


How Non-Member Functions Improve Encapsulation

When it comes to encapsulation, sometimes less is more.



I'll start with the punchline: If you're writing a function that can be implemented as either a member or as a non-friend non-member, you should prefer to implement it as a non-member function. That decision *increases* class encapsulation. When you think encapsulation, you should think non-member functions.

Surprised? Read on.

Background

When I wrote the first edition of *Effective C++* in 1991 [1], I examined the problem of determining where to declare a function that was related to a class. Given a class *C* and a function *f* related to *C*, I developed the following algorithm:

```
if (f needs to be virtual)
    make f a member function of C;
else if (f is operator>> or
        operator<<)
{
    make f a non-member function;
    if (f needs access to non-public
        members of C)
        make f a friend of C;
}
else if (f needs type conversions
        on its left-most argument)
{
    make f a non-member function;
    if (f needs access to
        non-public members of C)
        make f a friend of C;
}
else
    make f a member function of C;
```

This algorithm served me well through the years, and when I revised *Effective C++* for its second edition in

1997 [2], I made no changes to this part of the book.

In 1998, however, I gave a presentation at Actel, where Arun Kundu observed that my algorithm dictated that functions should be member functions even when they could be implemented as non-members that used only *C*'s public interface. Is that really what I meant, he asked me? In other words, if *f* could be implemented as a member function or a non-friend non-member function, did I really advocate making it a member function? I thought about it for a moment, and I decided that that was not what I meant. I therefore modified the algorithm to look like this [3]:

```
if (f needs to be virtual)
    make f a member function of C;
else if (f is operator>> or
        operator<<)
{
    make f a non-member function;
    if (f needs access to non-public
        members of C)
        make f a friend of C;
}
else if (f needs type conversions
        on its left-most argument)
{
    make f a non-member function;
    if (f needs access to non-public
        members of C)
        make f a friend of C;
}
else if (f can be implemented via C's
        public interface)
    make f a non-member function;
else
    make f a member function of C;
```

Since then, I've been battling programmers who've taken to heart the lesson that being object-oriented means putting functions inside the classes containing the data on which the functions operate. After all, they tell me, that's what encapsulation is all about. They are mistaken.

Encapsulation

Encapsulation is a *means*, not an end. There's nothing inherently desirable about encapsulation. Encapsulation is useful only because it yields other things in our software that we care about. In particular, it yields flexibility and robustness. Consider this struct, whose implementation I think we'll all agree is unencapsulated:

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


The weakness of this struct is that it's not flexible in the face of change. Once clients started using this struct, it would, practically speaking, be very hard to change it; too much client code would be broken. If we later decided we wanted to compute *x* and *y* instead of storing those values, we'd probably be out of luck. We'd be similarly thwarted if we decided a superior design would be to look *x* and *y* up in a database. This is the real problem with poor encapsulation: it precludes future implementation changes. Unencapsulated software is inflexible, and as a result, it's not very robust. When the world changes, the software is unable to gracefully change with it. (Remember that we're talking here about what is *practical*, not what is possible. It's clearly possible to change struct *Point*, but if enough code is dependent on it in its current form, it's not practical.)

Now consider a class with an interface that offers clients capabilities similar to

Scott Meyers is a recognized authority on C++; he provides consulting services to clients worldwide. He is the author of *Effective C++*, *Second Edition* (Addison-Wesley, 1998), *More Effective C++* (Addison-Wesley, 1996), and *Effective C++ CD* (Addison-Wesley, 1999). Scott received his Ph.D. in Computer Science from Brown University in 1993.

Power Tools

for measurement



National Instruments delivers the measurement tools to convert your Visual Basic or Visual C/C++ compiler into a high-performance Measurement Ready™ development environment.

