



C++ Library Reference

901 San Antonio Road
Palo Alto, , CA 94303-4900
USA 650 960-1300 fax 650 969-9131

Part No: 805-4957
Revision A, February 1999

Copyright Copyright 1999 Sun Microsystems, Inc. 901 San Antonio Road, Palo Alto, California 94303-4900 U.S.A. All rights reserved.

All rights reserved. This product or document is protected by copyright and distributed under licenses restricting its use, copying, distribution, and decompilation. No part of this product or document may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any.

Portions of this product may be derived from the UNIX® system, licensed from Novell, Inc., and from the Berkeley 4.3 BSD system, licensed from the University of California. UNIX is a registered trademark in the United States and in other countries and is exclusively licensed by X/Open Company Ltd. Third-party software, including font technology in this product, is protected by copyright and licensed from Sun's suppliers. RESTRICTED RIGHTS: Use, duplication, or disclosure by the U.S. Government is subject to restrictions of FAR 52.227-14(g)(2)(6/87) and FAR 52.227-19(6/87), or DFAR 252.227-7015(b)(6/95) and DFAR 227.7202-3(a).

Sun, Sun Microsystems, the Sun logo, SunDocs, SunExpress, Solaris, Sun Performance Library, Sun Performance WorkShop, Sun Performance WorkShop Fortran, Sun Visual WorkShop, Sun WorkShop, Sun WorkShop Compilers C, Sun WorkShop Compilers C++, Sun WorkShop Compilers Fortran, Sun WorkShop Memory Monitor, Sun WorkShop Professional, Sun WorkShop Professional C, and Sun WorkShop TeamWare are trademarks, registered trademarks, or service marks of Sun Microsystems, Inc. in the U.S. and other countries. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

The OPEN LOOK® and Sun™ Graphical User Interfaces were developed by Sun Microsystems, Inc. for its users and licensees. Sun acknowledges the pioneering efforts of Xerox Corporation in researching and developing the concept of visual or graphical user interfaces for the computer industry. Sun holds a nonexclusive license from Xerox to the Xerox Graphical User Interface, which license also covers Sun's licensees who implement OPEN LOOK GUIs and otherwise comply with Sun's written license agreements.

THIS PUBLICATION IS PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT.

Copyright 1999 Sun Microsystems, Inc., 901 San Antonio Road, Palo Alto, Californie 94303-4900 U.S.A. Tous droits réservés.

Ce produit ou document est protégé par un copyright et distribué avec des licences qui en restreignent l'utilisation, la copie et la décompilation. Aucune partie de ce produit ou de sa documentation associée ne peut être reproduite sous aucune forme, par quelque moyen que ce soit, sans l'autorisation préalable et écrite de Sun et de ses bailleurs de licence, s'il y en a.

Des parties de ce produit pourront être dérivées du système UNIX® licencié par Novell, Inc. et du système Berkeley 4.3 BSD licencié par l'Université de Californie. UNIX est une marque enregistrée aux Etats-Unis et dans d'autres pays, et licenciée exclusivement par X/Open Company Ltd. Le logiciel détenu par des tiers, et qui comprend la technologie relative aux polices de caractères, est protégé par un copyright et licencié par des fournisseurs de Sun.

Sun, Sun Microsystems, the Sun logo, SunDocs, SunExpress, Solaris, Sun Performance Library, Sun Performance WorkShop, Sun Performance WorkShop Fortran, Sun Visual WorkShop, Sun WorkShop, Sun WorkShop Compilers C, Sun WorkShop Compilers C++, Sun WorkShop Compilers Fortran, Sun WorkShop Memory Monitor, Sun WorkShop Professional, Sun WorkShop Professional C, et Sun WorkShop TeamWare sont des marques de fabrique ou des marques déposées, ou marques de service, de Sun Microsystems, Inc. aux Etats-Unis et dans d'autres pays. Toutes les marques SPARC sont utilisées sous licence et sont des marques de fabrique ou des marques déposées de SPARC International, Inc. aux Etats-Unis et dans d'autres pays. Les produits portant les marques SPARC sont basés sur une architecture développée par Sun Microsystems, Inc.

Les utilisateurs d'interfaces graphiques OPEN LOOK® et Sun™ ont été développés de Sun Microsystems, Inc. pour ses utilisateurs et licenciés. Sun reconnaît les efforts de pionniers de Xerox Corporation pour la recherche et le développement du concept des interfaces d'utilisation visuelle ou graphique pour l'industrie de l'informatique. Sun détient une licence non exclusive de Xerox sur l'interface d'utilisation graphique, cette licence couvrant aussi les licenciés de Sun qui mettent en place les utilisateurs d'interfaces graphiques OPEN LOOK et qui en outre se conforment aux licences écrites de Sun.

CETTE PUBLICATION EST FOURNIE "EN L'ETAT" SANS GARANTIE D'AUCUNE SORTE, NI EXPRESSE NI IMPLICITE, Y COMPRIS, ET SANS QUE CETTE LISTE NE SOIT LIMITATIVE, DES GARANTIES CONCERNANT LA VALEUR MARCHANDE, L'APTITUDE DES PRODUITS A REpondre A UNE UTILISATION PARTICULIERE OU LE FAIT QU'ILS NE SOIENT PAS CONTREFAISANTS DE PRODUITS DE TIERS.



Contents

Preface	vii
1. Introduction to C++ Libraries	1
Other Libraries	1
The C++ Standard Library	1
Tools.h++ Library	2
Man Pages for the Libraries	2
Sun WorkShop Memory Monitor	2
2. The Complex Arithmetic Library	5
The Complex Library	5
Using Complex Library	6
Type <code>complex</code>	6
Constructors of Class <code>complex</code>	6
Arithmetic Operators	7
Mathematical Functions	8
Error Handling	10
Input and Output	12
Mixed-Mode Arithmetic	12
Efficiency	13
Complex Man Pages	14

3. The Iostream Library	15
Predefined Iostreams	15
Basic Structure of Iostream Interaction	16
Iostreams	17
Output Using Iostream	18
Input Using Iostream	21
Defining Your Own Extraction Operators	21
Using the char* Extractor	22
Reading Any Single Character	22
Binary Input	23
Peeking at Input	23
Extracting Whitespace	23
Handling Input Errors	24
Using Iostreams with stdio	24
Creating Iostreams	25
Dealing with Files Using Class fstream	25
Assignment of Iostreams	28
Format Control	29
Manipulators	29
Using Plain Manipulators	30
Parameterized Manipulators	31
Strstreams: Iostreams for Arrays	33
Stdiobufs: Iostreams for stdio files	33
Streambufs	33
Working with Streambufs	33
Using Streambufs	34
Iostream Man Pages	35
Iostream Terminology	36

4.	Using <code>istream</code> in a Multithreaded Environment	39
	Multithreading	39
	Organization of the MT-safe <code>istream</code> Library	40
	Public Conversion Routines	41
	Compiling and Linking with the MT-safe <code>libc</code> Library	42
	MT-safe <code>istream</code> Restrictions	43
	Performance	44
	Interface Changes to the <code>istream</code> Library	46
	New Classes	46
	New Class Hierarchy	46
	New Functions	47
	Global and Static Data	48
	Sequence Execution	49
	Object Locks	49
	Class <code>stream_locker</code>	50
	MT-safe Classes	51
	Object Destruction	52
	An Example Application	52
A.	Associated Man Pages	55
	Index	73

Preface

The *C++ Library Reference* describes the C++ libraries, including:

- C++ Standard Library
- Tools.h++ Class Library
- Sun WorkShop Memory Monitor
- Iostream
- Complex

Who Should Use This Book

This is a reference manual intended for programmers with a working knowledge of the C++ language and some understanding of the Solaris™ system and UNIX® commands.

How This Book Is Organized

Chapter 1” gives an overview of the C++ libraries.

Chapter 2” explains the arithmetic operators and mathematical functions in the library.

Chapter 3” discusses the input and output facility used in C++.

Chapter 4” details how to use the `iostream` library for input and output in a multithreaded environment.

Appendix A” lists the man pages associated with the C++ libraries.

Multiplatform Release

Note - The name of the latest Solaris operating environment release is Solaris 7 but code and path or package path names may use Solaris 2.7 or SunOS 5.7.

The Sun™ WorkShop™ documentation applies to Solaris 2.5.1, Solaris 2.6, and Solaris 7 operating environments on:

- The SPARC™ platform
 - The x86 platform, where x86 refers to the Intel implementation of one of the following: Intel 80386, Intel 80486, Pentium, or the equivalent
-

Note - The term “x86” refers to the Intel 8086 family of microprocessor chips, including the Pentium, Pentium Pro, and Pentium II processors and compatible microprocessor chips made by AMD and Cyrix. In this document, the term “x86” refers to the overall platform architecture. Features described in this book that are particular to a specific platform are differentiated by the terms “SPARC” and “x86” in the text.

Related Books

The following Sun manuals and guides provide additional useful information.

C++ Books

- *C++ User's Guide* provides information on command-line options and how to use the compiler.
- *C++ Programming Guide* discusses issues relating to the use of templates, exception handling, and interfacing with FORTRAN 77.
- *C++ Migration Guide* describes migrations between compiler releases.
- *Standard C++ Library User's Guide* describes how to use the Standard C++ Library.
- *Standard C++ Class Library Reference* provides detail on the Standard C++ Library.

Other Programming Books

- *Tools.h++ User's Guide* provides details on the `Tools.h++` class library.
- *Tools.h++ Class Library Reference* discusses use of the C++ classes for enhancing the efficiency of your programs.
- *C User's Guide* describes compiler options, pragmas, and more.
- *FORTRAN 77 Language Reference Manual* provides a complete language reference.
- *Fortran User's Guide* provides information on command-line options and how to use the compilers.
- *Fortran Programming Guide* discusses issues relating to input/output, libraries, program analysis, debugging, and performance.
- *Fortran Library Reference* gives detail on the language and routines.
- *Numerical Computation Guide* details floating-point computation numerical accuracy issues.

Other Sun WorkShop Books

- *Sun WorkShop Quick Install* provides installation instructions.
- *Sun WorkShop Installation and Licensing Reference* provides supporting installation and licensing information.
- *Using Sun WorkShop* gives information on performing development operations through Sun WorkShop.
- *Debugging a Program With dbx* provides information on using `dbx` commands to debug a program.
- *Analyzing Program Performance With Sun WorkShop* describes the profiling tools; the LoopTool, LoopReport, and LockLint utilities; and use of the Sampling Analyzer to enhance program performance.
- *Sun WorkShop Performance Library Reference Manual* discusses the library of subroutines and functions to perform useful operations in computational linear algebra and Fourier transforms.
- *Sun WorkShop Visual User's Guide* describes how to use Visual to create C++ and Java[™] graphical user interfaces.

