

# Rules of thumb for the design of C++0x

Bjarne Stroustrup  
Texas A&M University  
<http://www.research.att.com/~bs>

## Abstract

A good programming language is far more than a simple collection of features. My ideal is to provide a set of facilities that smoothly work together to support design and programming styles of a generality beyond my imagination. Here, I outline rules of thumb (guidelines, principles) that are being applied in the design of C++0x. For example, generality is preferred over specialization, novices as well as experts are supported, library extensions are preferred over language changes, compatibility with C++98 is emphasized, and evolution is preferred over radical breaks with the past. Since principles cannot be understood in isolation, I very briefly present a few of the proposals being considered in the ISO C++ standards committee.

## Introduction

I get an email suggesting an improvement to C++ [ISO, 1998] [ISO, 2003] about once or twice a day. Many suggestions are good in the sense that if one became part of the language or the standard library it would make life easier for a large number of programmers. The suggestions are of course not all unique, so my collection of suggestions contains “only” a hundred or so good ideas. The count depends critically on how you cluster related suggestions. You can see an incomplete list of suggested language features on <http://www.research.att.com/~bs/evol-issues.html> and a wish list for standard library suggestions on <http://lafstern.org/matt/wishlist.html>. My basic opinion is that the language wish list is far too long and that the libraries wish list is far too modest.

So, imagine you were the ISO C++ standards committee and could decree what the next C++ standard would look like: What would you add? What would you remove? What would you change? Seriously: Why don't you take a break from reading right now and make a list? If you are still happy with your list when you have completed this article then please add arguments for your ideas and email your suggestions to me.

As we are revising the standard, what are we trying to achieve? We can't expect perfect agreement on this point; after all, who are “we”? The C++ community is huge (well over 3 million programmers according to IDC last year) and incredibly diverse. We use just about every computer and operation system there is, work in just about every application area there is, and in every country on earth. This implies a diversity of needs and wishes, and to me it implies that the further evolution must be guided by a set of general principles to avoid accidentally damaging a sub-community in the quest for improvements for some other sub-community.

The sections below illustrate the main principles that we use to guide the development of C++0x. Clearly they owe a lot to the “Rules of thumb” described in [Stroustrup, 1994]. That is no accident: My aim is to reinforce C++’s proven strengths rather than try to use the new standard to encourage dramatic innovation. To balance: It is not my aim to leave C++ standing still – as a living language, C++ must grow and adapt. It must become a more effective tool for solving current and future problems facing the community. I switched to first person singular here to emphasize that I’m just one individual member of a committee that contains many members. A huge range of the opinions from the C++ community are represented and what I say here appear to be backed by a strong consensus. However, the principles and their application to real problems are open to much discussion, interpretation, and experimentation. That is how it should be.

So, why don’t we just accept all reasonable suggestions to make everybody happy? There are too many suggestions to analyze and adequately specify. If we accepted all suggestions there would be 6 or 7 ways of doing anything in C++ rather than just 2 or 3. There would be far too many features to teach and most programmers would settle into a small zone of comfort eschewing most new features. It would be extremely difficult to maintain the zero-overhead principle, partly because the implementation effort would be immense, partly because the implementation of different features would interact, and partly because many suggestions don’t consider performance and rely on elaborate runtime mechanisms. This, of course, presupposes that the implementer community would accept such a standard. I wouldn’t bet on that. Please remember these concerns when you want to argue for adding “just my two new features.”

## Design and evolution

C++0x will be almost 100% compatible with existing Standard C++. Your existing code will with a very high probability be conforming C++0x if it was C++98. The best guarantee for that is that the members of the committee are collectively responsible for much more old code than you are. However, we care about the future, see the evolution of C++ as a necessity, and want future code to be easier to write, more elegant, easier to maintain, and possibly even better performing than is currently feasible.

We aim for compatibility but realize that there may be cases where a large advantage is worth paying for by small incompatibility. An obvious example is a new keyword. We try to avoid incompatibilities, but the cost of an absolute “no incompatibilities” rule would be “no change” or obscure syntax for all new features. I personally strongly dislike workarounds, such as a `__XXX` keyword supplemented by a macro

```
#define XXX __XXX
```

The extra complexity of the workaround compared to a keyword **XXX** serves users who do not want to change code containing an **XXX** at the inconvenience of the rest of the community, added complexity to learners, and the cheap amusement of people who don’t wish C++ well. There are of course cases where someone uses **XXX** and cannot modify

the code, but then there are alternatives: Don't upgrade to a new compiler or use a backwards compatibility switch.

