, for example, `x=y+{1,2}`. See §6.

This proposal offers initializer lists as a uniform notation for all initialization with a single semantics for all cases. The main proposal breaks no legal ISO C++ program except those that uses the proposed standard library name `initializer_list` in a way that could clash.

3 Problems: Four ways of providing an initializer

Initialization of objects is an important aspect of C++ programming. Consequently, a variety of facilities for initialization are offered and the rules for initialization have become complex. Can we simplify them? Consider How to initialize an object of type `X` with a value `v`:

```
X t1 = v;      // “copy initialization” possibly copy construction
X t2(v);      // direct initialization
X t3 = { v }; // initialize using initializer list
X t4 = X(v); // make an X from v and copy it to t4
```

We can define `X` so that for some `v`, 0, 1, 2, 3, or 4 of these definitions compile. For example:

```
int v = 7;
typedef vector<int> X;
X t1 = v;      // error: vector’s constructor for int is explicit
X t2(v);      // ok
X t3 = { v }; // error: vector<int> is not an aggregate
X t4 = X(v); // ok (make an X from v and copy it to t4; possibly optimized)
```

and

```
int v = 7;
typedef int X;
X t1 = v;      // ok
X t2(v);      // ok
X t3 = { v }; // ok; see standard 8.5; equivalent to “int t3 = v;”
X t4 = X(v); // ok
```

and

```
int v = 7;
typedef struct { int x; int y; } X;
X t1 = v;      // error
X t2(v);      // error
X t3 = { v }; // ok: X is an aggregate (“extra members” are default initialized)
```

```
X t4 = X(v); // error: we can't cast an int to a struct
```

and

```
int v = 7;
typedef int* X;
X t1 = v; // error
X t2(v); // error
X t3 = { v }; // error
X t4 = X(v); // ok: unfortunately this converts an int to an int* (see §7)
```

Our aim is a design where a single notation where for every (**X**,**v**) pair:

- Either all examples are legal or none are
- Where initialization is legal, all resulting values are identical

This proposal meets these goals for every use of initialization in the language except that it is still possible to disallow pass-by-value by disabling a copy constructor. For example:

```
X a= {v};
void f(X);
f({v});
```

Here the **X** value constructed for **v** for **a** and for **f()** will be the same for all types **X** and values **v**, but the by-value passing of that **X** to **f()** may be disallowed.

The next three subsections consider further problems with reaching a uniform initialization style and semantics without introducing new syntax. We conclude that we must live with different meanings for the existing different initialization syntaxes. It is possible that we have missed a satisfactory solution to this puzzle, but having looked repeatedly we haven't found one and we don't propose to spend more time on this.

3.1 Can we eliminate the different forms of initialization?

It would be nice if we didn't need four different ways of writing an initialization. Francis Glassborow explains this in greater detail in N1701. Unfortunately, we lose something if we eliminate the distinctions. Consider:

```
vector<int> v = 7; // error: the constructor is explicit
vector<int> v(7); // ok
```

If the two versions were given the same meaning, either

- both would be correct (and we would be back in “the bad old days” where all constructors were used as implicit conversions) or
- both would fail (and many programs using a **vector** or similar type would fail).

We consider both alternatives unacceptable. It follows that we cannot eliminate the distinction between copy initialization and direct initialization without serious compatibility problems.

Question: but why would anyone expect the `v = 7` notation to work? And if they did why would they expect it to have a different effect from the `v(7)`? Some people expect the `v = 7` example to initialize `v` with the single element `7`. Scripting languages supply a steady stream of people with that expectation.

3.2 *Can we eliminate differences caused by copying?*

In addition to the issue of implicit vs. explicit constructors, we have the issue of actual copying vs. construction “in place”. Assume that `X` is a type that we can initialize with an `int`; consider

```
void f(const X& v);
void g(X);
X v(1);      // copy not allowed
X v2 = 1;   // copy undesirable and easily avoided
f(1);       // copy undesirable and easily avoided
g(1);       // copy (almost) unavoidable
```

For the results of initialization to be exactly the same in all cases, we must either copy in all cases or in none. Copying in all cases is clearly undesirable because of the significant overhead it would impose compared to current C++. On the other hand, avoiding copying in every case is essentially impossible. Either choice would be incompatible, so we will have to live with copying in some cases (e.g. for call-by-value arguments) and not in others (e.g. in direct initialization of local variables). As ever, copy operations that just copy correctly are only visible as a factor in performance, a private copy constructor can cause an initialization to be illegal, and a copy constructor may “become visible” if it throws an exception.

There is currently one way of getting uniform initialization: Always use the most explicit form of initialization:

```
vector<int> v = vector<int>(7);    // copy?
X e3 = X(1);                      // copy?

template<class T> void f(T v);
f(vector<int>(7));                // copy
f(X(1));                          // copy
```

We cannot recommend that style for systematic use because it is unnecessarily verbose and also implies serious inefficiency unless compilers are guaranteed to eliminate most copy operations.

3.3 A constructor problem: explicit constructors

Explicit constructors can cause different behavior from different forms of initialization. Consider:

```

struct X {
    explicit X(int);
    X(double); // not explicit
};

X a = 1; // call f(double)
X b(1); // call f(int)

void f(X);
f(1); // call f(double)

```

The reason **f(double)** is called is that the explicit constructor is considered only in the case of direct initialization. We consider this backwards: what should happen is that the best matching constructor should be chosen, and the call then rejected if it is not legal. That would make the resolution of these cases identical to the cases where a constructor is rejected because it is private.

We don't make a proposal for that change here, but note this as a case where a difference in initialization behavior could be eliminated by a rule change. See also Section 6.1.1.

We furthermore conjecture that having both an explicit and a non-explicit constructor taking a single argument is poor class design.

4 Initializer lists

There is a widespread wish for more general use of initializer lists as a form of user-defined-type literal. The pressure for that comes not only from “native C++” wish for improvement but also from familiarity with similar facilities in languages such as C99, Java, C#, C++/CLI, and scripting languages. Our basic idea is to allow initializer lists for every initialization. What you lose by consistently using initializer lists are the possibilities of ambiguities inherent in = initialization (as opposed to the direct initialization using () and proposed { }).

Consider a few plausible examples:

```

X v = {1, 2, 3.14}; // as initializer
const X& r1 = {1, 2, 3.14}; // as initializer
X& r2 = {1, 2, 3.14}; // as lvalue initializer (will be an error)

void f1(X);
f1({1, 2, 3.14}); // as argument

```

```

void f2(const X&);
f2({1, 2, 3.14});           // as argument
void f3(X&);
f3({1, 2, 3.14});         // as lvalue argument (will be an error)

X g() { return {1, 2, 3.14}; } // as return value

class D : public X {
    X m;
    D() : X({1, 2, 3.14}), // base initializer
        m({1, 2, 3.14}) { } // member initializer
};
X* p = new X({1, 2, 3.14}); // make an X on free store X
                               // initialize it with {1,2,3.14}

void g(X);
void g(Y);
g({1, 2, 3.14});         // (how) do we resolve overloading?

X&& r = { 1, 2, 3 }; // rvalue reference

```

We must consider the cases where **X** is a scalar type, a class, a class without a constructor, a union, and an array. As a first idea, let's assume that all of the cases should be valid and see what that would imply and what would be needed to make it so. Our design makes these examples legal, with the exceptions of the lvalue examples. We don't propose to make initializers lvalues.

Note that this provides a way of initializing member arrays. For example:

```

class X {
    int a[3];
public:
    X() :a({1,2,3}) { } // or just :a{1,2,3}
};

```

Some people consider this important. Over the years, there has been a slow, but steady, stream of requests for some way of initializing member arrays.