Solaris Books

- The Solaris *Linker and Libraries Guide* gives information on linking and libraries.
- The Solaris *Programming Utilities Guide* provides information for developers about the special built-in programming tools available in the SunOS[™] system.

Ordering Sun Documents

The SunDocsSM program provides more than 250 manuals from Sun Microsystems, Inc. If you live in the United States, Canada, Europe, or Japan, you can purchase documentation sets or individual manuals using this program.

For a list of documents and how to order them, see the catalog section of the SunExpressTM Internet site at <http://www.sun.com/sunexpress>.

Accessing Sun Documents Online

Sun WorkShop documentation is available online from several sources:

- The `docs.sun.com` Web site
- AnswerBook2TM collections
- HTML documents
- Online help and release notes

Using the `docs.sun.com` Web site

The `docs.sun.com` Web site enables you to access Sun technical documentation online. You can browse the `docs.sun.com` archive or search for a specific book title or subject. The URL is <http://docs.sun.com>.

Accessing AnswerBook2 Collections

The Sun WorkShop documentation is also available using AnswerBook2 software. To access the AnswerBook2 collections, your system administrator must have installed the AnswerBook2 documents during the installation process (if the documents are not installed, see your system administrator or Chapter 3 of *Sun WorkShop Quick Install* for installation instructions). For information about accessing AnswerBook2 documents, see Chapter 6 of *Sun WorkShop Quick Install*, Solaris installation documentation, or your system administrator.

Note - To access AnswerBook2 documents, Solaris 2.5.1 users must first download AnswerBook2 documentation server software from a Sun Web page. For more information, see Chapter 6 of *Sun WorkShop Quick Install*.

Accessing HTML Documents

The following Sun Workshop documents are available online only in HTML format:

- Tools.h++ Class Library Reference
- Tools.h++ User's Guide
- *Numerical Computation Guide*
- Standard C++ Library User's Guide
- *Standard C++ Class Library Reference*
- *Sun WorkShop Performance Library Reference Manual*
- *Sun WorkShop Visual User's Guide*
- Sun WorkShop Memory Monitor User's Manual

To access these HTML documents:

1. Open the following file through your HTML browser:

install-directory/SUNWspr0/DOC5.0/lib/locale/C/html/index.html

Replace *install-directory* with the name of the directory where your Sun WorkShop software is installed (the default is /opt).

The browser displays an index of the HTML documents for the Sun WorkShop products that are installed.

2. Open a document in the index by clicking the document's title.

Accessing Sun WorkShop Online Help and Release Notes

This release of Sun WorkShop includes an online help system as well as online manuals. To find out more see:

- Online Help. A help system containing extensive task-oriented, context-sensitive help. To access the help, choose Help Help Contents. Help menus are available in all Sun WorkShop windows.
- Release Notes. The Release Notes contain general information about Sun WorkShop and specific information about software limitations and bugs. To access the Release Notes, choose Help Release Notes.

What Typographic Changes Mean

The following table describes the typographic changes used in this book.

TABLE P-1 Typographic Conventions

Typeface or Symbol	Meaning	Example
<i>AaBbCc123</i>	The names of commands, files, and directories; on-screen computer output	Edit your <code>.login</code> file. Use <code>ls -a</code> to list all files. <code>machine_name%</code> You have mail.
AaBbCc123	What you type, contrasted with on-screen computer output	<code>machine_name%</code> su Password:
<i>AaBbCc123</i>	Command-line placeholder: replace with a real name or value	To delete a file, type <code>rm filename</code> .
<i>AaBbCc123</i>	Book titles, new words or terms, or words to be emphasized	Read Chapter 6 in <i>User's Guide</i> . These are called <i>class</i> options. You <i>must</i> be root to do this.

Shell Prompts in Command Examples

The following table shows the default system prompt and superuser prompt for the C shell, Bourne shell, and Korn shell.

TABLE P-2 System Prompts

Shell	Prompt
C shell prompt	machine_name%
C shell superuser prompt	machine_name#
Bourne shell and Korn shell prompt	\$
Bourne shell and Korn shell superuser prompt	#

Introduction to C++ Libraries

The C++ class libraries are modular components of reusable code. Using class libraries, you can integrate blocks of code that have been previously built and tested.

A C++ library consists of one or more header files and an object library. The header files provide class and other definitions needed to access the library functions. The object library provides compiled functions and data that are linked with your program to produce an executable program.

This manual describes two class libraries provided with the C++ compiler:

- Complex numbers, described in Chapter 2
- Iostreams, described in Chapter 3

For a discussion of building shared and static libraries, and using libraries, see the *C++ User's Guide*.

Other Libraries

In addition to the Complex and Iostreams libraries, this release provides the C++ Standard library, the Tools.h++ library, and the Sun WorkShop Memory Monitor library.

The C++ Standard Library

The C++ Standard Library (libCstd) is based on the RogueWave™ Standard C++ Library, Version 2.0.1. This library is available only for the default mode (-compat=5) of the compiler and is not supported with use of the -compat or

-compat=4 options. For details, see the *Standard C++ Library User's Guide*, the *Standard C++ Library User's Guide*, and the *Standard C++ Class Library Reference*.

Tools.h++ Library

Tools.h++ is a C++ foundation class library from RogueWave. Version 7.0.7 of this library is provided with this release. For information on using this library in the default mode of the compiler, see the *C++ Users Guide*.

For further information about this library, see:

- *Tools.h++ User's Guide* (Version 7.0.7)
- *Tools.h++ Class Library Reference* (Version 7.0.7)

Man Pages for the Libraries

The man pages associated with the libraries described in this manual are located in:

install-directory/SUNWspro/man/man3

install-directory/SUNWspro/man/man3cc4

install-directory/SUNWspro/man/man3c++

where the default *install-directory* is /opt.

To access these man pages, ensure that your MANPATH includes *install-directory*/SUNWspro/man (see the *Sun WorkShop Quick Install*).

To access man pages for the Sun WorkShop Compiler C++ 4.2 libraries, type:

```
man -s 3CC4 library-name
```

To access man pages for the Sun WorkShop Compiler C++ 5.0 libraries, type:

```
man library-name
```

Sun WorkShop Memory Monitor

The Sun WorkShop Memory Monitor provides facilities to automatically report and fix memory leaks, memory fragmentation, and premature frees. It has three modes of operation:

- Debugging
- Deployment
- Garbage collection

These modes are dependent on the library you link your application with.

The components of the Sun WorkShop Memory Monitor are:

- `libgc`—the library used in garbage collection and deployment modes
- `libgc_dbg`—the library used in memory debugging mode
- `gcmonitor`—the daemon used in memory debugging mode

For complete documentation on the Sun WorkShop Memory Monitor, see `/install-directory/SUNWspr0/SC5.0/htmldocs/locale/C/gc/index.html`, where the default *install-directory* is `/opt`.

Man pages for the Sun WorkShop Memory Monitor are located in:

`install-directory/SUNWspr0/man/man1`

`install-directory/SUNWspr0/man/man3`

where the default *install-directory* is `/opt`.

The Complex Arithmetic Library

Complex numbers are numbers made up of a *real* and an *imaginary* part. For example:

```
3.2 + 4i  
1 + 3i  
1 + 2.3i
```

In the degenerate case, $0 + 3i$ is an entirely imaginary number generally written as $3i$, and $5 + 0i$ is an entirely real number generally written as 5 . You can represent complex numbers using the `complex` data type.

Note - The Complex Arithmetic library (`libcomplex`) is available only for `-compat=4`, and is not obsolete. For `-compat=5`, this is part of the C++ Standard Library, `libCstd`.

The Complex Library

The complex arithmetic library implements a complex number data type as a new data type and provides:

- Operators
- Mathematical functions (defined for the built-in numerical types)
- Extensions (for iostreams that allow input and output of complex numbers)
- Error handling mechanisms

Complex numbers can also be represented as an *absolute value* (or *magnitude*) and an *argument* (or *angle*). The library provides functions to convert between the real and

imaginary (Cartesian) representation and the magnitude and angle (polar) representation.

The *complex conjugate* of a number has the opposite sign in its imaginary part.

Using Complex Library

To use the complex library, include the header file `complex.h` in your program, and compile and link with the `-library=complex` option.

Type `complex`

The complex arithmetic library defines one class: class `complex`. An object of class `complex` can hold a single complex number. The complex number is constructed of two parts:

- The real part
- The imaginary part

```
class complex {
    double re, im;
};
```

The numerical values of each part are held in fields of type `double`. Here is the relevant part of the definition of `complex`:

The value of an object of class `complex` is a pair of `double` values. The first value represents the real part; the second value represents the imaginary part.

Constructors of Class `complex`

There are two constructors for `complex`. Their definitions are:

```
complex::complex() { re=0.0; im=0.0; }
complex::complex(double r, double i = 0.0) { re=r; im=i; }
```

If you declare a complex variable without parameters, the first constructor is used and the variable is initialized, so that both parts are 0. The following example creates a complex variable whose real and imaginary parts are both 0:

```
complex aComp;
```

You can give either one or two parameters. In either case, the second constructor is used. When you give only one parameter, it is taken as the value for the real part and the imaginary part is set to 0. For example:

```
complex aComp(4.533);
```

creates a complex variable with the value:

```
4.533 + 0i
```

If you give two values, the first is taken as the value of the real part and the second as the value of the imaginary part. For example:

```
complex aComp(8.999, 2.333);
```

creates a complex variable with the value:

```
8.999 + 2.333i
```

You can also create a complex number using the `polar` function, which is provided in the complex arithmetic library (see “Mathematical Functions ” on page 8). The `polar` function creates a complex value given a pair of polar coordinates, magnitude and angle.

There is no destructor for type `complex`.

Arithmetic Operators

The complex arithmetic library defines all the basic arithmetic operators. Specifically, the following operators work in the usual way and with the usual precedence:

```
+ - / * =
```

The operator `-` has its usual binary and unary meanings.

In addition, you can use the following operators in the usual way:

```
+= -= *= /=
```

However, these last four operators do not produce values that you can use in expressions. For example, the following does not work:

```
complex a, b;
...
if ((a+=2)==0) {...}; // illegal
b = a *= b; // illegal
```

You can also use the following equality operators in their regular meaning:

`==` `!=`

When you mix real and complex numbers in an arithmetic expression, C++ uses the complex operator function and converts the real values to complex values.

Mathematical Functions

The complex arithmetic library provides a number of mathematical functions. Some are peculiar to complex numbers; the rest are complex-number versions of functions in the standard C mathematical library.