Enhance your programming capabilities with:

- Data acquisition (DAQ)
- Analysis
- Graphing

Microsoft Certified
Solution Provider

Download your **FREE**
evaluation tools today



**NATIONAL
INSTRUMENTS™**

www.ni.com/tools

(800) 452-6202

Fax: (512) 683-9300 • info@ni.com

© Copyright 1999 National Instruments Corporation. All rights reserved. Product and company names listed are trademarks or trade names of their respective companies.



those afforded by the struct above, but with an encapsulated implementation:

```
class Point {
public:
    int getXValue() const;
    int getYValue() const;
    void setXValue(int newXValue);
    void setYValue(int newYValue);

private:
    ...           // whatever...
};
```

This interface supports the implementation used by the struct (storing *x* and *y* as ints), but it also affords alternative implementations, such as those based on computation or database lookup. This is a more flexible design, and the flexibility makes the resulting software more robust. If the class's implementation is found lacking, it can be changed without requiring changes to client code. Assuming the declarations of the public member functions remain unchanged, client source code is unaffected. (If a suitable implementation has been adopted [4], clients need not even recompile.)

Encapsulated software is more flexible than unencapsulated software, and, all other things being equal, that flexibility makes it the superior design choice.

Degrees of Encapsulation

The class above doesn't fully encapsulate its implementation. If the implementation changes, there's still code that might break. In particular, the member functions of the class might break. In all likelihood, they are dependent on the particulars of the data members of the class. Still, it seems clear that the class is more encapsulated than the struct, and we'd like to have a way to state this more formally.

It's easily done. The reason the class is more encapsulated than the struct is that more code might be broken if the (public) data members in the struct change than if the (private) data members of the class change. This leads to a reasonable approach to evaluating the relative encapsulations of two implementations: if changing one might lead to more broken code than would the corresponding change to the other, the former is less encapsulated than the latter. This definition is consistent with our intuition that if making a change is likely to break a lot of code, we're less likely to make that change than we would be to make a different change that affected less code. There is a direct relationship between

encapsulation (how much code might be broken if something changes) and practical flexibility (the likelihood that we'll make a particular change).

An easy way to measure how much code might be broken is to count the functions that might be affected. That is, if changing one implementation leads to more potentially broken functions than does changing another implementation, the first implementation is less encapsulated than the second. If we apply this reasoning to the struct above, we see that changing its data members may break an unknowably large number of functions — every function that uses the struct. In general, we can't count how many functions this is, because there's no way to locate all the code that uses a particular struct. This is especially true for library code. However, the number of functions that might be broken if the class's data members change is easy to determine: it's all the functions that have access to the private part of the class. That's just four functions (assuming none are declared in the private part of the class), and we know that because they're all conveniently listed in the class definition. Since they're the only functions that have access to the private parts of the class, they're the only functions that can be affected if those parts change.

Encapsulation and Non-Member Functions

We've now seen that a reasonable way to gauge the amount of encapsulation in a class is to count the number of functions that might be broken if the class's implementation changes. That being the case, it becomes clear that a class with *n* member functions is more encapsulated than a class with *n*+1 member functions. And that observation is what justifies my argument for preferring non-member non-friend functions to member functions: if a function *f* could be implemented as a member function or as a non-friend non-member function, making it a member would decrease encapsulation, while making it a non-member wouldn't. Since functionality is not at issue here (the functionality of *f* is available to class clients regardless of where *f* is located), we naturally prefer the more encapsulated design.

It's important that we're trying to choose between member functions and *non-friend non-member* functions. Just like member functions, friend functions may be broken when a class's implementation changes, so

the choice between member functions and friend functions is properly made on behavioral grounds. Furthermore, we now see that the common claim that “friend functions violate encapsulation” is not quite true. Friends don’t violate encapsulation, they just decrease it — in exactly the same manner as member functions.

This analysis applies to any kind of member functions, including static ones. Adding a static member function to a class when its functionality could be implemented as a non-friend non-member decreases encapsulation by exactly the same amount as does adding a non-static member function. One implication of this is that it’s generally a bad idea to move a free function into a class as a static member just to show that it’s related to the class. For example, if I have an abstract base class for Widgets and then use a factory function [4,5,6] to make it possible for clients to create Widgets, the following is a common, but inferior way to organize things:

```
// a design less encapsulated than it could be
class Widget {
    ...           // all the Widget stuff; may be
                // public, private, or protected

public:

    // could also be a non-friend non-member
    static Widget* make(/* params */);
};
```

A better design is to move make out of Widget, thus increasing the overall encapsulation of the system. To show that Widget and make are related, the proper tool is a namespace:

```
// a more encapsulated design
namespace WidgetStuff {
    class Widget { ... };
    Widget* make( /* params */ );
};
```

Alas, there is a weakness to this design when templates enter the picture. For details, see the accompanying sidebar.