4.1 The basic rule for initializer lists

The most general rule of the use of initializer lists is:

- Look for a sequence constructor and use it if we find a best one; if not
- Look for a constructor (excluding sequence constructors) and use it if we find a best one; if not
- Look to see if we can do traditional aggregate or built-in type initialization; if not

- It's an error

We propose to retain the slightly more restrictive rule “never use aggregate initialization if a constructor is declared”. Without this restriction, we would not be able to enforce invariants by defining constructors. Consequently, we consider this restriction necessary and get this modified basic rule:

- If a constructor is declared
 - Look for a sequence constructor and use it if we find a best one; if not
 - Look for a constructor (excluding sequence constructors) and use it if we find a best one; if not
 - It's an error
- If no constructor is declared
 - look to see if we can do traditional aggregate or built-in type initialization; if not
 - It's an error

This can (and should) be integrated into the overload resolution rules.

We could consider an even simpler rule: use initializer lists only for sequence constructors, but that would leave the problem of non-uniform initialization semantics unaddressed.

4.2 *Sequence constructors*

A sequence constructor is defined for a class like this:

```
class C {
    C(initializer_list<int>); // construct from a sequence of ints
    // ...
};
```

The **initializer_list** argument type indicates that the constructor is a sequence constructor. The type in <...> indicates the type of elements accepted. A sequence constructor is invoked for an array of values that can be accessed through the **initializer_list** argument. The **initializer_list** is a standard library class that offers three member functions to allow access to the sequence:

```
template<class E> class initializer_list {
    // representation (a pair of pointers or a pointer plus a length)
public:
    constexpr initializer_list(const E*, const E*);    // from [first,last)
    constexpr initializer_list(const E*, int);        // from [first, first+length)

    constexpr int size() const;                       // number of elements
    constexpr const T* begin() const;                // first element
```



```

        constexpr const T* end() const;    // one-past-the-last element
    };

```

The **constexpr** specifier (N1980=0050) indicates that if an **initializer_list** object happens to be **constexpr**, the operations on it will be usable as constant expressions.

The three member functions provide STL-style (**begin()**,**end()**) access or “Fortran-style” (**first()**,**size()**) access. It is essential that the sequence is immutable: A sequence constructor cannot modify its input sequence. A sequence constructor might look like this:

```

template<class E> class vector {
    E* elem;
public:
    vector (initializer_list<E> s) // construct from a sequence of Es
    {
        reserve(s.size());
        uninitialized_fill(s.begin(),s.end(),elem);
    }
    // ... as before ...
};

```

Consider:

```

std::vector<double> v = {1, 2, 3.14};

```

That’s easily done: **std::vector** has no sequence constructor (until we add the one above), so we try **{1, 2, 3.14}** as a set of arguments to other constructors, that is, we try **vector(1,2,3.14)**. That fails, so all of the examples fail to compile when **X** is **std::vector**.

Now add **vector(initializer_list<E>)** to **vector<E>** as shown above. Now, the example works. The initializer list **{1, 2, 3.14}** is interpreted as a temporary constructed like this:

```

double tmp[] = {double(1), double(2), 3.14 } ;
initializer_list<double> tmp(tmp,sizeof(tmp)/sizeof(double));
vector<double> v(tmp);

```

That is, the compiler constructs an array containing the initializers converted to the desired type (here, **double**). This array is passed to **vector**’s sequence constructor as an **initializer_list**. The sequence constructor then copies the values from the array into its own data structure for elements.

Note that an **initializer_list** is a small object (probably two words), so passing it by value makes sense. Passing by value also simplifies inlining of **begin()** and **end()** and constant expression evaluation of **size()**.

4.3 *The initializer list rewrite rule*

A simple way of understanding initializer list is in terms of a rewrite rule. Given

```
void f(initializer_list<int>);
f({1,2.0,'3'});
```

The compiler lays down an array

```
int a[] = {int(1), int(2.0), int('3') };
```

And rewrites the call to

```
f(initializer_list<int>(a,3)); // rewritten to use initializer_list
```

Assuming that `initializer_list` is in scope (§4.5.1), all is now well.

In general, given

```
X v = {1,2.0,'3'};
```

the compiler looks at `X` and if it finds a sequence constructor taking a `initializer_list<Y>`, it lays down an array

```
Y a[] = { Y(1), Y(2.0), Y('3') };
```

and rewrites the definition to

```
X v(initializer_list<Y>(a,3)); // rewritten to use initializer_list
```

Thus, from the point of view of the rest of the language an initializer list that is accepted by a sequence constructor is simply an invocation of the suitable constructor.

For the purpose of overloading, going from an initializer list to its `initializer_list` object counts as a built-in conversion (as opposed to a user-defined conversion), independently of what conversions were needed to generate the homogenous array.

4.4 *Syntax*

In the EWG there were strong support for the idea of the sequence constructor, but initially no consensus about the syntax needed to express it. There was a strong preference for syntax to make the “special” nature of a sequence constructor explicit. This could be done by a special syntax

```
class X {
    // ...
    X{(const int*, const int* ); // construct from
```

```

// a initializer list of ints
// ...
};

```

or a special (compiler recognized) argument type. For example:

```

class X {
    // ...
    X(initializer_list<int>);    // construct from a initializer list of ints
    // ...
};

```

Based on extensive discussions, we prefer the **X(initializer_list<int>)** design, because this “special compiler-recognized class name” approach

- Hides the representation of the object generated by the constructor and used by the sequence constructor. In particular, it does not expose pointers in a way that force teachers to introduce pointers before initializer lists.
- Is composable: We can use **initializer_list<initializer_list<int>>** to read a nested structure, such as { {1,2,3}, {3,4,5}, {6,7,8} } without introducing a name for the inner element type.
- The **initializer_list** type can be used for any argument that can accept an initializer list. For example **int f(int, initializer_list<int>, int)** can accept calls such as **f(1, {2,3}, 4)**. This eliminates the need for variable argument lists (... arguments) in many (most?) places.

Finding a syntax for sequence constructors was harder – much harder – than finding its semantics. Here are some alternatives. Consider these possible ways of expressing a sequence constructor for a class **C<E>**:

```

template<Forward_iterator For> C<E>::C(For first, For last);
template<int N> C<E>::C(E(&)[N]);
C<E>::C(const E*, const E*);
C<E>::C({} const E* first, const E* last);
C<E>::C(E ... seq);
C<E>::C(... E seq);
C<E>::C(... initializer_list<T> seq);
C<E>::C(... E* seq);
C<E>::C ({}<E> seq);
C<E>::C(E{} seq);
C<E>::C(E seq{});
C<E>::C(E[*] seq); // use sizeof to get number of elements
C<E>::C(E seq[*]);
C<E>::C(const E (&)[N]); // N “magically” becomes the number of elements

```

And more. None provided the three advantages of the **initializer_list<E>** approach without other problems.

The hardest part of the design was probably to pick a name for the “special compiler recognized class name”. Had we been designing C++ from scratch, we would probably have chosen **C::C(Sequence<int>)**. However, all the short good names have been taken (e.g., **Sequence**, **Range**, and **Seq**). Alternatives considered included **seqinit**, **seqref**, **seqaccess**, **seq_access** **seq_init**, and **Seq_init**. Our choice, **initializer_list**, seems the most descriptive and the least obnoxious name that has not already been widely used; we hope that the extravagant length is a protection. A quick check using google found only one occurrence with that capitalization, and that was in a Java program. We suggest **initializer_list** rather than **Initializer_list** because initial lower case is the norm in the standard library.