All of these functions produce a result for every possible argument. If a function cannot produce a mathematically acceptable result, it calls `complex_error` and returns some suitable value. In particular, the functions try to avoid actual overflow and call `complex_error` with a message instead. The following tables describe the remainder of the complex arithmetic library functions.

TABLE 2-1 Complex Arithmetic Library Functions

Complex Arithmetic Library Function	Description
<code>double abs(const complex)</code>	Returns the magnitude of a complex number.
<code>double ang(const complex)</code>	Returns the angle of a complex number.
<code>complex conj(const complex)</code>	Returns the complex conjugate of its argument.
<code>double imag(const complex&)</code>	Returns the imaginary part of a complex number.
<code>double norm(const complex)</code>	Returns the square of the magnitude of its argument. Faster than <code>abs</code> , but more likely to cause an overflow. For comparing magnitudes.

TABLE 2-1 Complex Arithmetic Library Functions *(continued)*

Complex Arithmetic Library Function	Description
<code>complex polar(double mag, double ang=0.0)</code>	Takes a pair of polar coordinates that represent the magnitude and angle of a complex number and returns the corresponding complex number.
<code>double real(const complex&)</code>	Returns the real part of a complex number.

TABLE 2-2 Complex Mathematical and Trigonometric Functions

Complex Arithmetic Library Function	Description
<code>complex acos(const complex)</code>	Returns the angle whose cosine is its argument.
<code>complex asin(const complex)</code>	Returns the angle whose sine is its argument.
<code>complex atan(const complex)</code>	Returns the angle whose tangent is its argument.
<code>complex cos(const complex)</code>	Returns the cosine of its argument.
<code>complex cosh(const complex)</code>	Returns the hyperbolic cosine of its argument.
<code>complex exp(const complex)</code>	Computes e^{*x} , where e is the base of the natural logarithms, and x is the argument given to <code>exp</code> .
<code>complex log(const complex)</code>	Returns the natural logarithm of its argument.
<code>complex log10(const complex)</code>	Returns the common logarithm of its argument.

TABLE 2-2 Complex Mathematical and Trigonometric Functions *(continued)*

Complex Arithmetic Library Function	Description
<code>complex pow(double b, const complex exp)</code>	Takes two arguments: <code>pow(b, exp)</code> . It raises <i>b</i> to the power of <i>exp</i> .
<code>complex pow(const complex b, int exp)</code>	
<code>complex pow(const complex b, double exp)</code>	
<code>complex pow(const complex b, const complex exp)</code>	
<code>complex sin(const complex)</code>	Returns the sine of its argument.
<code>complex sinh(const complex)</code>	Returns the hyperbolic sine of its argument.
<code>complex sqrt(const complex)</code>	Returns the square root of its argument.
<code>complex tan(const complex)</code>	Returns the tangent of its argument.
<code>complex tanh(const complex)</code>	Returns the hyperbolic tangent of its argument.

Error Handling

The complex library has these definitions for error handling:

```
extern int errno;
class c_exception { ... };
int complex_error(c_exception&);
```

The external variable `errno` is the global error state from the C library. `errno` can take on the values listed in the standard header `errno.h` (see the man page `perror(3)`). No function sets `errno` to zero, but many functions set it to other values.

To determine whether a particular operation fails:

1. Set `errno` to zero before the operation.

2. Test the operation.

The function `complex_error` takes a reference to type `c_exception` and is called by the following complex arithmetic library functions:

- `exp`
- `log`
- `log10`
- `sinh`
- `cosh`

The default version of `complex_error` returns zero. This return of zero means that the default error handling takes place. You can provide your own replacement function `complex_error` that performs other error handling. Error handling is described in the man page `cplxerr(3C++)`.

Default error handling is described in the man pages `cplxtrig(3C++)` and `cplxexp(3C++)`. It is also summarized in the following table.

TABLE 2-3 Complex Arithmetic Library Functions

Complex Arithmetic Library Function	Default Error Handling Summary
<code>exp</code>	If overflow occurs, sets <code>errno</code> to <code>ERANGE</code> and returns a huge complex number.
<code>log</code> , <code>log10</code>	If the argument is zero, sets <code>errno</code> to <code>EDOM</code> and returns a huge complex number.
<code>sinh</code> , <code>cosh</code>	If the imaginary part of the argument causes overflow, returns a complex zero. If the real part causes overflow, returns a huge complex number. In either case, sets <code>errno</code> to <code>ERANGE</code> .

Input and Output

The complex arithmetic library provides default *extractors* and *inserters* for complex numbers, as shown in the following example:

```
ostream& operator<<(ostream&, const complex&); //inserter
istream& operator>>(istream&, complex&); //extractor
```

See sections “Basic Structure of `Iostream` Interaction” on page 16 and “Output Using `Iostream`” on page 18 for basic information on extractors and inserters.

For input, the complex extractor `>>` extracts a pair of numbers (surrounded by parentheses and separated by a comma) from the input stream and reads them into a complex object. The first number is taken as the value of the real part; the second as the value of the imaginary part. For example, given the declaration and input statement:

```
complex x;
cin >> x;
```

and the input `(3.45, 5)`, the value of `x` is equivalent to $3.45 + 5.0i$. The reverse is true for inserters. Given `complex x(3.45, 5)`, `cout<<x` prints `(3.45, 5)`.

The input usually consists of a pair of numbers in parentheses separated by a comma; white space is optional. If you provide a single number, with or without parentheses and white space, the extractor sets the imaginary part of the number to zero. Do not include the symbol `i` in the input text.

The inserter inserts the values of the real and imaginary parts enclosed in parentheses and separated by a comma. It does not include the symbol `i`. The two values are treated as doubles.

Mixed-Mode Arithmetic

Type `complex` is designed to fit in with the built-in arithmetic types in mixed-mode expressions. Arithmetic types are silently converted to type `complex`, and there are `complex` versions of the arithmetic operators and most mathematical functions. For example:

```
int i, j;
double x, y;
complex a, b;
a = sin((b+i)/y) + x/j;
```

The expression `b+i` is mixed-mode. Integer `i` is converted to type `complex` via the constructor `complex::complex(double, double=0)`, the integer first being converted to type `double`. The result is to be divided by `y`, a `double`, so `y` is also converted to `complex` and the complex divide operation is used. The quotient is thus type `complex`, so the complex sine routine is called, yielding another `complex` result, and so on.

Not all arithmetic operations and conversions are implicit, or even defined, however. For example, complex numbers are not well-ordered, mathematically speaking, and complex numbers can be compared for equality only.

```
complex a, b;
a == b; // OK
a != b; // OK
a < b; // error: operator < cannot be applied to type complex
a >= b; // error: operator >= cannot be applied to type complex
```

Similarly, there is no automatic conversion from type `complex` to any other type, because the concept is not well-defined. You can specify whether you want the real part, imaginary part, or magnitude, for example.

```
complex a;
double f(double);
f(abs(a)); // OK
f(a); // error: no match for f(complex)
```

Efficiency

The design of the `complex` class addresses efficiency concerns.

The simplest functions are declared `inline` to eliminate function call overhead.

Several overloaded versions of functions are provided when that makes a difference. For example, the `pow` function has versions that take exponents of type `double` and `int` as well as `complex`, since the computations for the former are much simpler.

The standard C math library header `math.h` is included automatically when you include `complex.h`. The C++ overloading rules then result in efficient evaluation of expressions like this:

```
double x;
complex x = sqrt(x);
```

In this example, the standard math function `sqrt(double)` is called, and the result is converted to type `complex`, rather than converting to type `complex` first and then calling `sqrt(complex)`. This result falls right out of the overload resolution rules, and is precisely the result you want.

Complex Man Pages

The remaining documentation of the complex arithmetic library consists of the man pages listed in the following table.

TABLE 2-4 Man Pages for Type `complex`

Man Page	Overview
<code>cplx.intro(3C++)</code>	General introduction to the complex arithmetic library
<code>cartpol(3C++)</code>	Cartesian and polar functions
<code>cplxerr(3C++)</code>	Error-handling functions
<code>cplxexp(3C++)</code>	Exponential, log, and square root functions
<code>cplxops(3C++)</code>	Arithmetic operator functions
<code>cplxtrig(3C++)</code>	Trigonometric functions

The Iostream Library

C++, like C, has no built-in input or output statements. Instead, I/O facilities are provided by a library. The standard C++ I/O library is the `iostream` library.

This chapter provides an introduction to the `iostream` library and examples of its use. It does not provide a complete description of the `iostream` library. See the `iostream` library man pages for more details.

Note - This library is part of `libC` for `-compat=4`. For standard mode, we recommend using the `iostream` library from the standard C++ library. There is a separate library for `-compat=5` (`libiostream`).

This library is for compatibility with existing sources for `-compat=5`. The new `iostreams` library, conforming to the ISO standard, is available in the C++ standard library, `libCstd`. If you have sources that use the old `iostreams`, and if you want to compile your sources with `-compat=5` (default), use the `iostream` library as follows:

```
CC filename -library=iostream
```

or

```
CC -compat=5 filename -library=iostream
```

Make sure that your source file includes `iostream.h`.

Predefined Iostreams

There are four predefined `iostreams`:

- `cin`, connected to standard input

- `cout`, connected to standard output
- `cerr`, connected to standard error
- `clog`, connected to standard error

The predefined `iostreams` are fully buffered, except for `cerr`. See “Output Using `Iostream`” on page 18 and “Input Using `Iostream`” on page 21.

Basic Structure of `Iostream` Interaction

By including the `iostream` library, a program can use any number of input or output streams. Each stream has some source or sink, which may be one of the following:

- Standard input
- Standard output
- Standard error
- A file
- An array of characters

A stream can be restricted to input or output, or a single stream can allow both input and output. The `iostream` library implements these streams using two processing layers.

- The lower layer implements sequences, which are simply streams of characters. These sequences are implemented by the `streambuf` class, or by classes derived from it.
- The upper layer performs formatting operations on sequences. These formatting operations are implemented by the `istream` and `ostream` classes, which have as a member an object of a type derived from class `streambuf`. An additional class, `iostream`, is for streams on which both input and output can be performed.

Standard input, output, and error are handled by special class objects derived from class `istream` or `ostream`.

The `ifstream`, `ofstream`, and `fstream` classes, which are derived from `istream`, `ostream`, and `iostream` respectively, handle input and output with files.

The `istrstream`, `ostrstream`, and `strstream` classes, which are derived from `istream`, `ostream`, and `iostream` respectively, handle input and output to and from arrays of characters.