Syntax Issues

If you’re like many people with whom I’ve discussed this issue, you’re likely to have reservations about the syntactic implications of my advice that non-friend non-member functions should be preferred to member functions, even if you buy my argument about encapsulation. For example, suppose a class Wombat supports the functionality of both eating and sleeping. Further suppose that the eating functionality must be implemented as a member function, but the sleeping functionality could be implemented as a member or as a non-friend non-member function. If you follow my advice from above, you’d declare things like this:

```
class Wombat {
public:
    void eat(double tonsToEat);
    ...
};
```

```
void sleep(Wombat& w, double hoursToSnooze);
```

That would lead to a syntactic inconsistency for class clients, because for a Wombat w, they’d write

```
w.eat(.564);
```

to make it eat, but they would write

```
sleep(w, 2.57);
```

to make it sleep. Using only member functions, things would look much neater:

```
class Wombat {
public:
    void eat(double tonsToEat);
    void sleep(double hoursToSnooze);
    ...
};
```

```
w.eat(.564);
w.sleep(2.57);
```

Ah, the uniformity of it all! But this uniformity is misleading, because there are more functions in the world than are dreamt of by your philosophy.

To put it bluntly, non-member functions happen. Let us continue with the Wombat example. Suppose you write software to model these fetching creatures, and imagine that one of the things you frequently need your Wombats to do is sleep for precisely half an hour. Clearly, you could litter your code with calls to w.sleep(.5), but

New Release!
More Features!

SftTree 4.0

Tree Control for Windows

From simple listbox with bitmaps to multi-line, multi-column, hierarchical data display, SftTree delivers!

#	Printer / Files	Description	Ready	
3	Xerox DocuTech 135 (A)	1st Floor, Printer Ro...	<input checked="" type="checkbox"/> Ready	
5	HP LaserJet III PS	2nd Floor, Room 228	<input checked="" type="checkbox"/> Ready	
6	TREECTL.CPP	Sample Program	<input checked="" type="checkbox"/> Printing	
7	TREEX.RC	Resource Script	<input checked="" type="checkbox"/> Ready	
8	BONUS.DOC	Additional document...	<input checked="" type="checkbox"/> Ready	
9	HP LaserJet 4	2nd Floor, Room 206	<input type="checkbox"/> Printer Down	
10	URGENT.DOC	Most urgent output	<input type="checkbox"/> Not ready	

- multiple columns
- multiple text lines per item
- drag & drop with auto-scroll
- single/multiple selection
- Tool Tips
- cell editing

- columns with titles/buttons
- resizable/reorderable columns
- selectable column alignment
- sorting
- standard or 3D look
- expand/collapse buttons

Use the DLL version with C, C++ (MFC, OWL). Our OCX version supports all popular environments including VB, Delphi, etc. Also available as VCL-style control for Delphi and C++ Builder.

Supports Windows 3.1, Windows NT/2000 and Windows 95/98

SftTree/DLL 4.0 (with C++ classes)	\$279	<small>VISA/MC/AMEX accepted. S&H additional. Site licensing available.</small>
SftTree/OCX 4.0	\$299	
SftTree/VCL 4.0	\$199	

Call today for your free demo!

softtel vdm

NOT THE INDUSTRY STANDARD

318 Tamiami Trail, Suite 14
Punta Gorda, FL 33950

(941) 505-8600
FAX (941) 505-8555

http://www.softelvdm.com

that would be a lot of .5s to type, and at any rate, what if that magic value were to change? There are a number of ways to deal with this issue, but perhaps the simplest is to define a function that encapsulates the details of what you want to do. Assuming you're not the author of Wombat, the function will necessarily have to be a non-member, and you'll have to call it as such:

```
// might be inline, but it doesn't matter
void nap(Wombat& w) { w.sleep(.5); }
```

```
Wombat w;
...
nap(w);
```

And there you have it, your dreaded syntactic inconsistency. When you want to feed your wombats, you make member function calls, but when you want them to nap, you make non-member calls.