The name **initializer_list** is not a keyword. Rather, it is assumed to be in namespace **std**, so you can use it for something unrelated. For example:

```
int initializer_list = 7;
```

Doing so is would probably not be a good idea, though, once people get used to the standard (library) meaning.

4.5 *The initializer_list class*

Some obvious questions:

- Is **initializer_list** a keyword? No, but.
- Must I **#include** a header to use **initializer_list**? Yes, **#include<initializer_list>**
- Why don't we use a constructor that takes a general STL sequence?
- Why don't we use a general standard library class (e.g. **Range** or **Array**)?
- Why don't we use **T(&)[N]**?
- Can the **size()** be a constant expression? Yes.

More detailed answers and reasoning follows.

4.5.1 Keyword?

Is **initializer_list** a keyword? No; it is a name in the standard library namespace and the compiler will use it. In particular, if you declare an argument of type **initializer_list<int>** and pass an initializer list to it, the compile will generate a call **std::initializer_list<int>(p,s)**, where **p** is the pointer to the start of the initializer list array and **s** is its number of elements. For example:

```
// won't compile unless std::initializer_list is in scope:
```

```
void f(std::initializer_list<int> s);
```

```
void g()
```

```

{
    int initializer_list = 7;
    f({1,2,3});    // ok: use std::initializer_list
}

```

If you don't declare `initializer_list` (e.g., by including `<initializer_list>`), you get compile-time errors.

4.5.2 Include header?

Must I `#include` a header to use it? Yes, you must include `<initializer_list>`.

4.5.3 Why don't we use `T(&)[N]`?

Using “a notation” would save us a keyword (or the moral equivalent of a keyword: a frequently used name in `std`, such as `initializer_list`) and make it clear that a core language facility was used. Using `T(&)[N]` in particular would make it clear that we were dealing with a fixed length homogenous list (that is, an array).

We have an aesthetic problem with `T(&)[N]`, which would transform into an educational problem an myths about its rationale. However, the critical problem is that relying on this would turn every function that takes an initializer list as an argument into a template. For example, we might have a simple function:

```
void f(int,int);
```

We might want to generalize this to deal with `N` integers:

```
template<int N> void f(int (&)[N]);
```

Unfortunately, each different argument list size generates its own specialization. For example:

```

f({1});
f({1,2});
f({1,2,3});

```

Each calls a different function. This implies code replication, inability to use `f()` in a dynamically linked library, and problems with overloading: no use of `f()` as a virtual function, for callbacks etc. That's too high a price to pay for solving a naming problem. This is especially so, as the fact that the length of the list is a constant is rarely particularly useful.

4.5.4 Why don't we use a constructor that takes a general STL sequence?

For example, for `vector`, why don't we just deem

```
template<class For> vector(For first, For last);
```

to be the sequence constructor for **vector**? First of all, it doesn't support the use of initializer lists for arbitrary arguments. For example

```
void f(int, int*, int*,int);
```

This should not be sufficient clue that **f()** was willing to accept **f(1, {2,3,4,5,6},7)** as a call. To avoid chaos, we need something more explicit.

Secondly, the overload resolution rules can't work as described unless a sequence constructor is distinguishable from other constructors (and we can't eliminate current uses of these "iterator constructors"). It would also be odd to accept the "iterator constructor" above as a sequence constructor for any sequence of **T** while rejecting a constructor taking two **int*** arguments as a sequence constructor. However,

```
X::X(int*,int*);
```

Just might be taking two unrelated integers, rather than a sequence. For example:

```
X a(new int(7), new int(9));
```

Finally, pairs of iterators are not trivially composable. For example, handling **{{1},{2,3}, {3,4,5}}** would require an intermediate named type with a sequence constructor to handle the sub-sequences **{1}**, **{2,3}**, and **{4,5,6}**.

4.5.5 General (std::) class?

Why don't we use a general standard library class (e.g. **vector**, **Range**, or **Array**)? The compiler-generated array that is the in-memory representation of the initializer list must be immutable. If not, we could be back to "the good old days of Fortran 2 where you could change the value of the literal **1** to **2**". For example, imagine that **initializer_list** allowed modification of the array:

```
int f()
{
    Odd_vector<int> v = { 1, 2, 3 };
    return v[0];
}
```

We would certainly expect **f()** always to return **1**. But consider

```
template<class T> class Odd_vector { // very odd
    // ...
    Odd_vector(initializer_list<T> s)
```

```

        {
            // copy from the array into the vector
            *s.begin() += 1;    // illegal, but imagine what if
        }
    }

```

Assuming (reasonably, according to the simple memory model presented in §4.3) that `{1,2,3}` defines a single array with initial value `{1,2,3}` repeatedly accessed by the sequence constructor, we can get

```

    cout << f(); // write 1
    cout << f(); // write 2
    cout << f(); // write 3
    ...

```

As each invocation of the sequence constructor modifies that array's first element. It follows that we cannot accept anything as our accessor to the underlying array unless it can keep the array immutable.

4.5.6 Constant expression?

Can the `size()` be a constant expression? Yes, but only when a use of `size()` is in the same translation unit as the initializer list and after it. Consider:

```

template<class T> class initializer_list {
    // ...
    constexpr int size() const { /* ... */ }
};

// ...

initializer_list<int> s = {1,2,3};

char a[s.size()];           // ok: size is a constant expression

```

Clearly, there is enough information to deduce that `s.size()` is `3`. Equally clearly, making `s.size()` a constant expression requires a special rule. The proposal for generalizing constant expressions (N1920=05-0180) shows how this can work. We are not sure whether this is really important or just something people thought interesting. However, it follows directly from the definition for `constexpr`. Similarly, the “in the same translation unit” restriction follows from the `constexpr` definition, which in turn simply reflects the underlying realities of separate compilation. For example:

```

// file 1:
void f(initializer_list<int>);

```

```

// ...
f({1,2,3});

// file 2:
void f(initializer_list<int> s)
{
    char a[s.size()];    // error: size is not a constant expression
    // ...
}

```

There simply isn't sufficient information in file2 to evaluate `s.size()` at compile time.

4.6 *Initializer lists and ordinary constructors*

When a class has both a sequence constructor and an “ordinary” constructor, a question can arise about which to choose. The resolution outlined in §4.1 is that the sequence constructor is chosen if the initializer list can be considered as an array of elements of the type required by the sequence constructor (possibly after conversions of elements). If not, we try the elements of the list as arguments to the “ordinary” constructors. The former (“use the sequence constructor”) matches the traditional use of initializer lists for arrays. The latter (“use an ordinary constructor”) mirrors the traditional use of initializer lists for **structs** (initializing constructor arguments rather than **struct** members). Appendix B discusses the decision to give priority to sequence constructors over “ordinary constructors” in quite some detail.

4.7 *Initializer lists, aggregates, and built-in types*

So what happens if a type has no constructors? We have three cases to consider: an array, a class without constructors, and non-composite built-in type (such as an **int**). First consider a type without constructors:

```

struct S { int a; double v; };
S s = { 1, 2.7 };

```

This has of course always worked and it still does. Its meaning is unchanged: initialize the members of `s` in declaration order with the elements from the initializer list in order, etc.