When you open an input or output stream, you create an object of one of these types, and associate the `streambuf` member of the stream with a device or file. You generally do this association through the stream constructor, so you don't work with the `streambuf` directly. The `iostream` library predefines stream objects for the

standard input, standard output, and error output, so you don't have to create your own objects for those streams.

You use operators or `iostream` member functions to insert data into a stream (output) or extract data from a stream (input), and to control the format of data that you insert or extract.

When you want to insert and extract a new data type—one of your classes—you generally overload the insertion and extraction operators.

Iostreams

To use `iostream` routines, you must include the header files for the part of the library you need. The header files are described in the following table.

TABLE 3-1 `iostream` Routine Header Files

Header File	Description
<code>iostream.h</code>	Declares basic features of <code>iostream</code> library.
<code>fstream.h</code>	Declares <code>istream</code> s and <code>streambuf</code> s specialized to files. Includes <code>iostream.h</code> .
<code>strstream.h</code>	Declares <code>istream</code> s and <code>streambuf</code> s specialized to character arrays. Includes <code>iostream.h</code> .
<code>iomanip.h</code>	Declares manipulators: values you insert into or extract from <code>istream</code> s to have different effects. Includes <code>iostream.h</code> .
<code>stdiostream.h</code>	(obsolete) Declares <code>istream</code> s and <code>streambuf</code> s specialized to use <code>stdio</code> <code>FILE</code> s. Includes <code>iostream.h</code> .
<code>stream.h</code>	(obsolete) Includes <code>iostream.h</code> , <code>fstream.h</code> , <code>iomanip.h</code> , and <code>stdiostream.h</code> . For compatibility with old-style streams from C++ version 1.2.

You usually don't need all of these header files in your program. Include only the ones that contain the declarations you need. The `iostream` library is part of `libc`, and is linked automatically by the `CC` driver.

Output Using `ostream`

Output using `ostream` usually relies on the overloaded left-shift operator (`<<`) which, in the context of `ostream`, is called the insertion operator. To output a value to standard output, you insert the value in the predefined output stream `cout`. For example, given a value `someValue`, you send it to standard output with a statement like:

```
cout << someValue;
```

The insertion operator is overloaded for all built-in types, and the value represented by `someValue` is converted to its proper output representation. If, for example, `someValue` is a `float` value, the `<<` operator converts the value to the proper sequence of digits with a decimal point. Where it inserts `float` values on the output stream, `<<` is called the float inserter. In general, given a type `X`, `<<` is called the `X` inserter. The format of output and how you can control it is discussed in the `ios(3C++)` man page.

The `ostream` library does not support user-defined types. If you define types that you want to output in your own way, you must define an inserter (that is, overload the `<<` operator) to handle them correctly.

The `<<` operator can be applied repetitively. To insert two values on `cout`, you can use a statement like the one in the following example:

```
cout << someValue << anotherValue;
```

The output from the above example will show no space between the two values. So you may want to write the code this way:

```
cout << someValue << " " << anotherValue;
```

The `<<` operator has the precedence of the left shift operator (its built-in meaning). As with other operators, you can always use parentheses to specify the order of action. It is often a good idea to use parentheses to avoid problems of precedence. Of the following four statements, the first two are equivalent, but the last two are not.

```
cout << a+b;      // + has higher precedence than <<
cout << (a+b);
cout << (a&y);   // << has precedence higher than &
cout << a&y;     // probably an error: (cout << a) & y
```


Defining Your Own Insertion Operator

The following example defines a `string` class:

```
#include <stdlib.h>
#include <iostream.h>

class string {
private:
    char* data;
    size_t size;

public:
    // (functions not relevant here)

    friend ostream& operator<<(ostream&, const string&);
    friend istream& operator>>(istream&, string&);
};
```

The insertion and extraction operators must in this case be defined as friends because the data part of the `string` class is private.

```
ostream& operator<< (ostream& ostr, const string& output) { return ostr << output.data; }
```

Here is the definition of `operator<<` overloaded for use with strings.

```
cout << string1 << string2;
```

`operator<<` takes `ostream&` (that is, a reference to an `ostream`) as its first argument and returns the same `ostream`, making it possible to combine insertions in one statement.

Handling Output Errors

Generally, you don't have to check for errors when you overload `operator<<` because the `iostream` library is arranged to propagate errors.

When an error occurs, the `iostream` where it occurred enters an error state. Bits in the `iostream`'s state are set according to the general category of the error. The inserters defined in `iostream` ignore attempts to insert data into any stream that is in an error state, so such attempts do not change the `iostream`'s state.

In general, the recommended way to handle errors is to periodically check the state of the output stream in some central place. If there is an error, you should handle it in some way. This chapter assumes that you define a function `error`, which takes a

string and aborts the program. `error` is not a predefined function. See “Handling Input Errors ” on page 24 for an example of an error function. You can examine the state of an `iostream` with the operator `!`, which returns a nonzero value if the `iostream` is in an error state. For example:

```
if (!cout) error( "output error");
```

There is another way to test for errors. The `ios` class defines operator `void *()`, so it returns a `NULL` pointer when there is an error. You can use a statement like:

```
if (cout << x) return ; // return if successful
```

You can also use the function `good`, a member of `ios`:

```
if ( cout.good() ) return ; // return if successful
```

The error bits are declared in the enum:

```
enum io_state { goodbit=0, eofbit=1, failbit=2, badbit=4, hardfail=0x80} ;
```

For details on the error functions, see the `iostream` man pages.

Flushing

As with most I/O libraries, `iostream` often accumulates output and sends it on in larger and generally more efficient chunks. If you want to flush the buffer, you simply insert the special value `flush`. For example:

```
cout << "This needs to get out immediately." << flush ;
```

`flush` is an example of a kind of object known as a *manipulator*, which is a value that can be inserted into an `iostream` to have some effect other than causing output of its value. It is really a function that takes an `ostream&` or `istream&` argument and returns its argument after performing some actions on it (see “Manipulators ” on page 29).

Binary Output

To obtain output in the raw binary form of a value, use the member function `write` as shown in the following example. This example shows the output in the raw binary form of `x`.

```
cout.write((char*)&x, sizeof(x));
```

The previous example violates type discipline by converting `&x` to `char*`. Doing so is normally harmless, but if the type of `x` is a class with pointers, virtual member functions, or one that requires nontrivial constructor actions, the value written by the above example cannot be read back in properly.

Input Using `Iostream`

Input using `iostream` is similar to output. You use the extraction operator `>>` and you can string together extractions the way you can with insertions. For example:

```
cin >> a >> b ;
```

This statement gets two values from standard input. As with other overloaded operators, the extractors used depend on the types of `a` and `b` (and two different extractors are used if `a` and `b` have different types). The format of input and how you can control it is discussed in some detail in the `ios(3C++)` man page. In general, leading whitespace characters (spaces, newlines, tabs, form-feeds, and so on) are ignored.

Defining Your Own Extraction Operators

When you want input for a new type, you overload the extraction operator for it, just as you overload the insertion operator for output.

Class `string` defines its extraction operator in the following code example:

CODE EXAMPLE 3-1 `string` Extraction Operator

```
istream& operator>> (istream& istr,
string& input)
{
    const int maxline = 256;
    char holder[maxline];
    istr.get(holder, maxline, '\n');
    input = holder;
    return istr;
}
```

The `get` function reads characters from the input stream `istr` and stores them in `holder` until `maxline-1` characters have been read, or a new line is encountered, or EOF, whichever happens first. The data in `holder` is then null-terminated. Finally, the characters in `holder` are copied into the target string.

By convention, an extractor converts characters from its first argument (in this case, `istream& istr`), stores them in its second argument, which is always a reference, and returns its first argument. The second argument must be a reference because an extractor is meant to store the input value in its second argument.

Using the `char*` Extractor

This predefined extractor is mentioned here because it can cause problems. Use it like this:

```
char x[50];
cin >> x;
```

This extractor skips leading whitespace and extracts characters and copies them to `x` until it reaches another whitespace character. It then completes the string with a terminating null (0) character. Be careful, because input can overflow the given array.

You must also be sure the pointer points to allocated storage. For example, here is a common error:

```
char * p; // not initialized cin >> p;
```

There is no telling where the input data will be stored, and it may cause your program to abort.

Reading Any Single Character

In addition to using the `char` extractor, you can get a single character with either form of the `get` member function. For example:

```
char c;
cin.get(c); // leaves c unchanged if input fails

int b;
b = cin.get(); // sets b to EOF if input fails
```

Note - Unlike the other extractors, the `char` extractor does not skip leading whitespace.

Here is a way to skip only blanks, stopping on a tab, newline, or any other character:

```
int a;
do {
    a = cin.get();
}
while( a == ' ' );
```

Binary Input

If you need to read binary values (such as those written with the member function `write`), you can use the `read` member function. The following example shows how to input the raw binary form of `x` using the `read` member function, and is the inverse of the earlier example that uses `write`.

```
cin.read((char*)&x, sizeof(x));
```

Peeking at Input

You can use the `peek` member function to look at the next character in the stream without extracting it. For example:

```
if (cin.peek() != c) return 0;
```

Extracting Whitespace

By default, the `iostream` extractors skip leading whitespace. You can turn off the *skip* flag to prevent this from happening. The following example turns off whitespace skipping from `cin`, then turns it back on:

```
cin.unsetf(ios::skipws); // turn off whitespace skipping
. . .
cin.setf(ios::skipws); // turn it on again
```

You can use the `istream` manipulator `ws` to remove leading whitespace from the `istream`, whether or not skipping is enabled. The following example shows how to remove the leading whitespace from `istream` `istr`:

```
istr >> ws;
```

Handling Input Errors

By convention, an extractor whose first argument has a nonzero error state should not extract anything from the input stream and should not clear any error bits. An extractor that fails should set at least one error bit.

As with output errors, you should check the error state periodically and take some action, such as aborting, when you find a nonzero state. The `!` operator tests the error state of an `istream`. For example, the following code produces an input error if you type alphabetic characters for input:

```
#include <unistd.h>
#include <iostream.h>
void error (const char* message) {
    cerr << message << "\n" ;
    exit(1);
}
int main() {
    cout << "Enter some characters: ";
    int bad;
    cin >> bad;
    if (!cin) error("aborted due to input error");
    cout << "If you see this, not an error." << "\n";
    return 0;
}
```

Class `ios` has member functions that you can use for error handling. See the man pages for details.

Using Iostreams with `stdio`

You can use `stdio` with C++ programs, but problems can occur when you mix `iostreams` and `stdio` in the same standard stream within a program. For example, if you write to both `stdout` and `cout`, independent buffering occurs and produces unexpected results. The problem is worse if you input from both `stdin` and `cin`, since independent buffering may turn the input into trash.