If you reflect a bit and are honest with yourself, you'll admit that you have this alleged inconsistency with all the non-trivial classes you use, because no class has every function desired by every client. Every client adds at least a few convenience functions of their own, and these functions are always non-members. C++ programmers are used to this, and they think nothing of it. Some calls use

member syntax, and some use non-member syntax. People just look up which syntax is appropriate for the functions they want to call, then they call them. Life goes on. It goes on especially in the STL portion of the Standard C++ library, where some algorithms are member functions (e.g., `size`), some are non-member functions (e.g., `unique`), and some are both (e.g., `find`). Nobody blinks. Not even you.

Interfaces and Packaging

Herb Sutter has explained that the "interface" to a class (roughly speaking, the functionality provided by the class) includes the non-member functions related to the class, and he's shown that the name lookup rules of C++ support this meaning of "interface" [7,8]. This

**GET INFORMATION
ABOUT ANY TOOL
OR PRODUCT
DISPLAYED IN THIS
ISSUE - INSTANTLY!**

Check out:

www.cuj.com

and select:

Development Tool
Information

All tools are categorized by
function for *QUICK* access!

**FIND THE TOOLS
YOU NEED - NOW!**

C/C++ Users Journal™
Advanced Solutions for C/C++ Programmers



Look Beyond HTML

multi edit

One editing environment supporting the latest internet programming languages as well as workhorses like C++ and Delphi (more than 40 languages supported). Advanced web-editing and site management technology, integral FTP and remote telnet compilation give you the cutting edge tools you need to develop for the online world! There is much more, check out our website for details on the latest release of Multi-Edit!

Try Multi-Edit Today, Free!
www.multiedit.com
The Editor that Thinks Like You Do

Here are just a few of the latest features:

- Enhanced embedded scripting support
- Advanced HTML 4.0 support and online reference
- Improvements to the project manager, hex editing, collapsing and VCS support
- Visual Studio integration
- Over a dozen new languages supported
- Color Printing

Multi-Edit is a trademark of American Cybernetics, Inc., and is backed by an unconditional 30-day guarantee. Other products are the trademarks of their respective publishers

American Cybernetics Inc.
Sales: (800) 899-0100
Voice: (480) 968-1945
Fax: (480) 966-1654
sales@multiedit.com

Outside the US & Canada contact:
Soft/Export
www.softexport.com
info@softexport.com
International Inquiries: +353 1 294 2121
UK: 0800 973098
France: 0800 905823
Germany: 0130 860341

AMERICAN CYBERNETICS

is wonderful news for my “non-friend non-members are better than members” argument, because it means that the decision to make a class-related function a non-friend non-member instead of a member need not even change the interface to that class! Moreover, the liberation of the functions in a class’s interface from the confines of the class definition leads to some wonderful packaging flexibility that would otherwise be unavailable. In particular, it means that the interface to a class may be split across multiple header files.

Suppose the author of the Wombat class discovered that Wombat clients often need a number of convenience functions related to eating, sleeping, and breeding. Such convenience functions are by definition not strictly necessary. The same functionality could be obtained via other (albeit more cumbersome) member function calls. As a result, and in accord with my advice in this article, each convenience function should be a non-friend non-member. But suppose the clients of the convenience functions for eating rarely needed the convenience functions for sleeping or breeding. And suppose the clients of the sleeping and breeding convenience functions also rarely needed the convenience functions for eating and, respectively, breeding and sleeping.

Rather than putting all Wombat-related functions into a single header file, a preferable design would be to partition the Wombat interface across four separate headers, one for core Wombat functionality (primarily the class definition), and one each for convenience functions related to eating, sleeping, and breeding. Clients then include only the headers they need. The resulting software is not only clearer, it also contains fewer gratuitous compilation dependencies [4,9]. This multiple-header approach was adopted for the standard library. The contents

of namespace `std` are spread across 50 different headers. Clients `#include` the headers declaring the parts of the library they care about, and they ignore everything else.