Arrays can also be initialized as ever. For example:

```

int d[] = { 1, 2, 3, 5, 8 };

```

What happens if we use an initializer list for a non-aggregate? Consider:

```

int a = { 2 };    // ok: a==2
                  // (as currently: there is a single value in the initializer list)

```



```

int b = { 2, 3 };      // error: two values in the initializer list
int c = {};          // ok: default initialization: c==int()

```

In line with our ideal of allowing initializer lists just about everywhere – and following existing rules – we can initialize a non-aggregate with an initializer list with 0 or 1 element. The empty initializer list gives value initialization. The reason to extend the use of initializer lists in this direction is to get a uniform mechanism for initialization. In particular, we don't have to worry about whether a type is implemented as a built-in or a user-defined type and we don't have to depart from the direct initialization to avoid the unfortunate syntax clash between () initialization and function declaration. For example:

```

X a = { v };
X b = {};

```

This works for every type **X** that can be initialized by a **v** and has a default constructor. The alternatives have well known problems:

```

X a = v;           // not direct initialization (e.g. consider a private copy constructor)
X b;               // different syntax needed (with context sensitive semantics!)
X c = X();         // different syntax, repeating the type name

X a2(v);           // use direct initialization
X b2();            // oops!

```

It appears that {} initialization is not just more general than the previous forms, but also less error prone.

We do not propose that surplus initializers be allowed:

```

int a = { 1, 2 };   // error no second element
struct S { int a; };
S s = { 1, 2 };    // error no second element

```

Allowing such constructs would simply open the way for unnecessary errors.

Discussion: Discussion: The standard currently says (12.6.1/2) that when an object is initialized with a brace-enclosed initializer list, elements are initialized through “copy-initialization” semantics. For uniformity and consistency of the initialization rules this should be changed to “direct-initialization” semantics. That will not change the semantics of current well-formed programs; it will make legal examples where the only problem was a private copy constructors.

5 Initializer list technicalities

As the saying goes “the devil is in the details”, so let's consider a few technical details to try to make sure that we are not blindsided.

5.1 Sequence constructors

Can a class have more than one sequence constructor? Yes. An initializer list that would be a valid for two (or more) sequence constructors is ambiguous. For example:

```

Class V {      // poor design that's asking for trouble
    V(initializer_list<int>);
    V(initializer_list<double>);
    V(int,double);
    // ...
};

V v = { 1, 2.1 };      // error

```

In other words, we don't look further after finding an ambiguity among sequence constructors (only if no sequence constructor matches).

Can a sequence constructor be a template? Yes. Note that a “yes” here implies that more than one sequence constructor is possible.

Can a sequence constructor be invoked for a sequence that isn't an initializer list? No. For example, there is no way that **f(1,2,3)** can invoke a sequence constructor for an argument type the way **f({1,2,3})** can.

5.2 What really is an initializer list?

The simplest model is an array of values placed in memory by the compiler. That would make an initializer list a modifiable lvalue. It would also require that every initializer list be placed in memory and that if an initializer list appears 10 times then 10 copies must be present. So, we propose that all initializer lists be rvalues. That enables optimizations:

- Identical initializer lists need at most be stored once (though of course that optimization isn't required).
- An initializer list need not be stored at all. For example, **z=complex{1,2}** may simply generate two assignments to **z**.
- An initializer list that consists exclusively of constant expressions can be stored in read-only memory.

The second optimization would require a clever compiler or literal constructors (§5).

Note that an initializer list that is to be read by a sequence constructor must be placed in an array. The element type is determined by the sequence constructor. Sometimes, it will be necessary to apply constructors to construct that array.

Initializer lists that are used for aggregates and argument lists can be heterogeneous and need rarely be stored in memory.

Must initializer lists contain only constants? No, variables are allowed (as in current initializer lists); we just use a lot of literals because that's the easiest in small examples.

Can we nest initializer lists? Yes (as in current initializer lists). For example:

```
vector<vector<int>> v = { {1,2,3}, {4,5,6}, {7,8,9} };    // a 3 by 3 matrix
```

A more interesting example might be

```
Map<string,int> m = { {“ardwark”,91}, {“bison”, 43} };
```

Assuming that map has a sequence constructor for from a **pair**<**string,int**>, this will work, correctly converting the literal strings to **strings**.

5.3 Ambiguities and deduction

An initializer list is simply a sequence of values. If we considered it to have a type, it is would the list of its element types. For example, the type of **{1,2.0}** would be **{int,double}**. This implies that we can easily create examples that are – or at least appears to be – ambiguous. We can create ambiguities among sequence constructors of a single class:

```
class X {
    X(initializer_list<int>);    // sequence constructor
    X(initializer_list<double>); // sequence constructor
    // ...
};

X x1 = { 1, 2.0 };    // error: ambiguous
X x2 = { 1, 2 };    // X(initializer_list<int>);
X x3 = { 1.0, 2.0 }; // X(initializer_list<double>);
```

The resolution rule for sequence constructors is the same as for (other) function arguments. Roughly: there has to be a function for which the element is a best match in all cases and there has to be an element for which that element is a better mach for one list element type than for all other list element types.

Once we have found the best match (if any) for a single class, we can also have ambiguities among sequence constructors for different classes:

```
class X {
    X(initializer_list<int>);    // sequence constructor
    // ...
};
```

```

};

class Y {
    Y(initializer_list<int>);    // sequence constructor
    // ...
};

class Z {
    Z(int,int);    // not a sequence constructor
    // ...
};

void f(X);
void f(Y);

void g(Y);
void g(Z);

f({1,2,3});    // error: ambiguous (f(X) and f(Y)?)
g({1,2,3});    // ok: g(Y)
g({1,2});    // ok: g(Y) (note: not g(Z));
g({1});    // ok

```

The overload resolution rules are basically unchanged: try to match all functions in scope and pick the best match if there is a best match. Note that a match to a sequence constructor in one class still takes priority over a match to an ordinary constructor in another class.

When an ambiguity is found, we need a way to resolve it, How do we resolve ambiguity errors from an initializer list? By saying what we mean; in other words by stating our intended type of the initializer list:

```

f(X{1,2,3});    // ok: f(X)
g(Z{1,2});    // ok: g(Z)

```

Apart from using { } rather than (), it's the same idea as the current techniques of using explicit constructor calls.

Discussion: We do not propose to allow an “unqualified initializer list” to be used as an initializer for a variable declared `auto` or a template argument. For example:

```

auto x = {1, 2, 3.14};    // error
template<class T> void ff(T);
ff({1, 2, 3.14});    // error

```

There is no strong reason not to allow this, but we don't want to propose a feature until we have a practical use in mind. If we wanted to allow this, we could simply "remember" the type of the initializer list and use it when the **auto** variable or template argument is used. In this case, the type of **x** would be **{int,int,double}** which can be converted into a named type when necessary. For example:

```

auto x = {1, 2, 3.14};           // remember x' is a {int,int,double}
vector<int> v = x;             // initialize v {1, 2, 3.14};
g(x);                         // as above

```

It's comforting to know that the concepts extend nicely even if we have no use for the extension.

5.4 *Initializer lists and templates*

Can an initializer list be used as a template argument? Consider:

```

template<class T> void f(const T&);

f({ });                       // error
f({1});
f({1,2,3,4,5,6});
f({1,2.0});                   // error
f(X{1,2.0});                 // ok: T is X

```

There is obviously no problem with the last call (provided **X{1,2.0}** itself is valid) because the template argument is an **X**. Since we are not introducing arbitrary lists of types (product types), we cannot deduce **T** to be **{int,double}** for **f({1,2.0})**, so that call is an error. Plain **{ }** does not have a type, so **f({ })** is also an error.