To eliminate this problem with standard input, standard output and standard error, use the following instruction before performing any input or output. It connects all the predefined `iostreams` with the corresponding predefined `stdio` `FILES`.

```
ios::sync_with_stdio();
```

Such a connection is not the default because there is a significant performance penalty when the predefined streams are made unbuffered as part of the connection. You can use both `stdio` and `iostreams` in the same program applied to different files. That is, you can write to `stdout` using `stdio` routines and write to other files attached to `iostreams`. You can open `stdio` `FILES` for input and also read from `cin` so long as you don't also try to read from `stdin`.

Creating Iostreams

To read or write a stream other than the predefined `iostreams`, you need to create your own `iostream`. In general, that means creating objects of types defined in the `iostream` library. This section discusses the various types available.

Dealing with Files Using Class `fstream`

Dealing with files is similar to dealing with standard input and standard output; classes `ifstream`, `ofstream`, and `fstream` are derived from classes `istream`, `ostream`, and `iostream`, respectively. As derived classes, they inherit the insertion and extraction operations (along with the other member functions) and also have members and constructors for use with files.

Include the file `fstream.h` to use any of the `fstreams`. Use an `ifstream` when you only want to perform input, an `ofstream` for output only, and an `fstream` for a stream on which you want to perform both input and output. Use the name of the file as the constructor argument.

For example, copy the file `thisFile` to the file `thatFile` as in the following example:

```
ifstream fromFile("thisFile");
if (!fromFile)
    error("unable to open 'thisFile' for input");
ofstream toFile ("thatFile");
if ( !toFile )
    error("unable to open 'thatFile' for output");
char c ;
while (toFile && fromFile.get(c)) toFile.put(c);
```

This code:

- Creates an `ifstream` object called `fromFile` with a default mode of `ios::in` and connects it to `thisFile`. It opens `thisFile`.
- Checks the error state of the new `ifstream` object and, if it is in a failed state, calls the `error` function, which must be defined elsewhere in the program.
- Creates an `ofstream` object called `toFile` with a default mode of `ios::out` and connects it to `thatFile`.
- Checks the error state of `toFile` as above.
- Creates a `char` variable to hold the data while it is passed.
- Copies the contents of `fromFile` to `toFile` one character at a time.

Note - It is, of course, undesirable to copy a file this way, one character at a time. This code is provided just as an example of using `fstreams`. You should instead insert the `streambuf` associated with the input stream into the output stream. See “Streambufs” on page 33, and the man page `sbufpub(3C++)`.

Open Mode

The mode is constructed by `or`-ing bits from the enumerated type `open_mode`, which is a public type of class `ios` and has the definition:

```
enum open_mode {binary=0, in=1, out=2, ate=4, app=8, trunc=0x10,
  nocreate=0x20, noreplace=0x40};
```

Note - The `binary` flag is not needed on Unix, but is provided for compatibility with systems that do need it. Portable code should use the `binary` flag when opening binary files.

You can open a file for both input and output. For example, the following code opens file `someName` for both input and output, attaching it to the `fstream` variable `inoutFile`.

```
fstream inoutFile("someName", ios::in|ios::out);
```


Declaring an `fstream` Without Specifying a File

You can declare an `fstream` without specifying a file and open the file later. For example, the following creates the `ofstream toFile` for writing.

```
ofstream toFile;
toFile.open(argv[1], ios::out);
```

Opening and Closing Files

You can close the `fstream` and then open it with another file. For example, to process a list of files provided on the command line:

```
ifstream infile;
for (char** f = &argv[1]; *f; ++f) {
    infile.open(*f, ios::in);
    ...;
    infile.close();
}
```

Opening a File Using a File Descriptor

If you know a file descriptor, such as the integer `1` for standard output, you can open it like this:

```
ofstream outfile;
outfile.attach(1);
```

When you open a file by providing its name to one of the `fstream` constructors or by using the `open` function, the file is automatically closed when the `fstream` is destroyed (by a `delete` or when it goes out of scope). When you attach a file to an `fstream`, it is not automatically closed.

Repositioning within a File

You can alter the reading and writing position in a file. Several tools are supplied for this purpose.

- `streampos` is a type that can record a position in an `iostream`.
- `tellg` (`tellp`) is an `istream` (`ostream`) member function that reports the file position. Since `istream` and `ostream` are the parent classes of `fstream`, `tellg` and `tellp` can also be invoked as a member function of the `fstream` class.

- `seekg` (`seekp`) is an `istream` (`ostream`) member function that finds a given position.
- The `seek_dir` enum specifies relative positions for use with `seek`.

```
enum seek_dir { beg=0, cur=1, end=2 }
```

For example, given an `fstream` `aFile`:

```
streampos original = aFile.tellp(); //save current position
aFile.seekp(0, ios::end); //reposition to end of file
aFile << x; //write a value to file
aFile.seekp(original); //return to original position
```

`seekg` (`seekp`) can take one or two parameters. When it has two parameters, the first is a position relative to the position indicated by the `seek_dir` value given as the second parameter. For example:

```
aFile.seekp(-10, ios::end);
```

moves to 10 bytes from the end while

```
aFile.seekp(10, ios::cur);
```

moves to 10 bytes forward from the current position.

Note - Arbitrary seeks on text streams is not portable, but you can always return to a previously saved `streampos` value.

Assignment of `Iostreams`

`Iostreams` does not allow assignment of one stream to another.

The problem with copying a stream object is that there are then two versions of the state information, such as a pointer to the current write position within an output file, which can be changed independently. As a result, problems could occur.

Format Control

Format control is discussed in detail in the in the man page `ios(3C++)`.

Manipulators

Manipulators are values that you can insert into or extract from `iostreams` to have special effects.

Parameterized manipulators are manipulators that take one or more parameters.

Because manipulators are ordinary identifiers, and therefore use up possible names, `iostream` doesn't define them for every possible function. A number of manipulators are discussed with member functions in other parts of this chapter.

There are 13 predefined manipulators, as described in Table 3-2. When using that table, assume the following:

- `i` has type `long`.
- `n` has type `int`.
- `c` has type `char`.
- `istr` is an input stream.
- `ostr` is an output stream.

TABLE 3-2 `Iostream` Predefined Manipulators

Predefined Manipulator	Description
<code>std::dec</code>	Makes the integer conversion base 10.
<code>std::endl</code>	Inserts a newline character (' <code>\n</code> ') and invokes <code>ostream::flush()</code> .
<code>std::ends</code>	Inserts a null (<code>0</code>) character. Useful when dealing with <code>stringstreams</code> .
<code>std::flush</code>	Invokes <code>ostream::flush()</code> .
<code>std::hex</code>	Makes the integer conversion base 16.
<code>std::oct</code>	Make the integer conversion base 8.

TABLE 3-2 `iostream` Predefined Manipulators (continued)

Predefined Manipulator	Description
<code>ws</code>	Extracts whitespace characters (skips whitespace) until a non-whitespace character is found (which is left in <code>istr</code>).
<code>setbase(n)</code>	Sets the conversion base to <code>n</code> (0, 8, 10, 16 only).
<code>setw(n)</code>	Invokes <code>ios::width(n)</code> . Sets the field width to <code>n</code> .
<code>resetiosflags(i)</code>	Clears the flags bitvector according to the bits set in <code>i</code> .
<code>setiosflags(i)</code>	Sets the flags bitvector according to the bits set in <code>i</code> .
<code>setfill(c)</code>	Sets the fill character (for padding a field) to <code>c</code> .
<code>setprecision(n)</code>	Sets the floating-point precision to <code>n</code> digits.

To use predefined manipulators, you must include the file `iosmanip.h` in your program.

You can define your own manipulators. There are two basic types of manipulator:

- Plain manipulator—Takes an `istream&`, `ostream&`, or `ios&` argument, operates on the stream, and then returns its argument.
- Parameterized manipulator—Takes an `istream&`, `ostream&`, or `ios&` argument, one additional argument (the parameter), operates on the stream, and then returns its stream argument.

Using Plain Manipulators

A plain manipulator is a function that:

- Takes a reference to a stream
- Operates on it in some way
- Returns its argument

The shift operators taking (a pointer to) such a function are predefined for `istream`s, so the function can be put in a sequence of input or output operators.

The shift operator calls the function rather than trying to read or write a value. An example of a tab manipulator that inserts a tab in an ostream is:

```
ostream& tab(ostream& os) {
    return os << '\t' ;
}
...
cout << x << tab << y ;
```

This is an elaborate way to achieve the following:

```
const char tab = "\t";
...
cout << x << tab << y;
```

Here is another example, which cannot be accomplished with a simple constant. Suppose you want to turn whitespace skipping on and off for an input stream. You can use separate calls to `ios::setf` and `ios::unsetf` to turn the `skipws` flag on and off, or you could define two manipulators:

```
#include <iostream.h>
#include <iomanip.h>
istream& skipon(istream &is) {
    is.setf(ios::skipws, ios::skipws);
    return is;
}
istream& skipoff(istream& is) {
    is.unsetf(ios::skipws);
    return is;
}
...
int main ()
{
    int x,y;
    cin >> skipon >> x >> skipoff >> y;
    return 1;
}
```

Parameterized Manipulators

One of the parameterized manipulators that is included in `iomanip.h` is `setfill`. `setfill` sets the character that is used to fill out field widths. It is implemented as shown in the following example::

```
//file setfill.cc
#include<iostream.h>
#include<iomanip.h>

//the private manipulator
```

```

static ios& sfill(ios& i, int f) {
    i.fill(f);
    return i;
}
//the public applicator
smanip_int setfill(int f) {
    return smanip_int(sfill, f);
}

```

A parameterized manipulator is implemented in two parts:

- The *manipulator*. It takes an extra parameter. In the previous code example, it takes an extra `int` parameter. You cannot place this manipulator function in a sequence of input or output operations, since there is no shift operator defined for it. Instead, you must use an auxiliary function, the applicator.
- The *applicator*. It calls the manipulator. The applicator is a global function, and you make a prototype for it available in a header file. Usually the manipulator is a static function in the file containing the source code for the applicator. The manipulator is called only by the applicator, and if you make it static, you keep its name out of the global address space.

Several classes are defined in the header file `iomanip.h`. Each class holds the address of a manipulator function and the value of one parameter. The `iomanip` classes are described in the man page `manip(3C++)`. The previous example uses the `smanip_int` class, which works with an `ios`. Because it works with an `ios`, it also works with an `istream` and an `ostream`. The previous example also uses a second parameter of type `int`.