The idea of language evolution is often contrasted to that of language design. That's a false dichotomy. Language evolution is design. It differs from supposedly "from scratch" design by being more constrained by respect for existing code and by benefiting from more direct application of experience (feedback). In return for the difficulties imposed by compatibility, we gain the advantages of relative ease of adoption by a huge community and the smooth interoperation of old and new features. By choosing evolution we also avoid the problem of accidentally eliminating useful programming techniques that just happened to be outside our experience or focus.

Language design can become obsessed with minor notational issues. However, a language feature is most significant if it makes a new style of programming effective to its users. In the context of C++, that means that the aim of language and standard library changes must be to bring new programming and design techniques into mainstream (industrial and educational) use. Only features that change the way we think are really worth providing.

To sum up: The aim for C++0x is evolution constrained by a strong need for compatibility. The aim of that evolution is to provide major real-world improvements.

## **Generality and specialization**

Requests for specialized facilities and minor notational improvements are very common. When provided, they rarely fail to win applause. After all, if a feature is a direct solution to a problem and doesn't significantly interact with other facilities, it is easy to explain, often easy to implement, and typically have a logically minimal expression leading to very concise pieces of code. People comparing languages using lists love such features. The snag is that the problems we face are essentially infinite, so we need an infinity of such specialized features. Examples are Pascal procedure parameters and C# delegates. The alternative traditionally offered by C++ (and before that by K&R C) is to provide very general features from which a good programmer can construct solutions to a wide variety of problems. Examples are pointers and classes.

C++'s emphasis on general features (notably classes) has been its main strength and often its lack of specialized features (such as "properties" and threads) have been seen as its main weakness. Obviously, both observations can be simultaneously true. Nevertheless, we must keep the focus on general features aimed at efficient abstraction; the huge diversity of the C++ community requires that. Features specifically tailored for, say Windows application building or embedded systems programming, would become a huge liability if they did only what they were specialized for. C++0x will not be a "Windows language" or a "Web language" or an "Embedded systems language"; it will be a general-purpose language that supports those applications areas – and more – using a common set of facilities.

One important reason to favor general mechanisms over specialized solutions to specific current problems is that the general mechanisms are likely to help with yet-un-thought-of problems: they are insurance against nasty surprises in the future. I don't want a language that can express only what I specifically designed it for.

The obvious areas where C++ could be improved to for greater generality is through better support for generic programming and more flexible initialization/construction mechanisms (see example section). It is also obvious that some support for concurrency is needed as many forms of concurrent, parallel, and distributed programming are becoming common. The diversity of such approaches and techniques implies that no single mechanism can adequately cover all applications. Thus, the obvious approach is to provide very simple language mechanisms supported by libraries (built generic and object-oriented techniques).

To sum up: The aim for C++0x is to supply general language mechanisms that can be used freely in combination and to deliver more specialized features as standard library facilities built from language features available to all.

## Experts and novices

C++ has drifted towards becoming an “expert friendly” language. In a gathering (in person or on the web) of experts, it is hard to build a consensus (or even interest) for something that “just” helps novices. The general opinion (in such a gathering) is typically that the best we can do for novices is helping them to become experts. But it takes time to become an expert and most people need to be reasonably productive during the time it takes. More interesting, many C++ novices have no wish to or need to become experts in C++. If you are a physicist needing to do a few calculations a week, an expert in some business processes involving software, or a student learning to program, you want to learn only as many language facilities as you need to get your job done; you don't want to become a language expert – you want to be (or become) an expert in your own field and know enough of some programming language to get your work done. When supported by suitable libraries, C++ can be used like that – it is widely used like that. However, there are traps, pitfalls, and educational approaches that make such “occasional use” of C++ unnecessarily difficult. With a modest effort, C++0x can do much better in this area.