This leaves the homogeneous lists. Should **f({1})** and **f({1,2,3,4,5,6})** be accepted? If so, with what meaning? If so, the answer must be that the deduced type, **T**, is **initializer_list<int>**. Unless someone comes up with at least one good use of this simple feature (a homogeneous list of elements of type **E** is deduced to be an **initializer_list<E>**), we won't propose it and all the examples will be errors: No template argument can be deduced from an (unqualified) initializer list. One reason to be cautious here is that we can imagine someone getting confused about the possible interpretations of single-element lists. For example, could **f({1})** invoke **f<int>(1)**? No, that would be quite inconsistent.

5.5 *C99 style initializers with casts*

If we wanted to increase C99 compatibility, we could additionally accept the more verbose version:

```
f((X){1,2,3}); // ok: f(X)
g(((Z){1,2})); // ok: g(Z)
```

This is not something we propose. The C semantics require the initializer list to be an lvalue with weird results. Here is an example from the C99 standard [6.5.2.5 Compound literals]:

EXAMPLE 8 Each compound literal creates only a single object in a given scope:

```
struct s { int i; };
int f (void)
{
    struct s *p = 0, *q;
    int j = 0;
again:
    q = p, p = &((struct s){ j++ });
    if (j < 2) goto again;
    return p == q && q->i == 1;
}
```

The function `f()` always returns the value 1.

17 Note that if an iteration statement were used instead of an explicit `goto` and a labeled statement, the lifetime of the unnamed object would be the body of the loop only, and on entry next time around `p` would have an indeterminate value, which would result in undefined behavior.

There is a danger that the “semi-compatible” syntax might become popular in C++ just as “the abomination” `f(void)`. Also, there would be subtle incompatibilities between the C99 definition of such as `construct` and any consistent C++ view (see N1509).

Using the `X{x,y}` syntax rather than the `(X){x,y}` syntax will require people to introduce a `typedef` if they want to resolve to a built-in type such as `char*`. We think that is preferable to introducing an additional syntax with a semantics that subtly differs from C’s.

5.6 Refining the syntax

So far, we have used initializer lists after `=` in definitions (as always) and as function arguments. The aim is to allow an initializer list wherever an expression is allowed. In addition, we propose to allow the programmer to leave out the `=` in a declaration:

```
auto x1 = X{1,2};
X x2 = {1,2};
X x3{1,2};
X x4({1,2});
X x5(1,2);
```

These five declarations are equivalent (except for the name of the variables) and all variables get the same type (`X`) and value (`{1,2}`). Similarly, we can leave out the parentheses in an initializer after `new`:

```
X* p1 = new X({1,2});
X* p2 = new X{1,2};
```

It is never ideal to have several ways of saying something, but if we can't limit the syntactic diversity we can in this case at least reduce the semantics variation. We could eliminate these forms:

```
X x3{1,2};
X* p2 = new X{1,2};
```

However, since **X{1,2}** must exist as an expression, the absence of these two syntactic forms would cause confusion, and they are the least verbose forms. Note that **new X{1,2}** must be interpreted as “an **X** allocated on the free store initialized by **{1,2}**” rather than “**new** applied to the expression **X{1,2}**”. This is equivalent to the current rule for **new X(1,2)**.

Note that if we add a sequence constructor to **std::vector**, each of these definitions will create a **vector** of one element with the value **7.0**:

```
vector<double> v1 = { 7 };
vector<double> v2 { 7 };
vector<double> v3 ({ 7 });

auto p1 = new vector<double>{ 7 };
auto p2 = new vector<double>({ 7 });
```

We don't propose a special syntax for saying “this is a sequence: don't treat is as a constructor argument list”. The general resolution mechanism for resolving initializer list ambiguities will do. For example:

```
vector<double> v3 (initializer_list<int>{ 7 }); // redundant qualification
```

However, this qualification is redundant and should never be needed; see Appendix B.

Discussion: we think that the most likely confusion and common error from the new syntax will (as with the old initialization syntax) be related to initializer lists with very few (0, 1, or 2) arguments. Consider:

```
vector<double> v2 { 7 };
```

A naïve reader will have no way of knowing that this creates a **vector** of one **double** initialized to **7.0** and not a **vector** of seven **doubles**. Obviously, making the second interpretation the correct one would be even worse. Consider

```
vector<double> v0 { }; // a vector with no elements
vector<double> v1 { 7 }; // a vector with one element
```

```

// (not a vector with 7 elements initialized to 0)
vector<double> v2 { 7, 8 }; // a vector with two elements
// (not a vector seven elements initialized to 8)
vector<double> v3 { 7, 8, 9 }; // a vector with three elements

```

We feel that this must work as stated. See Appendix B for a detailed discussion of design choices in this area.

6 Initializer lists in expressions

We have discussed initializer lists in the context of initialization. However, we could imagine them used elsewhere. Logically, an initializer list could appear in any place where an expression could. We would need a reason to prohibit that.

6.1 Assignments

Assignments and initializations are closely related. For example, there is no real implementation difference between them for built-in types. Consider:

```

X v = {1,2};
v = {3,4};

```

Having accepted the initialization, it would be hard to argue that the assignment was illegal. After all, we define $\mathbf{x}=\mathbf{y}$ as (something like) $\mathbf{x.operator}=(\mathbf{y})$. For some suitable type \mathbf{X} , we could write the assignment as $\mathbf{v.operator}(\{3,4\})$ and have it work because now $\{3,4\}$ is an initializer. Provided that there is no problem with the syntax, this example must be accepted.

6.2 General expressions

Consider more general uses of initializer lists. For example:

```

v = v+{3,4};
v = {6,7}+v;

```

When we consider operators as syntactic sugar for functions, we naturally consider the above equivalent to

```

v = operator+(v,{3,4});
v = operator+({6,7},v);

```

It is therefore natural to extend the use of initializer lists to expressions. We have not explored the grammar for this in detail and suggest that it should be explored. We see no obvious problems with this general use of initializer lists and suspect that people will expect it to work if the simpler uses work. In particular, the grammar will have to be explored.

6.3 *Function calls*

Consider:

```
void f(const vector<int>&);
// ...
f({1,2,3});
```

Why two pairs of parentheses? Why not just use one:

```
f(1,2,3);    // usual parentheses with new semantics
f{1,2,3};    // new parentheses with new semantics
```

Without major changes this wouldn't make sense. The `{...}` syntax specifies arguments for creating a single object (here `a`, `vector`) whereas in general, `(...)` specifies initializes for a sequence of objects. The first alternative would obviously lead to a confusing mess, and we don't see how we could craft rules that don't change the meaning of existing programs. The second alternative is more tempting. Obviously, there would not be compatibility problems related to the semantics. However, we don't see sufficient gain so we don't propose this.

6.4 *Lists on the left-hand side*

Whether we should allow lists on the right hand side of an assignment is a separate issue. For example:

```
{a,b} = x;
```

We make no proposal or recommendation about this. It is a separate question. Obviously, if we allowed that, "initializer lists" on the left-hand side would have to contain lvalues.

7 Casting

When a user-defined type is involved, we can define the meaning of C-style casting `(T)v` and functional style construction `T(v)` through constructors and conversion operators. However, we cannot change the meaning of a new-style cast and `T(v)` is by definition an old-style cast so its default meaning implies really nasty casts (incl. `reinterpret_cast`) for some built-in type combinations. For example, `int(p)` will convert a pointer `p` to an `int`. This leads to two common suggestions:

- Allow user-defined `static_cast`, etc.
- Default `T(v)` to mean `static_cast<T>(v)` rather than `(T)v`.