The applicator creates and returns a class object. In the previous code example the class object is an `smanip_int`, and it contains the manipulator and the `int` argument to the applicator. The `iomanip.h` header file defines the shift operators for this class. When the applicator function `setfill` appears in a sequence of input or output operations, the applicator function is called, and it returns a class. The shift operator acts on the class to call the manipulator function with its parameter value, which is stored in the class.

In the following example, the manipulator `print_hex`:

- Puts the output stream into the hex mode.
- Inserts a `long` value into the stream.
- Restores the conversion mode of the stream.

The class `omanip_long` is used because this code example is for output only, and it operates on a `long` rather than an `int`:

```

#include <iostream.h>
#include <iomanip.h>
static ostream& xfield(ostream& os, long v) {
    long save = os.setf(ios::hex, ios::basefield);
    os << v;
    os.setf(save, ios::basefield);
}

```

```
        return os;
    }
   omanip_long print_hex(long v) {
        returnomanip_long(xfield, v);
    }
}
```

Strstreams: Iostreams for Arrays

See the `strstream(3C++)` man page.

Stdiobufs: Iostreams for stdio files

See the `stdiobuf(3C++)` man page.

Streambufs

`Iostreams` are the formatting part of a two-part (input or output) system. The other part of the system is made up of `streambufs`, which deal in input or output of unformatted streams of characters.

You usually use `streambufs` through `iostreams`, so you don't have to worry about the details of `streambufs`. You can use `streambufs` directly if you choose to, for example, if you need to improve efficiency or to get around the error handling or formatting built into `iostreams`.

Working with Streambufs

A `streambuf` consists of a stream or sequence of characters and one or two pointers into that sequence. Each pointer points between two characters. (Pointers cannot actually point between characters, but it is helpful to think of them that way.) There are two kinds of `streambuf` pointers:

- A *put* pointer, which points just before the position where the next character will be stored

- A *get* pointer, which points just before the next character to be fetched

A `streambuf` can have one or both of these pointers.

Position of Pointers

The positions of the pointers and the contents of the sequences can be manipulated in various ways. Whether or not both pointers move when manipulated depends on the kind of `streambuf` used. Generally, with queue-like `streambufs`, the `get` and `put` pointers move independently; with file-like `streambufs` the `get` and `put` pointers always move together. A `strstream` is an example of a queue-like stream; an `fstream` is an example of a file-like stream.

Using Streambufs

You never create an actual `streambuf` object, but only objects of classes derived from class `streambuf`. Examples are `filebuf` and `strstreambuf`, which are described in man pages `filebuf(3c++)` and `ssbuf(3)`, respectively. Advanced users may want to derive their own classes from `streambuf` to provide an interface to a special device or to provide other than basic buffering. Man pages `sbufpub(3C++)` and `sbufprot(3C++)` discuss how to do this.

Apart from creating your own special kind of `streambuf`, you may want to access the `streambuf` associated with an `iostream` to access the public member functions, as described in the man pages referenced above. In addition, each `iostream` has a defined inserter and extractor which takes a `streambuf` pointer. When a `streambuf` is inserted or extracted, the entire stream is copied.

Here is another way to do the file copy discussed earlier, with the error checking omitted for clarity:

```
ifstream fromFile("thisFile");
ofstream toFile ("thatFile");
toFile << fromFile.rdbuf();
```

We open the input and output files as before. Every `iostream` class has a member function `rdbuf` that returns a pointer to the `streambuf` object associated with it. In the case of an `fstream`, the `streambuf` object is type `filebuf`. The entire file associated with `fromFile` is copied (inserted into) the file associated with `toFile`. The last line could also be written like this:

```
fromFile >> toFile.rdbuf();
```

The source file is then extracted into the destination. The two methods are entirely equivalent.

Iostream Man Pages

A number of C++ man pages give details of the `iostream` library. The following table gives an overview of what is in each man page.

TABLE 3-3 Iostream Man Pages Overview

Man Page	Overview
<code>filebuf</code>	Details the public interface for the class <code>filebuf</code> , which is derived from <code>streambuf</code> and is specialized for use with files. See the <code>sbufpub(3C++)</code> and <code>sbufprot(3C++)</code> man pages for details of features inherited from class <code>streambuf</code> . Use the <code>filebuf</code> class through class <code>fstream</code> .
<code>fstream</code>	Details specialized member functions of classes <code>ifstream</code> , <code>ofstream</code> , and <code>fstream</code> , which are specialized versions of <code>istream</code> , <code>ostream</code> , and <code>iostream</code> for use with files.
<code>ios</code>	Details parts of class <code>ios</code> , which functions as a base class for <code>iostreams</code> . It contains state data common to all streams.
<code>ios.intro</code>	Gives an introduction to and overview of <code>iostreams</code> .
<code>istream</code>	Details the following: <ul style="list-style-type: none">• Member functions for class <code>istream</code>, which supports interpretation of characters fetched from a <code>streambuf</code>• Input formatting• Positioning functions described as part of class <code>ostream</code>.• Some related functions• Related manipulators
<code>manip</code>	Describes the input and output manipulators defined in the <code>iostream</code> library.
<code>ostream</code>	Details the following: <ul style="list-style-type: none">• Member functions for class <code>ostream</code>, which supports interpretation of characters written to a <code>streambuf</code>• Output formatting• Positioning functions described as part of class <code>ostream</code>• Some related functions• Related manipulators

TABLE 3-3 `Iostream` Man Pages Overview (continued)

Man Page	Overview
<code>sbufprot</code>	Describes the interface needed by programmers who are coding a class derived from class <code>streambuf</code> . Also refer to the <code>sbufpub(3C++)</code> man page because some public functions are not discussed in the <code>sbufprot(3C++)</code> man page.
<code>sbufpub</code>	Details the public interface of class <code>streambuf</code> , in particular, the public member functions of <code>streambuf</code> . This man page contains the information needed to manipulate a <code>streambuf</code> -type object directly, or to find out about functions that classes derived from <code>streambuf</code> inherit from it. If you want to derive a class from <code>streambuf</code> , also see the <code>sbufprot(3C++)</code> man page.
<code>ssbuf</code>	Details the specialized public interface of class <code>strstreambuf</code> , which is derived from <code>streambuf</code> and specialized for dealing with arrays of characters. See the <code>sbufpub(3C++)</code> man page for details of features inherited from class <code>streambuf</code> .
<code>stdiobuf</code>	Contains a minimal description of class <code>stdiobuf</code> , which is derived from <code>streambuf</code> and specialized for dealing with <code>stdio</code> FILES. See the <code>sbufpub(3C++)</code> man page for details of features inherited from class <code>streambuf</code> .
<code>strstream</code>	Details the specialized member functions of <code>strstreams</code> , which are implemented by a set of classes derived from the <code>iostream</code> classes and specialized for dealing with arrays of characters.

Iostream Terminology

The `iostream` library descriptions often use terms similar to terms from general programming, but with specialized meanings. The following table defines these terms as they are used in discussing the `iostream` library.

TABLE 3-4 `Iostream` Terminology

<code>Iostream</code> Term	Definition
Buffer	<p>A word with two meanings, one specific to the <code>iostream</code> package and one more generally applied to input and output.</p> <p>When referring specifically to the <code>iostream</code> library, a buffer is an object of the type defined by the class <code>streambuf</code>.</p> <p>A buffer, generally, is a block of memory used to make efficient transfer of characters for input or output. With buffered I/O, the actual transfer of characters is delayed until the buffer is full or forcibly flushed.</p> <p>An unbuffered buffer refers to a <code>streambuf</code> where there is no buffer in the general sense defined above. This chapter avoids use of the term buffer to refer to <code>streambufs</code>. However, the man pages and other C++ documentation do use the term buffer to mean <code>streambufs</code>.</p>
Extraction	The process of taking input from an <code>iostream</code> .
<code>Fstream</code>	An input or output stream specialized for use with files. Refers specifically to a class derived from class <code>iostream</code> when printed in <code>courier</code> font.
Insertion	The process of sending output into an <code>iostream</code> .
<code>Iostream</code>	Generally, an input or output stream.
<code>Iostream</code> library	The library implemented by the include files <code>iostream.h</code> , <code>fstream.h</code> , <code>strstream.h</code> , <code>iomanip.h</code> , and <code>stdiostream.h</code> . Because <code>iostream</code> is an object-oriented library, you should extend it. So, some of what you can do with the <code>iostream</code> library is not implemented.
Stream	An <code>iostream</code> , <code>fstream</code> , <code>strstream</code> , or user-defined stream in general.
<code>Streambuf</code>	A buffer that contains a sequence of characters with a put or get pointer, or both. When printed in <code>courier</code> font, it means the particular class. Otherwise, it refers generally to any object of class <code>streambuf</code> or a class derived from <code>streambuf</code> . Any stream object contains an object, or a pointer to an object, of a type derived from <code>streambuf</code> .
<code>Strstream</code>	An <code>iostream</code> specialized for use with character arrays. It refers to the specific class when printed in <code>courier</code> font.

Using Iostreams in a Multithreaded Environment

This chapter describes how to use the `iostream` classes of the `libc` and `libiostream` libraries for input-output (I/O) in a multithreaded environment. It also provides examples of how to extend functionality of the library by deriving from the `iostream` classes. This chapter is *not* a guide for writing multithreaded code in C++, however.

The discussion here applies only to the old `iostreams` (`libc` and `libiostream`) and does not apply to `libcstd`, the new `iostream` that is part of the C++ Standard Library.

Multithreading

Multi-threading (MT) is a powerful facility that can speed up applications on multiprocessors; it can also simplify the structuring of applications on both multiprocessors and uniprocessors. The `iostream` library has been modified to allow its interfaces to be used by applications in a multithreaded environment by programs that utilize the multithreading capabilities when running Solaris version 2.2, 2.3, 2.4, 2.5, 2.5.1, 2.6, or 7. Applications that utilize the single-threaded capabilities of previous versions of the library are not affected by the behavior of the modified `iostream` interfaces.

A library is defined to be MT-safe if it works correctly in an environment with threads. Generally, this “correctness” means that all of its public functions are reentrant. The `iostream` library provides protection against multiple threads that attempt to modify the state of objects (that is, instances of a C++ class) shared by

more than one thread. However, the scope of MT-safety for an `iostream` object is confined to the period in which the object's public member function is executing.



Caution - An application is *not* automatically guaranteed to be MT-safe because it uses MT-safe objects from the `libc` library. An application is defined to be MT-safe only when it executes as expected in a multithreaded environment.