Consider a couple of trivial examples. Have you ever written something like this?

```
vector<vector<double>> v;
```

or this?

```
int i = extract_int(s); // s is a string, e.g. “12.37”
```

or this?

```
vector<int>::iterator p =
```

```
find(tbl.begin(),tbl.end(),x); // tbl is a const vector<int>
```

If not, I suspect you have either been using a different style of C++ with its own problems, or written very little C++. Admitting >> as the end of two template argument lists will solve the first problem. Providing a standard library operation of parse a numeric value in a string would save the novice the (significant) bother of discovering the (obvious to experts) definition of a function like **extract\_int()**. Allowing the type of **p** to be deduced from its initializer would solve the third problem:

```
auto p = find(tbl.begin(),tbl.end(),x); // tbl is a const vector<int>
                                     // p becomes a vector<int>::const_iterator
```

The >> and **auto** solutions have been approved in principle for C++0x. Attacking the problem of supporting “novices of all backgrounds” requires work on both the language and the standard library. Concerns for education will be central for that; see for example [Stroustrup, 1999].

Overloading based on concepts (see example section), would allow a further simplification

```
auto p = find(tbl,x); // tbl is a const vector<int>
                       // p becomes a vector<int>::const_iterator
```

Features added for novices (of all levels of expertise) should not be isolated from the rest of the language, creating some sort of “novices’ ghetto”. Like all features, they should scale to use in major systems, interact smoothly with other features, and provide a path for learning the complete language and standard library.

To sum up: C++0x must support novices of all backgrounds much better than does current C++ – both through less error-prone language features and through more supportive libraries.

## Type safety

The key to elegant and efficient code is type safety using a flexible and extensible type system. C++0x will not be able to make C++ completely type safe – that would require banning arrays, uninitialized pointers, unions, narrowing conversions, C-style linking, and much more. Doing so would also cause problems with hardware access as needed in many embedded systems applications. So what can we do? Lack of type safety is the root cause of many problems with correctness and with performance. For example:

```
void get_input(char* p)
{
    char ch;
    while (cin.get(ch) && !iswhite(ch)) *p++ = ch;
    *p = 0;
}
```

That should send chills down your spine; it really is scary. Similarly, if you are concerned with performance you should be most unhappy with this style of generic linked list:

```

struct Link {
    Link* link;
    void* data;
};

void my_clear(Link* p, int sz)    // clear data of size sz
{
    for (Link* q = p; q!=0; q = q->link) memset(q->data,0,sz);
}

```

What can we do? Basically, all we can do is to provide alternatives to the unsafe practices and rely on tools (a la lint, but with better knowledge of the C++ type system) to detect unsafe (“traditional”) use. For example:

```

string s;
cin >> s;

```

This alternative completely bypasses the problems of that low-level and sloppy **get\_input()** while being easier to use and (at least potentially) as fast. Similarly:

```

template<class In> void my_stl_clear(In first, In last)
{
    while (first!=last) *first++ = 0;
}

```

This is a complete alternative to **my\_clear()**. What’s wrong with **my\_clear()**? It is a bit longer than **my\_stl\_clear()**, but that’s not the point when we consider performance. The problem is lack of type information. **Link::data** is a **void\***, so it may point to any memory location. This means that the compiler must assume that the write to **q->data** (in **memset()**) may change the contents of **\*(q->link)** or even **q->link** itself. That basically disables the optimizer. Had **link** and **data** had different static types, the compiler would have assumed that they weren’t aliases. So, given a reasonable optimizer, **my\_stl\_clear()** runs several times faster than **my\_clear()**. In addition, the use of **memset()** is an example of the workarounds that we must use to access data through a **void\***: We often end up using a function or even an indirectly called function. That too can be very expensive [Stroustrup, 1999].

The most direct way of addressing the problems caused by lack of type safety is to provide a range checked standard library based on statically typed containers and base teaching of C++ on that. This will not eliminate type errors – there will always be programmers who will decide to use the “hand-coded primitive” style, programmers who must (for a variety of reasons) turn off checking when they ship, and there will always be

many, many, areas that are not covered by the standard library. The last problem will be addressed by the standard library setting a standard for other libraries to meet.

A systematic use of range checking for a version of the standard library will imply the systematic use of exceptions, but that's ok for code that doesn't involve hard real time. By now we know how to deal with exceptions; for example, see Appendix E of [Stroustrup, 2000]

To sum up: C++0x will not be able to close all the loopholes in the C++ type system, but it will not introduce new holes and it will provide ways of avoiding inherently unsafe facilities – primarily through the standard library providing (compile-time or run-time) type safe alternatives.