The two suggestions are related because often the reason for wishing $\mathbf{T}(\mathbf{v})$ to mean `static_cast<T>(v)` is to be able to define it as a range-checked operation for some built-in type \mathbf{T} .

We have also heard the suggestion that $\mathbf{T}(\mathbf{v})$ should be “proper construction” and thus not allow narrowing conversions (e.g. `char(123456)`). However, the functional notation is used to be explicit about narrowing, so even though we like the idea, we consider banning narrowing by default would too radical.

We don’t propose to allow overloading of the new-style casts. If you want a different cast, you can define one using the same notational pattern, such as `lexical_cast<T>(v)`. The $\mathbf{T}(\mathbf{v})$ problem is worse: it basically defeats attempts to make casting safer and more visible. It also, takes the ideal syntax for the least desirable semantics. Unfortunately, it appears to be widely used for “nasty casts” (in correct code). For example:

```
typedef char* Pchar;
int i;
// ...
Pchar p = Pchar(i); // would usually require an obviously nasty reinterpret_cast
```

Basically, this means that we cannot change the meaning of $\mathbf{T}(\mathbf{v})$. This is really nasty for several reasons:

- Consider:

```
Pchar p = Pchar(i);
```

This looks innocent, but hides nasty code.

- When we write generic code, there is no other general syntax for construction:

```
template<class T, class V> void f(T t, V v)
{
    T x;
    // ...
    x = T(v);    // construct (but for some types it casts)
    // ...
}
```

We consider that a serious problem. The `{ }` syntax can be used as a remedy:

```
T{v}
```

Means “(direct) initialize \mathbf{v} to type \mathbf{T} ”. That is, $\mathbf{T}(\mathbf{v})$ will have the same value as the variable \mathbf{x} after `T x{v}`. Note that if \mathbf{T} has a sequence constructor, `T{v}` means “make a \mathbf{T} with a single element \mathbf{v} ”.

7.1 Can we ban narrowing for $T\{v\}$?

It is extremely tempting to outlaw narrowing in a $T\{v\}$ cast. However, we can't do that by itself. We must maintain the uniformity of $\{ \}$ initialization. After all, one of the main aims of generalizing initializer lists and encouraging their use is to address the problems with the diversity of meanings of other initialization notations. In particular, consider:

```
T{v}
T x{v};
T y = {v};
T a[] = {v};
```

The values of $T\{v\}$, x , y , and $a[0]$ must be identical.

That is, to get $T\{v\}$ as a “safe” cast, we would have to disallow narrowing in all such initialization. That's still very tempting because the amount of code affected will be “relatively minor”. However, remember that a “relatively minor” fraction of hundreds of million lines of C++ code could easily be far too much. Given the advantages of addressing the problem with narrowing we will explore this possibility. Please note that this proposal the ban narrowing for $\{ \}$ initialization (only) is separate for the main proposal for dealing with initializer lists.

First note that banning narrowing conversions for $\{ \}$ initialization cannot lead to “silent” change of meaning; it will simply cause previously legal C++ programs to be rejected by the compiler. For example:

```
char x = { 1 };           // error: 1 is an int
char a[] = { 'a', 'b', 'c', 0 }; // error: 0 is an int
```

This problem could be remedied by requiring the compiler to verify that no narrowing actually occurs:

```
char x = { 69 };        // ok
char y = { 1234 };     // error (assuming 8-bit chars)
```

For initializers that are literals, that's trivial and some current compilers already warn. That's the rule we propose. Note that whether narrowing would occur (if allowed) is often implementation defined.

That leaves initializer lists where the initializers are variables, such as:

```
void f(int a, int b, int c)
{
    char x = { a };           // error: a is an int
    char a[] = { a, b, c, 0 }; // error: a, b, c are ints
```

```

        // ...
    }

```

The proposal to ban narrowing is based on the conjecture that such cases are rare and has a high enough incidence of errors, especially portability errors, that the community would be willing to accept (not silent) errors.

7.1.1 Narrowing of function argument values

Consider

```

struct X {
        X(int);
};

X a(2.1);    // ok
X b = 2.1;  // ok
X c{2.1};   // error: narrowing

void f(X);
f(2.1);     // ok
f({2.1});   // error: narrowing

```

This would follow from a ban of narrowing where ever we use {...}. This is backwards in the sense that the default (no use of { } in ordinary calls) is less safe than the “odd” use with {...}. However, not doing it that way would break a lot of code.

7.1.2 History: why do we have the narrowing problem?

Are there any inherent benefits of implicit narrowing? Yes, consider:

```

void f(int i, double d)
{
        char c = i;
        int i2 = d;
        // ...
}

```

This is shorter than equivalent using casts (C-style):

```

void f(int i, double d)
{
        char c = (char)i;
        int i2 = (int)d;
        // ...
}

```

Or (C++ style):

```
void f(int i, double d)
{
    char c = static_cast<char>(i);
    int i2 = static_cast<int>(d);
    // ...
}
```

Some implicit casts, such as **double->int** and **int->char**, have traditionally been considered a significant – even invaluable – notational convenience. Others, such as **double->char** and **int*->bool**, are widely considered embarrassments. When Bjarne once asked around in the Unix room why implicit narrowing had actually been allowed. Nobody argued that there were a fundamental technical reason, someone pointed out the obvious potential for errors and all agreed that the reason was simply historical: Dennis Ritchie added floating point before Steve Johnson added casts. Thus, the use of implicit narrowing was well established before explicit casting became an option.

Bjarne tried to ban implicit narrowing in “C with Classes” but found that a combination of existing practice (especially relating to the use of chars) and existing code made that infeasible. Cfront, however, stamped out the **double->int** conversions for early generations of C++ programmers by providing long, ugly, and non-suppressible warnings.

Please note that the suggestion to ban narrowing does not actually touch these common examples. It relies on explicit use of { }.

8 Variadic templates

N1704 proposes a general and type safe method of passing both homogenous and heterogenous lists. Why don't we just use that proposal?

The major reason is that N1704 is a proposal for templates. We do not want to require that every variadic function should be a template. Doing so would imply the problems of code replication and the problems with defining virtual functions and (other) callbacks.

In addition, we worry that the heavy use of templates might make the proposal unsuitable for long initializer lists. For example,

```
vector<int> v = { 1,2,3, .... 1001, 1002, 1003 };
```

Consequently, we are of the opinion that the proposals address different problems and this is not the place for a details discussion of variadic templates.

9 Acknowledgements

Obviously, much of this initializer list and constructor design came from earlier papers and discussions. The main papers are listed in §1.

10 Appendix A: Suggested working paper changes

<<Incomplete pending further discussion of the proposal and design alternatives>>

Here are working paper changes for the main proposal and two subsidiary proposals. The two subsidiary proposals make sense only if the main proposal is accepted, but the main proposal does not depend on the subsidiary proposals.

10.1 Main proposal

We propose to allow initializer lists wherever an initializer can appear.

10.2 Narrowing proposal

We propose to ban narrowing conversions of values in initializer lists.

10.3 Syntax proposal

We propose to accept initializer lists as expressions.

10.4 Containers

We propose that each standard library container is provided with a sequence constructor for its element type.

11 Appendix B: Initializer lists and argument passing

This appendix is a discussion of the design alternatives to what we consider a key example: initializing a **vector**.

11.1 The problems

The design for initializer lists has two main aims:

- To provide a uniform initializer syntax and semantics
- To provide variable-length homogeneous lists of initializers

Any solution is constrained by the essential third aim

- Compatibility: Don't break old code

It is easy to meet any one of these three aims. Meeting all three simultaneously is a difficult puzzle.