In addition, this approach is extensible. When the declarations for the functions making up a class’s interface are spread across multiple header files, it becomes natural for clients creating application-specific

Templates and Factory Functions at Namespace Scope

In the main article, I argue that static member functions should be made non-members whenever that is possible, because that increases class encapsulation. I consider these two possible implementations for a factory function:

```
// the less encapsulated design
class Widget {
    ...
public:
    static Widget* make(/* params */);
};

// the more encapsulated design
namespace WidgetStuff {
    class Widget { ... };
    Widget* make( /* params */ );
};
```

Andrew Koenig pointed out that the first design (where `make` is static inside the class) enables one to write a template function that invokes `make` without knowing the type of what is being made:

```
template<typename T>
void doSomething( /* params */ )
{
    // invoke T's factory function
    T *pt = T::make( /* params */ );
    ...
}
```

This isn’t possible with the namespace-based design, because there’s no way from inside a template to identify the namespace in which a type parameter is located. That is, there’s no way to figure out what `???` is in the pseudocode below:

```
template<typename T>
void doSomething( /* params */ )
{
    // there's no way to know T's containing namespace!
    T *pt = ???::make( /* params */ );
    ...
}
```

For factory functions and similar functions which can be given uniform names, this means that maximal class encapsulation and maximal template utility are at odds. In such cases, you have to decide which is more important and cater to that. However, for static member functions with class-specific names, the template issue fails to arise, and encapsulation can again assume precedence. □

Zentropic Computing

The RTOS of Champions
HARD REAL-TIME!

REALTIME LINUX

Kernel 2.2 Install CD
10-20 μ sec latency
Industrial Support
SMP & POSIX
LIP-SMACKIN' GOOD!