## Performance and machine model

Performance based on a simple mapping of language features to hardware has been crucial for C++'s success (as it was for C). With the notable exception of exceptions, C++ requires minimal run-time support. Support for RTTI (**dynamic\_cast** and **typeid**) and free store (**new** and **delete**) needs to be included only if you directly or indirectly use those facilities. The possibility of eliminating potentially expensive features is important in many application areas, notably embedded systems programming. Exception handling is getting quite efficient these days, but unfortunately it is not predictable enough for hard-real time programs. Where necessary, exception handling can be disabled through compiler switches.

Nothing in C++0x will change this state of affairs. C++ is and C++0x will remain directly applicable to the most performance and resource critical applications; see [ISO 2004]. If C++0x adds facilities that require run-time support, those features will be designed to ensure that the additional support is required only by code that actually use them. That is, the zero-overhead principle: “what you don't use, you don't pay for” and “what you use can be implemented without overhead compared with hand coding” is still the bedrock of C++.

The C++ model of hardware is simple: basic types map directly to entities recognized by hardware, such as bytes, words, and registers. Sequences of objects map directly into the hardware's. Operations directly reflect machine instructions and just about everything can be done without spurious allocation, indirection, or circumlocution. The challenge will be to fit concurrency into this picture.

To sum up: C++0x will follow the zero-overhead principle and preserve a machine model that allow direct and efficient use of essentially all hardware.

## Language and libraries

A language cannot support everything, but conceivably a large set of libraries could. Not every library should be standard. The C++ community has ample room for both a large standard library and a libraries industry (both commercial and open source). The committee's libraries technical report [ISO 2005] raises the bar by providing facilities such as regular expression matching and hash tables. The Boost community shows other examples of what might become standard. However, even a reasonable portion of that is not enough for my taste. I'd like to see the C++ standard library become a much more extensive platform for systems programming, including concurrency and support for geographical distribution. "Distribution" is quite possibly beyond the committee's capacity – we are after all a group of volunteers with "day jobs" and without funds to sustain development. However, we can and must dream.

"Libraries" is an area where the committee can afford to be aggressive and opportunistic. Here the community can provide significant (and essential) support. A library differs from a language feature in the crucial respect that it is not exclusive: If you don't like a library, you can use another. Also, a library can be developed, tested, and even deployed without the involvement of compiler vendors. C++'s mechanisms for expressing libraries – even in performance critical and resource constrained areas – should not be underestimated. For current examples, see Boost ([www.boost.org](http://www.boost.org)). C++0x should be even more effective for library building.

The most commonly asked for new feature for C++ is a standard GUI. The technical, economical, and political odds against that happening are immense. However, so were the odds against a good container and algorithms library for C++98. Then the STL appeared. I hope for a miracle and dream of a simple and elegant standardized interface to a variety of commercial and open-source GUI facilities. This is not a reasonable dream, but (as many has pointed out) the world is not changed by reasonable people.

To sum up: library extensions will be preferred to language extensions. Whereas we'll be cautious and skeptical about language extensions, we'll be aggressive and opportunistic when it comes to new libraries – especially for libraries that extend the range of portable support for systems programming.

## The standard and the real world

The ISO standard is important to the C++ community, but it is obviously just a small part of what influences C++ program development. I don't mention other languages here, but coexistence with a huge number of systems and libraries is part of C++'s nature. That influences our considerations about language features and libraries. Furthermore, most suggestions for language and library extensions start out as more or less disguised requests to "do what language X does". I guess that every good feature in any language has at some point in time been suggested for C++. As professionals, most committee members are quite familiar with a wide range of languages and our experiences do guide our design decisions. For the particularly tricky issue of C++'s relationship to ISO C, see my series of articles [Stroustrup, 2002].



The development of C++0x must be timely and to gain acceptance from vendors, new language features must be easier to implement than the most difficult C++98 features [ISO, 1998].