The central problem with uniform initializer syntax is that the underlying semantics cannot be completely uniform, leaving open the possibility of ambiguity or (potentially surprising) implicit resolution of what could have been considered an ambiguity. Consider our key example:

- Create a **vector<int>** with 1 element initialized to the value 2.
- Create a **vector<int>** with the values 1 and 2 as its initial elements

How would we express that? The obvious answer seems to be:

```
vector<int> v1(1,2);           // C++03
vector<int> v2 = { 1, 2 };    // extend the C aggregate initialize list syntax
```

Remember that the declaration of **v1** is surprising to most until they have it explained and get used to it; it also causes its own problems for the type system (a clash with the template constructor for sequences, which will only be satisfactorily resolved by applying concepts). There is nothing fundamental about this solution. It relies on familiarity to let (...) hint at argument passing and ={...} hint at assignment to elements. In reality, both cases involve passing of arguments to a constructor and the “assignment to element” is of course initialization (rather than assignment). Obviously, this solution does not provide a uniform syntax for initialization. Furthermore, the syntactic differences don’t indicate semantic differences as clearly as we would like.

Consider a related example illustrating a compatibility concern:

```
struct S1 {
    int x,
    int y;
};

struct S2 {
    int x,
    int y;
    S2(int xx, int yy) : x(xx), y(yy) { }
};

S1 s1 = { 1, 2 };
S2 s2(1,2);
```

The initialization of **s1** uses {...} and the initialization of **s2** uses (...) to achieve exactly the same end. The syntax emphasizes a difference in the way that initialization is achieved even though a compiler might very well generate the exact same code in both cases. This is out character for C++ where we don’t usually introduce different syntax to distinguish between user-defined and built-in operations (semantics). For example, we use + for both built in add and a user-defined add and = for both built-in copy and a user-defined copy. We could (if we so chose) describe **S1** as having “the default member constructor” and allow

```
S1 s1(1,2); // not C++03 and not proposed here
```

This would be perfectly consistent with our treatment of default and copy constructors. Using parentheses throughout was one line of exploration for a uniform syntax, but it doesn't extend cleanly and compatibly to arrays, to variable-length argument lists or return values.

Now return to the key problem: How do we distinguish between a `vector<int>` with 1 element initialized to the value 2 and a `vector<int>` with the values 1 and 2 as its initial elements? Assume first that we use a uniform syntax for initialization:

```
vector<int> v2 { 1, 2 }; // one or two elements?
```

For this example, we could have used `vector<int> v1(1,2)` but – as mentioned – we found that to be a dead end. What choices do we have for this example?

- Sequence constructors take priority: It's a variable-length initializer list that happens to have two elements
- "Ordinary constructors" take priority: It's an argument list that matches one of vector's constructors
- All constructors are equal: It's an ambiguity error.

Deming it ambiguous saves people from some unpleasant surprises, so that's the first alternative to explore. For "ambiguous" to be an acceptable answer, the ambiguities must be relatively rare and easily resolved by the user. In terms of numbers of classes, this ambiguity is going to be rare. Most classes won't have sequence constructors and many of those that do, won't have constructors that can clash with a sequence constructor. Note that even `vector` doesn't suffer the problem for many (most?) element types. For example:

```
vector<int*> vp1 { 1, &obj }; // ordinary constructor
vector<int*> vp2 { &obj }; // sequence constructor
vector<int*> vp3 { 1 }; // ordinary constructor (initialize to 0)
```

Unfortunately, the examples that can/will cause problems are frequent and important: containers of elements of numeric types. Obviously, this importance is also why it would be unwise to accept a poor solution to the problem. So, assuming this is ambiguous:

```
vector<int> v2 { 1, 2 }; // one or two elements?
```

How do we resolve the ambiguity? We would need to have notations for resolving it both ways. Consider:

```
vector<int> v21 = initializer_list<int> { 1, 2 }; // two elements
vector<int> v22 ( 1, 2 ); // one element
```

This solution follows from compatibility and general principles. It suffers from two problems:

- The resolution to initializer list is verbose
- The resolution to argument list brings us back to “the old world”.

By defaulting the resolution to either of these alternatives, we can trade one of these problems for surprises to someone who expected the other resolution.

11.2 Notation choices

At this point, most people will think, “why not simply have one notation for variable-length initializer lists and another for function arguments?” That is:

```
vector<int> v21 { 1, 2 };    // initializer list: elements 1 and 2
vector<int> v22 ( 1, 2 );  // argument list: 1 element initialized to 2
```

In other words, do we really need a uniform initialization syntax that includes variable-length initializer lists? First let us consider this question in the abstract; that is, independently of compatibility concerns:

- Do we really want to distinguish syntactically between initializing elements with values and initializing elements by passing values to a constructor? We think not. The **S1** and **S2** example shows the feebleness of that distinction. Often, a constructor simply checks the values given to it (or puts them on a “normal form”) before assigning them to members. That’s not a fundamental logical distinction and in C++ we don’t usually syntactically distinguish between user-defined and built-in operations. Fundamentally, it should be possible for the user to take the point of view that “I don’t really care exactly how the object is initialized”. Conversely, it should be possible for the write of a class to take control over initialization, however expressed (e.g. for checking).
- Any syntactic distinction that doesn’t reflect a semantic distinction becomes a problem in the context of generic programming. For example, a template cannot without serious workarounds distinguish between argument types that require the one syntax for initialization from types that need the other. Uniformity of syntax is an important ideal here.
- The syntax above appears to distinguish between elements to be placed in the initialized object and arguments to a constructor. However,
 - sometimes, the arguments to a constructor are exactly values used to initialize members and
 - sometimes, the members initialized by an initializer list are just pointers to the “real elements” of the class – stored elsewhere and accessed indirectly through the members.

Thus, the syntax doesn’t necessarily reflect anything fundamental.

- Some homogeneous initializer lists (of values) are fixed-length, such as the list of coordinates for a point. Conversely, the list of arguments to a function can be variable length (we tend to simulate that with default arguments or overloading). Thus, the association of fixed-length with (...) and the association of variable-length with {...} is largely bogus.
- Some initializer lists are homogeneous, but many are not. A **pair** is a good example; so are most traditional initializer lists for **structs**. Some argument lists

are heterogeneous, but many are not. Thus the association of {...} with homogeneity and the association of (...) with heterogeneity is largely bogus. We conclude that from a fundamental point of view, we would prefer a uniform initialization syntax that could be used to provide values for either a sequence constructor or other constructors. In particular, a uniform syntax will help generic programming and support the C++ design aim of equal support for user-defined and built-in types.

Now consider compatibility. C introduced initializer lists for arrays and **structs** long before C++ came onto the field. It also introduced separate syntax for initialization by non-aggregates (e.g., =7) and for providing arguments to functions (e.g., (1,2)). C++ adopted all that and added to possibility of using the function call method of specifying arguments to object initialization. The result is a mess. For C, the mess can be excused because uniformity of notation wasn't a C design goal, but for C++ the non-uniformity has become a serious problem (educationally and in writing more generic/reusable code). For compatibility, we must

- accept both {...} and (...) initialization in the language.
- accept {...} lists to be variable-length and (sometimes) heterogeneous.
- accept (...) lists to (sometimes) be variable length (think “**printf**”) and heterogeneous
- not make any significant changes to the (...) semantics (note the subsidiary/additional proposal to ban narrowing in initializations; Section 7).
- note that (...) often appear in contexts where it is not an initializer (e.g. as a list of (argument) types, a comma expression or (more generally) as a sub-expression.