Organization of the MT-safe `iostream` Library

The organization of the MT-safe `iostream` library is slightly different from other versions of the `iostream` library. The exported interface of the library refers to the public and protected member functions of the `iostream` classes and the set of base classes available, and is consistent with other versions; however, the class hierarchy is different. See “Interface Changes to the `iostream` Library” on page 46 for details.

The original core classes have been renamed with the prefix `unsafe_`. Table 4-1 lists the classes that are the core of the `iostream` package.

TABLE 4-1 Core Classes

Class	Description
<code>stream_MT</code>	The base class for MT-safe classes.
<code>streambuf</code>	The base class for buffers.
<code>unsafe_ios</code>	A class that contains state variables that are common to the various stream classes; for example, error and formatting state.
<code>unsafe_istream</code>	A class that supports formatted and unformatted conversion from sequences of characters retrieved from the <code>streambufs</code> .
<code>unsafe_ostream</code>	A class that supports formatted and unformatted conversion to sequences of characters stored into the <code>streambufs</code> .
<code>unsafe_iostream</code>	A class that combines <code>unsafe_istream</code> and <code>unsafe_ostream</code> classes for bi-directional operations.

Each MT-safe class is derived from the base class `stream_MT`. Each MT-safe class, except `streambuf`, is also derived from the existing `unsafe_` base class. Here are some examples:

```
class streambuf: public stream_MT { ... };
class ios: virtual public unsafe_ios, public stream_MT { ... };
class istream: virtual public ios, public unsafe_istream { ... };
```

The class `stream_MT` provides the mutual exclusion (mutex) locks required to make each `iostream` class MT-safe; it also provides a facility that dynamically enables and disables the locks so that the MT-safe property can be dynamically changed. The basic functionality for I/O conversion and buffer management are organized into the `unsafe_` classes; the MT-safe additions to the library are confined to the derived classes. The MT-safe version of each class contains the same protected and public member functions as the `unsafe_` base class. Each member function in the MT-safe version class acts as a wrapper that locks the object, calls the same function in the `unsafe_` base class, and unlocks the object.

Note - The class `streambuf` is *not* derived from an unsafe class. The public and protected member functions of class `streambuf` are reentrant by locking. Unlocked versions, suffixed with `_unlocked`, are also provided.

Public Conversion Routines

A set of reentrant public functions that are MT-safe have been added to the `iostream` interface. A user-specified buffer is an additional argument to each function. These functions are described as follows.

TABLE 4-2 Reentrant Public Functions

Function	Description
<code>char *oct_r (char *buf, int buflen, long num, int width)</code>	Returns a pointer to the ASCII string that represents the number in octal. A width of nonzero is assumed to be the field width for formatting. The returned value is not guaranteed to point to the beginning of the user-provided buffer.
<code>char *hex_r (char *buf, int buflen, long num, int width)</code>	Returns a pointer to the ASCII string that represents the number in hexadecimal. A width of nonzero is assumed to be the field width for formatting. The returned value is not guaranteed to point to the beginning of the user-provided buffer.

TABLE 4-2 Reentrant Public Functions (continued)

Function	Description
<code>char *dec_r (char *buf, int buflen, long num, int width)</code>	Returns a pointer to the ASCII string that represents the number in decimal. A width of nonzero is assumed to be the field width for formatting. The returned value is not guaranteed to point to the beginning of the user-provided buffer.
<code>char *chr_r (char *buf, int buflen, long num, int width)</code>	Returns a pointer to the ASCII string that contains character <code>chr</code> . If the width is nonzero, the string contains <code>width</code> blanks followed by <code>chr</code> . The returned value is not guaranteed to point to the beginning of the user-provided buffer.
<code>char *form_r (char *buf, int buflen, long num, int width)</code>	Returns a pointer of the string formatted by <code>sprintf</code> , using the format string <code>format</code> and any remaining arguments. The buffer must have sufficient space to contain the formatted string.



Caution - The public conversion routines of the `iostream` library (`oct`, `hex`, `dec`, `chr`, and `form`) that are present to ensure compatibility with an earlier version of `libc` are *not* MT-safe.

Compiling and Linking with the MT-safe `libc` Library

When you build an application that uses the `iostream` classes of the `libc` library to run in a multithreaded environment, compile and link the source code of the application using the `-mt` option. This option passes `-D_REENTRANT` to the preprocessor and `-lthread` to the linker.

Note - Use `-mt` (rather than `-lthread`) to link with `libc` and `libthread`. This option ensures proper linking order of the libraries. Using `-lthread` improperly could cause your application to work incorrectly.

Single-threaded applications that use `iostream` classes do not require special compiler or linker options. By default, the compiler links with the `libc` library.

MT-safe `istream` Restrictions

The restricted definition of MT-safety for the `istream` library means that a number of programming idioms used with `istream` are unsafe in a multithreaded environment using shared `istream` objects.

Checking Error State

To be MT-safe, error checking must occur in a critical region with the I/O operation that causes the error. The following example illustrates how to check for errors:

CODE EXAMPLE 4-1 Checking Error State

```
#include <istream.h>
enumm iostate { IOok, IOeof, IOfail };

iostate read_number(istream& istr, int& num)
{
    stream_locker sl(istr, stream_locker::lock_now);

    istr >> num;

    if (istr.eof()) return IOeof;
    if (istr.fail()) return IOfail;
    return IOok;
}
```

In this example, the constructor of the `stream_locker` object `sl` locks the `istream` object `istr`. The destructor of `sl`, called at the termination of `read_number`, unlocks `istr`.

Obtaining Characters Extracted by Last Unformatted Input Operation

To be MT-safe, the `gcount` function must be called within a thread that has exclusive use of the `istream` object for the period that includes the execution of the last input operation and `gcount` call. The following example shows a call to `gcount`:

CODE EXAMPLE 4-2 Calling `gcount`

```
#include <istream.h>
#include <rlocks.h>
void fetch_line(istream& istr, char* line, int& linecount)
{
    stream_locker sl(istr, stream_locker::lock_defer);

    sl.lock(); // lock the stream istr
    istr >> line;
    linecount = istr.gcount();
    sl.unlock(); // unlock istr
    ...
}
```

```
}
```

In this example, the `lock` and `unlock` member functions of class `stream_locker` define a mutual exclusion region in the program.

User-Defined I/O Operations

To be MT-safe, I/O operations defined for a user-defined type that involve a specific ordering of separate operations must be locked to define a critical region. The following example shows a user-defined I/O operation:

CODE EXAMPLE 4-3 User-Defined I/O Operations

```
#include <rlocks.h>
#include <iostream.h>
class mystream: public istream {

    // other definitions...
    int getRecord(char* name, int& id, float& gpa);
};

int mystream::getRecord(char* name, int& id, float& gpa)
{
    stream_locker sl(this, stream_locker::lock_now);

    *this >> name;
    *this >> id;
    *this >> gpa;

    return this->fail() == 0;
}
```

Performance

Using the MT-safe classes in this version of the `libc` library results in some amount of performance overhead, even in a single-threaded application; however, if you use the `unsafe_` classes of `libc`, this overhead can be avoided.

The scope resolution operator can be used to execute member functions of the base `unsafe_` classes; for example:

```
cout.unsafe_ostream::put("4");

cin.unsafe_istream::read(buf, len);
```

Note - The `unsafe_` classes cannot be safely used in multithreaded applications.

Instead of using `unsafe_` classes, you can make the `cout` and `cin` objects `unsafe` and then use the normal operations. A slight performance deterioration results. The following example shows how to use `unsafe cout` and `cin`:

CODE EXAMPLE 4-4 Disabling `mt-safety`

```
#include <iostream.h>
cout.set_safe_flag(stream_MT::unsafe_object); //disable mt-safety
cin.set_safe_flag(stream_MT::unsafe_object); //disable mt-safety
cout.put('4');
cin.read(buf, len);
```

When an `iostream` object is MT-safe, mutex locking is provided to protect the object's member variables. This locking adds unnecessary overhead to an application that only executes in a single-threaded environment. To improve performance, you can dynamically switch an `iostream` object to and from MT-safety. The following example makes an `iostream` object MT-unsafe:

CODE EXAMPLE 4-5 Switching to MT-unsafe

```
fs.set_safe_flag(stream_MT::unsafe_object); // disable MT-safety
.... do various i/o operations
```

You can safely use an MT-unsafe stream in code where an `iostream` is *not* shared by threads; for example, in a program that has only one thread, or in a program where each `iostream` is private to a thread.

If you explicitly insert synchronization into the program, you can also safely use MT-unsafe `iostreams` in an environment where an `iostream` is shared by threads. The following example illustrates the technique:

CODE EXAMPLE 4-6 Using Synchronization with MT-unsafe Objects

```
generic_lock() ;
fs.set_safe_flag(stream_MT::unsafe_object) ;
... do various i/o operations
generic_unlock() ;
```

where the `generic_lock` and `generic_unlock` functions can be any synchronization mechanism that uses such primitives as mutex, semaphores, or reader/writer locks.

Note - The `stream_locker` class provided by the `libc` library is the preferred mechanism for this purpose.

See “Object Locks” on page 49 for more information.

Interface Changes to the `iostream` Library

This section describes the interface changes made to the `iostream` library to make it MT-safe.

New Classes

The following table lists the new classes added to the `libC` interfaces.

CODE EXAMPLE 4-7 New Classes

```
stream_MT
stream_locker
unsafe_ios
unsafe_istream
unsafe_ostream
unsafe_iostream
unsafe_fstreambase
unsafe_strstreambase
```

New Class Hierarchy

The following table lists the new class hierarchy added to the `iostream` interfaces.

CODE EXAMPLE 4-8 New Class Hierarchy

```
class streambuf : public stream_MT { ... };
class unsafe_ios { ... };
class ios : virtual public unsafe_ios, public stream_MT { ... };
class unsafe_fstreambase : virtual public unsafe_ios { ... };
class fstreambase : virtual public ios, public unsafe_fstreambase
{ ... };class unsafe_strstreambase : virtual public unsafe_ios
{ ... };
class strstreambase : virtual public ios, public unsafe_strstreambase
{ ... };
class unsafe_istream : virtual public unsafe_ios { ... };
class unsafe_ostream : virtual public unsafe_ios { ... };
class istream : virtual public ios, public unsafe_istream { ...
};
class ostream : virtual public ios, public unsafe_ostream { ...
};
class unsafe_iostream : public unsafe_istream, public unsafe_ostream
{ ... };
```

New Functions

The following table lists the new functions added to the `iostream` interfaces.