As programmers and especially as programmers of many systems, we dislike dialects. However, there always has been and there always will be dialects. Bindings to operating systems, data bases, middleware, etc. are a major source of non-standard features. Support for pre-standard facilities is another. A third source is support for “minority” needs; that is, facilities that are essential to a small community, but irrelevant to a large majority of the C++ community. My suggestion is for programmers to avoid non-standard features whenever that’s possible. When that is not possible – as is the case in parts of most major systems – I suggest localizing the use of the non-standard features and accessing them through interfaces written in ISO Standard C++. One of the aims of improving the general abstraction mechanisms of C++ is exactly to make such encapsulation of non-standard code easier. The alternative is vendor lock-in.

To sum up: We can’t do everything. What we do will be guided by practical software development concerns as well as compatibility and language design constraints.

## An example: Support for generic programming

Consider:

```
template<class T> class vector {
    // ...
    void push_back(const T&) { /* ... */ }
    // ...
};

vector<double> v;
v.push_back(1.2);
v.push_back(2.3);
v.push_back(3.4);
```

This is basically an example of the use of templates and the STL. From extensive real-world use and many experiments, we can consider the language facilities and programming techniques involved successful and amazingly flexible; see the STL and Boost for examples. In particular, the use of templates has become standard where performance is essential. So, how can we do better for C++0x? By “better” I mean extending what can be elegantly expressed without loss of performance compared to C++98 and which remedies problems with current use. In particular, how can that **vector** example be improved? The repeated **push\_back()** is verbose and ugly, the lack of specification of the element type is a weakness leading to spectacularly bad error messages, and ideally I don’t want to expose the implementation of **push\_back()**. A better version of that code would read:

```

template<Value_type T> class vector {
    // ...
    void push_back(const T&); // just the declaration
    // ...
};

vector<double> v;
v.push_back(1.2);
v.push_back(2.3);
v.push_back(3.4);

```

The **Value\_type** used to specify the element type **T** is a “concept”; it specifies what the vector assumes about **T**. Given that we can verify that **double** is a **Value\_type**, we can type check the definition of **v** without seeing the definition of **push\_back()**. We would need the definition of **push\_back()** if we wanted to inline, but that now becomes an implementation detail. Concepts could give us separate checking of translation units without imposing type hierarchies on template arguments or performance penalties on template users.

How could we eliminate the repeated call of **push\_back()**? We could allow a vector to take an initializer list as its argument. That would require the definition of a constructor taking such an initializer list. For example:

```

template<Value_type T> class vector {
    // ...
    vector(const T*, const T*); // sequence constructor
    // ...
};

vector<double> v = { 1.2, 2.3, 3.4 };

```

The detailed design of concepts and generalized use of initializer lists are still under intense discussion. The point here is not the details, but that I consider such facilities within the realistic grasp of C++0x.

## Will this happen?

I think C++0x will clearly reflect the rules of thumb outlined here. Lack of resources (time, people, etc.) will limit what we can do and obviously (given such a large and complex task) we’ll make a few mistakes. Also, some “random extensions” will slip through the net and become “odd and isolated” features in the language (much as enumerations are in C and C++). However, there is reason to expect that C++0x will be a significant improvement over current C++ for essentially all of its users and for many more users still to come.

## References

- [ISO, 1998] “Standard for the C++ Programming Language”. ISO/IEC 14882.
- [ISO, 2003] “The C++ Standard” (ISO/IEC 14882:2002). Wiley 2003. ISBN 0-470-84674-7.
- [ISO, 2004] “Technical Report on C++ Performance”. ISO/IEC PDTR 18015. Link on <http://www.research.att.com/~bs/C++.html>.
- [ISO, 2005] “Technical Report on C++ Standard Library Extensions”. ISO/IEC PDTR 19768.
- [Stroustrup, 1994] B. Stroustrup: “The Design and Evolution of C++”. Addison-Wesley. 1994. ISBN 0-201-54330-3.
- [Stroustrup, 1999] B. Stroustrup: “Learning Standard C++ as a New Language”. C/C++ Users Journal. May 1999. Link on <http://www.research.att.com/~bs/papers.html>.
- [Stroustrup, 2000] B. Stroustrup: “The C++ Programming Language (Special Edition)”. Addison Wesley. Reading Mass. USA. February 2000. ISBN 0-201-70073-5.
- [Stroustrup, 2002] B. Stroustrup: “C and C++: Siblings”, “C and C++: A Case for Compatibility”, “C and C++: Case Studies in Compatibility”. The C/C++ Users Journal. July, August, and September 2002. Link on <http://www.research.att.com/~bs/papers.html>.