www.zentropix.com

**SPECIAL
OFFER—
CALL
NOW!**

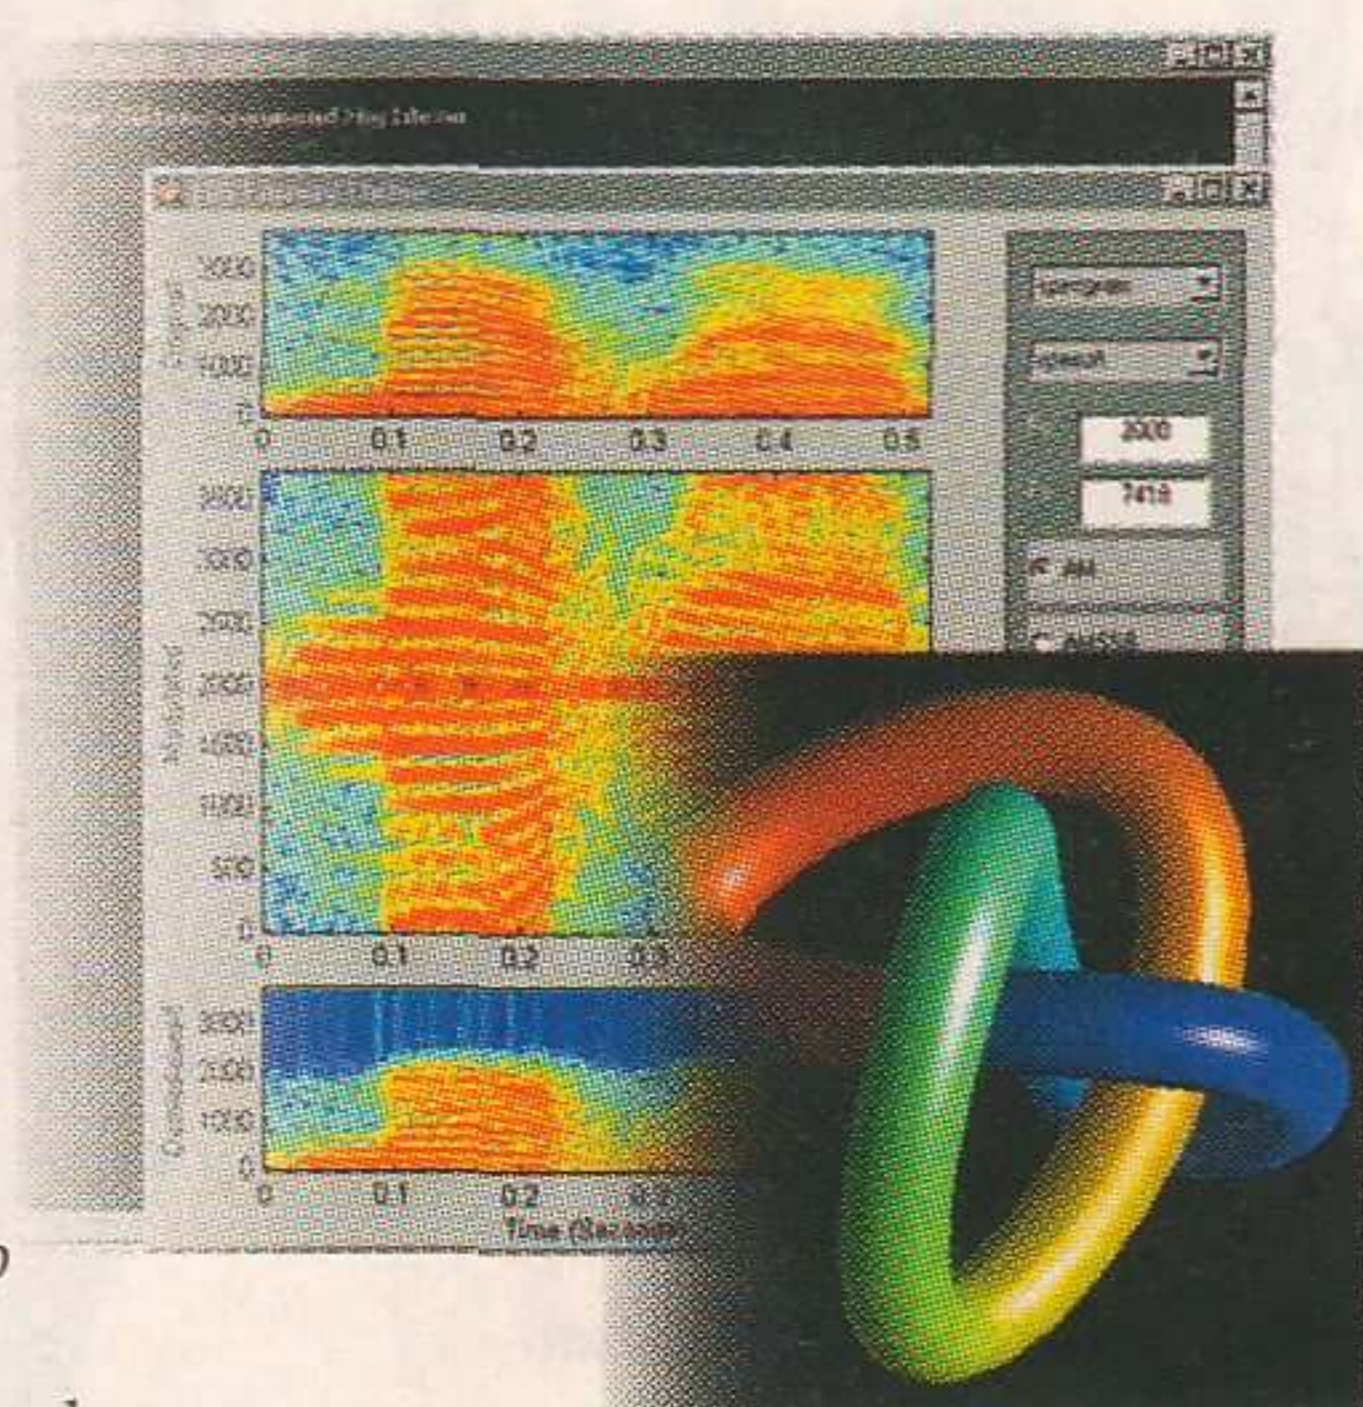
Zero to C/C++ in no time.

Use industry-standard MATLAB to design, develop, and refine your applications. Then use the C/C++ Compiler Suite to convert them into standalone C or C++ code.

MATLAB C/C++ Compiler Suite includes:

- MATLAB Compiler – automatically generates C/C++ code from programs written in MATLAB
- C/C++ Math Library – contains 350 MATLAB math functions accessible from C/C++
- C/C++ Graphics Library – embeds MATLAB graphics and GUIs in your C/C++ code.