Does this have implications on whether we should have a separate syntax for initializer lists and argument lists? Not really, but as ever compatibility constrains solutions and eliminates the possibility of simplifying by removing existing syntax.

We conclude that existing uses of {...} and (...) don't seem to significantly bias or guide the choice of notation for our **vector** example. In particular, both {...} and (...) are already used for both fixed-length and variable-length lists and for both homogeneous and heterogeneous lists. The wider use of (...) compared to {...} pushes us towards basing the unified initialization syntax on {...} rather than (...), where a “green field” design might have had a free choice.

This example shows how a uniform notation helps:

```
Map<string,int> m = { {"ardwark",91}, {"bison", 43} };
```

11.3 Disambiguation

Back to our example: Assume that we use a uniform syntax for initialization:

```
vector<int> v2 { 1, 2 }; // one or two elements?
```

For this example, we have three choices:

- Sequence constructors take priority: It's a variable-length initializer list that happens to have two elements
- "Ordinary constructors" take priority: It's an argument list that happens to match one of vector's constructors
- All constructors are equal: It's an ambiguity error.

In §11.1, we saw that if we take the third alternative, we need two ways of disambiguating. For example:

```
vector<int> v21 = initializer_list<int> { 1, 2 }; // two elements
vector<int> v22 ( 1, 2 ); // one element
```

If we pick either of the first two alternatives, we need only one way to disambiguate. First alternative (prefer sequence constructors):

```
vector<int> v2 { 1, 2 }; // two elements
vector<int> v22 ( 1, 2 ); // one element
```

Alternative two (prefer "ordinary constructors"):

```
vector<int> v2 { 1, 2 }; // one element
vector<int> v21 = initializer_list<int> { 1, 2 }; // two elements
```

The second alternative seems backwards as well as verbose, but it also has more fundamental problems:

```
vector<int> v2 { }; // no elements (default constructor)
vector<int> v2 { 3 }; // three elements (initialized to 0)
vector<int> v2 { 1, 2 }; // one element (initialized to 2)
vector<int> v2 { 1, 2, 3 }; // three elements with values 1, 2, 3
```

This is awful! Also:

```
vector<int*> vp1 { &i1 }; // ok (one element)
vector<int*> vp1 { &i1, &i2 }; // initialize using the sequence [&i1,&i2)
vector<int*> vp1 { &i1 , &i2, &i3 }; // ok (three elements)
```

Basically, giving priority to "ordinary constructors" implies a need for disambiguation that unpredictably appears depending on the number and types of elements. In other words, we can't use initializer lists uniformly even for a particular type. We will not pursue this alternative further. So we are left with two alternatives:

- Sequence constructors take priority: It's a variable-length initializer list that happens to have two elements.
- All constructors are equal: It's an ambiguity error.

The "ambiguity" resolution shares the non-uniformity problems we just mentioned with the "ordinary constructors take priority" approach. First we have:

```

vector<int> v2 { }; // ambiguous: default constructor or empty initializer?
vector<int> v2 { 3 }; // ambiguous
vector<int> v2 { 1, 2 }; // ambiguous
vector<int> v2 { 1, 2, 3 }; // three elements with values 1, 2, 3

vector<int*> vp1 { &i1 }; // ok (one element)
vector<int*> vp1 { &i1, &i2 }; // ambiguous
vector<int*> vp1 { &i1, &i2, &i3 }; // ok (three elements)

```

We can resolve the ambiguities either way. First, let's resolve to use ordinary constructors:

```

vector<int> v2; // no elements (default constructor)
vector<int> v2(3); // three elements (initialized to 0)
vector<int> v2(1, 2); // one element (initialized to 2)
vector<int> v2 { 1, 2, 3 }; // three elements with values 1, 2, 3

vector<int*> vp1 { &i1 }; // ok (one element)
vector<int*> vp1 (&i1, &i2); // initialize using the sequence [&i1,&i2]
vector<int*> vp1 { &i1, &i2, &i3 }; // ok (three elements)

```

So much for uniform syntax! Alternatively, we can resolve the examples to use the sequence constructor:

```

vector<int> v2 = initializer_list<int>{ };
vector<int> v2 = initializer_list<int>{ 3 };
vector<int> v2 = initializer_list<int>{ 1, 2 };
vector<int> v2 { 1, 2, 3 };

vector<int*> vp1 { &i1 };
vector<int*> vp1 = initializer_list<int*>{ &i1, &i2 };
vector<int*> vp1 { &i1, &i2, &i3 };

```

It's verbose, but at least we didn't have to modify the initializer lists themselves.

We strongly dislike that verbose disambiguation because initializer lists of built-in types are going to be common. Worse, the `initializer_list<int>` doesn't just disambiguate, it also casts. Consider:

```

void f(const vector<double>&);
// ...
f(initializer_list<int> { x, y }); // resolve to two elements

```

If `y` was a floating point number, we *explicitly* introduced an error. Disambiguation by prefixing an initializer list with its desired type is "natural" and general, but it is not

minimal or ideal. What would be minimal and ideal? Something that simply said “this initializer lists may not be used as arguments for an ordinary constructor”. For example, we could make a = significant in declarations

```
vector<int> v21 { 1, 2 };    // potentially ambiguous
```

or have to disambiguate:

```
vector<int> v21 = { 1, 2 }; // definitively a set of elements
```

That is, we could have `{...}` mean “initializer list (either arguments of values)” and `= {...}` mean “initializer list; do not use as constructor arguments”. This is a complete solution: It can be used in every initialization context. However, we consider it “too cute”. Note that is purely an aesthetic judgment. However, the minute we consider using initializer lists within general expressions that `= {...}` notation starts to look seriously weird:

```
void f(const vector<double>&);
// ...
f({x,y});           // potentially ambiguous
f(={ x ,y });      // ok

v1 = { 1, 2};           // potentially ambiguous
v2 = vector<int>{1,2}; // also potentially ambiguous!
v3 = initializer_list<int>{ 1, 2 }; // ok
v4 = ={1,2};         // ok
```

For the assignments, we are looking at the initializer list as the argument to `vector<int>::operator=()`.

11.4 Conclusion

So, how do we decide between the remaining two alternatives (“ambiguity” and “sequence constructors take priority over ordinary constructors)? Our proposal gives sequence constructors priority because

- Looking for ambiguities among all the constructors leads to too many “false positives”; that is, clashes between apparently unrelated constructors. See examples below.
- Disambiguation is itself error-prone (as well as verbose). See examples in §11.3.
- Using exactly the same syntax for every number of elements of a homogeneous list is important – disambiguation should be done for ordinary constructors (that do not have a regular pattern of arguments). See examples in §11.3.

The simplest example of a false positive is the default constructor:

```
vector<int> v;
vector<int> v { };    // potentially ambiguous
```

```
void f(vector<int>&);  
// ...  
f({ });           // potentially ambiguous
```

It is possible to think of classes where initialization with no members is semantically distinct from default initialization, but we wouldn't complicate the language to provide better support for those cases than for the more common case where they are semantically the same.

Giving priority to sequence constructors breaks argument checking into more comprehensible chunks and gives better locality.

```
void f(const vector<double>&);  
// ...  
struct X { X(int); /* ... */ };  
void f(X);  
// ...  
f(1);           // call f(X); vector's constructor is explicit  
f({1});        // potentially ambiguous: X or vector?  
f({1,2});      // potentially ambiguous: 1 or 2 elements of vector
```

Here, giving priority to sequence constructors eliminates the interference from **X**. Picking **X** for **f(1)** is a variant of the problem with explicit shown in §3.3.