Compile
complete
MATLAB
applications,
like this
signal
processing
example, into



C and C++ code.

Information and demos on the Web.

For product information, or to buy online now, visit www.mathworks.com/ccj.



MATLAB®

508-647-7000 • www.mathworks.com/ccj

sets of convenience functions to cluster those functions into a new header file and to `#include` that file as appropriate. In other words, to treat the application-specific convenience functions just like they treat the convenience functions provided by the author of the class. This is as it should be. After all, they're all just convenience functions.

Minimalness and Encapsulation

In *Effective C++*, I argued for class interfaces that are complete and minimal [10]. Such interfaces allow class clients to do anything they might reasonably want to do, but classes contain no more member functions than are absolutely necessary. Adding functions beyond the minimum necessary to let clients get their jobs done, I wrote, decreases the class's comprehensibility and maintainability, plus it increases compilation times for clients. Jack Reeves has written that the addition of member functions beyond those truly required violates the open/closed principle, yields fat class interfaces, and ultimately leads to software rot [11]. That's a fair number of arguments for minimizing the number of member functions in a class, but now we have an additional reason: failure to do so decreases a class's encapsulation.

Of course, a minimal class interface is not necessarily the best interface. I remarked in *Effective C++* that adding functions beyond those truly necessary may be justifiable if it significantly improves the performance of the class, makes the class easier to use, or prevents likely client errors [10]. Based on his work with various string-like classes, Jack Reeves has observed that some functions just don't "feel" right when made non-members, even if they could be non-friend non-members [12]. The "best" interface for a class can be found only by balancing many competing concerns, of which the degree of encapsulation is but one.

Still, the lesson of this article should be clear. Conventional wisdom notwithstanding, use of non-friend non-member functions *improves* a class's encapsulation, and a preference for such functions over member functions makes it easier to design and develop classes with interfaces that are complete and minimal (or close to minimal). Arguments about the naturalness of the resulting calling syntax are generally unfounded, and adoption of a predilection for non-friend non-member functions leads

to packaging strategies for a class's interface that minimize client compilation dependencies while maximizing the number of convenience functions available to them.

It's time to abandon the traditional, but inaccurate, ideas of what it means to be object-oriented. Are you a true encapsulation believer? If so, I know you'll embrace non-friend non-member functions with the fervor they deserve. □

Acknowledgements

Thanks to Arun Kundu for asking the question that led to this article. Thanks also to Jack Reeves, Herb Sutter, Dave Smallberg, Andrei Alexandrescu, Bruce Eckel, Bjarne Stroustrup, and Andrew Koenig for comments on pre-publication drafts that weren't as good as they should have been. (That's why they were drafts.) Finally, great thanks to Adela Novak for organizing the C++ seminars in Lucerne (Switzerland) that led to the many hours on planes and trains that allowed me to write the initial draft of this article.

Notes and References

- [1] Scott Meyers. *Effective C++: 50 Specific Ways to Improve Your Programs and Designs, First Edition* (Addison-Wesley, 1991), Item 19.
- [2] Scott Meyers. *Effective C++, Second Edition* (Addison-Wesley, 1998).
- [3] The algorithm remains unchanged in current printings of *Effective C++*, because I'd have to also add the requisite reasoning (this article), and making such a substantial change to a book already in production simply isn't practical.
- [4] *Effective C++*, Item 34.
- [5] Erich Gamma et al. *Design Patterns, Elements of Reusable Object-Oriented Software* (Addison-Wesley, 1995). Also known as the GOF or "Gang of Four" book.
- [6] James O. Coplien. *Advanced C++: Programming Styles and Idioms* (Addison-Wesley, 1991).
- [7] Herb Sutter. "Sutter's Mill: What's in a Class?" *C++ Report*, March 1998.
- [8] Herb Sutter. *Exceptional C++* (Addison-Wesley, 1999), Items 31-34.
- [9] John Lakos. *Large-Scale C++ Software Design* (Addison-Wesley, 1996).
- [10] *Effective C++*, Item 18.
- [11] Jack Reeves. "(B)leading Edge: How About Namespaces?," *C++ Report*, April 1999.
- [12] Jack Reeves. Personal communication.