CODE EXAMPLE 4-9 New Functions

```
class streambuf {
public:
    int sgetc_unlocked();
    void sgetn_unlocked(char *, int);
    int snextc_unlocked();
    int sbumpc_unlocked();
    void stoss_unlocked();
    int in_avail_unlocked();
    int sputbackc_unlocked(char);
    int sputc_unlocked(int);
    int sputn_unlocked(const char *, int);
    int out_waiting_unlocked();
protected:
    char* base_unlocked();
    char* ebuf_unlocked();
    int blen_unlocked();
    char* pbase_unlocked();
    char* eback_unlocked();
    char* gp_ptr_unlocked();
    char* egp_ptr_unlocked();
    char* pp_ptr_unlocked();
    void setp_unlocked(char*, char*);
    void setg_unlocked(char*, char*, char*);
    void pbump_unlocked(int);
    void gbump_unlocked(int);
    void setb_unlocked(char*, char*, int);
    int unbuffered_unlocked();
    char *ep_ptr_unlocked();
    void unbuffered_unlocked(int);
    int allocate_unlocked(int);
};

class filebuf : public streambuf {
public:
    int is_open_unlocked();
    filebuf* close_unlocked();
    filebuf* open_unlocked(const char*, int, int =
        filebuf::openprot);

    filebuf* attach_unlocked(int);
};

class strstreambuf : public streambuf {
public:
    int freeze_unlocked();
    char* str_unlocked();
};

unsafe_ostream& endl(unsafe_ostream&);
unsafe_ostream& ends(unsafe_ostream&);
unsafe_ostream& flush(unsafe_ostream&);
unsafe_istream& ws(unsafe_istream&);
```

```

unsafe_ios& dec(unsafe_ios&);
unsafe_ios& hex(unsafe_ios&);
unsafe_ios& oct(unsafe_ios&);

char* dec_r (char* buf, int buflen, long num, int width)
char* hex_r (char* buf, int buflen, long num, int width)
char* oct_r (char* buf, int buflen, long num, int width)
char* chr_r (char* buf, int buflen, long chr, int width)
char* str_r (char* buf, int buflen, const char* format, int width
            = 0);
char* form_r (char* buf, int buflen, const char* format, ...)

```

Global and Static Data

Global and static data in a multithreaded application are not safely shared among threads. Although threads execute independently, they share access to global and static objects within the process. If one thread modifies such a shared object, all the other threads within the process observe the change, making it difficult to maintain state over time. In C++, class objects (instances of a class) maintain state by the values in their member variables. If a class object is shared, it is vulnerable to changes made by other threads.

When a multithreaded application uses the `iostream` library and includes `iostream.h`, the standard streams—`cout`, `cin`, `cerr`, and `clog`—are, by default, defined as global shared objects. Since the `iostream` library is MT-safe, it protects the state of its shared objects from access or change by another thread while a member function of an `iostream` object is executing. However, the scope of MT-safety for an object is confined to the period in which the object's public member function is executing. For example,

```

int c;
cin.get(c);

```

gets the next character in the `get` buffer and updates the buffer pointer in *ThreadA*. However, if the next instruction in *ThreadA* is another `get` call, the `libc` library does not guarantee to return the next character in the sequence. It is not guaranteed because, for example, *ThreadB* may have also executed the `get` call in the intervening period between the two `get` calls made in *ThreadA*.

See “Object Locks” on page 49” for strategies for dealing with the problems of shared objects and multithreading.

Sequence Execution

Frequently, when `iostream` objects are used, a sequence of I/O operations must be MT-safe. For example, the code:

```
cout << " Error message:" << strerror(err_number) << "\n";
```

involves the execution of three member functions of the `cout` stream object. Since `cout` is a shared object, the sequence must be executed atomically as a critical section to work correctly in a multithreaded environment. To perform a sequence of operations on an `iostream` class object atomically, you must use some form of locking.

The `libc` library now provides the `stream_locker` class for locking operations on an `iostream` object. See “Object Locks” on page 49” for information about the `stream_locker` class.

Object Locks

The simplest strategy for dealing with the problems of shared objects and multithreading is to avoid the issue by ensuring that `iostream` objects are local to a thread. For example,

- Declare objects locally within a thread’s entry function.
- Declare objects in thread-specific data. (For information on how to use thread specific data, see the `thr_keycreate(3T)` man page.)
- Dedicate a stream object to a particular thread. The object thread is private by convention.

However, in many cases, such as default shared standard stream objects, it is not possible to make the objects local to a thread, and an alternative strategy is required.

To perform a sequence of operations on an `iostream` class object atomically, you must use some form of locking. Locking adds some overhead even to a single-threaded application. The decision whether to add locking or make `iostream` objects private to a thread depends on the thread model chosen for the application: Are the threads to be independent or cooperating?

- If each independent thread is to produce or consume data using its own `iostream` object, the `iostream` objects are private to their respective threads and locking is not required.

- If the threads are to cooperate (that is, they are to share the same `iostream` object), then access to the shared object must be synchronized and some form of locking must be used to make sequential operations atomic.

Class `stream_locker`

The `iostream` library provides the `stream_locker` class for locking a series of operations on an `iostream` object. You can, therefore, minimize the performance overhead incurred by dynamically enabling or disabling locking in `iostream` objects.

Objects of class `stream_locker` can be used to make a sequence of operations on a stream object atomic. For example, the code shown in the example below seeks to find a position in a file and reads the next block of data.

CODE EXAMPLE 4-10 Example of Using Locking Operations

```
#include <fstream.h>
#include <rlocks.h>

void lock_example (fstream& fs)
{
    const int len = 128;
    char buf[len];
    int offset = 48;
    stream_locker s_lock(fs, stream_locker::lock_now);
    . . . . // open file
    fs.seekg(offset, ios::beg);
    fs.read(buf, len);
}
```

In this example, the constructor for the `stream_locker` object defines the beginning of a mutual exclusion region in which only one thread can execute at a time. The destructor, called after the return from the function, defines the end of the mutual exclusion region. The `stream_locker` object ensures that both the seek to a particular offset in a file and the read from the file are performed together, atomically, and that *ThreadB* cannot change the file offset before the original *ThreadA* reads the file.

An alternative way to use a `stream_locker` object is to explicitly define the mutual exclusion region. In the following example, to make the I/O operation and subsequent error checking atomic, `lock` and `unlock` member function calls of a `stream_locker` object are used.

CODE EXAMPLE 4-11 Making I/O Operation and Error Checking Atomic

```
{
    ...
    stream_locker file_lck(openfile_stream,
                          stream_locker::lock_defer);
    ....
}
```



```
file_lck.lock(); // lock openfile_stream
openfile_stream << "Value: " << int_value << "\n";
if(!openfile_stream) {
    file_error("Output of value failed\n");
    return;
}
file_lck.unlock(); // unlock openfile_stream
}
```

For more information, see the `stream_locker(3C++)` man page.

MT-safe Classes

You can extend or specialize the functionality of the `ostream` classes by deriving new classes. If objects instantiated from the derived classes will be used in a multithreaded environment, the classes must be MT-safe.

Considerations when deriving MT-safe classes include:

- Making a class object MT-safe by protecting the internal state of the object from multiple-thread modification. To do this, serialize access to member variables in public and protected member functions with mutex locks.
- Making a sequence of calls to member functions of an MT-safe base class atomic, using a `stream_locker` object.
- Avoiding locking overhead by using the `_unlocked` member functions of `streambuf` within critical regions defined by `stream_locker` objects.
- Locking the public virtual functions of class `streambuf` in case the functions are called directly by an application. These functions are: `xsggetn`, `underflow`, `pbackfail`, `xspbtn`, `overflow`, `seekoff`, and `seekpos`.
- Extending the formatting state of an `ios` object by using the member functions `isword` and `pword` in class `ios`. However, a problem can occur if more than one thread is sharing the same index to an `isword` or `pword` function. To make the threads MT-safe, use an appropriate locking scheme.
- Locking member functions that return the value of a member variable greater in size than a `char`.

Object Destruction

Before an `iostream` object that is shared by several threads is deleted, the main thread must verify that the subthreads are finished with the shared object. The following example shows how to safely destroy a shared object.

CODE EXAMPLE 4-12 Destroying a Shared Object

```
#include <fstream.h>
#include <thread.h>
fstream* fp;

void *process_rtn(void*)
{
    // body of sub-threads which uses fp...
}

multi_process(const char* filename, int numthreads)
{
    fp = new fstream(filename, ios::in); // create fstream object
                                        // before creating threads.
    // create threads
    for (int i=0; i<numthreads; i++)
        thr_create(0, STACKSIZE, process_rtn, 0, 0, 0);

    ...
    // wait for threads to finish
    for (int i=0; i<numthreads; i++)
        thr_join(0, 0, 0);

    delete fp; // delete fstream object after
    fp = NULL; // all threads have completed.
}
```

An Example Application

The following code provides an example of a multiply-threaded application that uses `iostream` objects from the `libc` library in an MT-safe way.

The example application creates up to 255 threads. Each thread reads a different input file, one line at a time, and outputs the line to an output file, using the standard output stream, `cout`. The output file, which is shared by all threads, is tagged with a value that indicates which thread performed the output operation.

CODE EXAMPLE 4-13 Using Iostream Objects in a MT-safe Way

```
// create tagged thread data
// the output file is of the form:
// <tag><string of data>\n
// where tag is an integer value in a unsigned char.
// Allows up to 255 threads to be run in this application
// <string of data> is any printable characters
// Because tag is an integer value written
as char,
// you need to use od to look at the output
file, suggest:
// od -c out.file |more

#include <stdlib.h>
#include <stdio.h>
#include <iostream.h>
#include <fstream.h>
#include <thread.h>

struct thread_args {
    char* filename;
    int thread_tag;
};

const int thread_bufsize = 256;

// entry routine for each thread
void* ThreadDuties(void* v) {
    // obtain arguments for this thread
    thread_args* tt = (thread_args*)v;
    char ibuf[thread_bufsize];
    // open thread input file
    ifstream instr(tt->filename);
    stream_locker lockout(cout, stream_locker::lock_defer);
    while(1) {
        // read a line at a time
        instr.getline(ibuf, thread_bufsize - 1, '\n');
        if(instr.eof())
            break;
        // lock cout stream so the i/o operation is atomic
        lockout.lock();
        // tag line and send to cout
        cout << (unsigned char)tt->thread_tag << ibuf << "\n";
        lockout.unlock();
    }
    return 0;
}

int main(int argc, char** argv) {
    // argv: 1+ list of filenames per thread
    if(argc < 2) {
        cout << "\usage: " << argv[0] << " <files..>\n";
        exit(1);
    }
    int num_threads = argc - 1;
    int total_tags = 0;

    // array of thread_ids
    thread_t created_threads[thread_bufsize];
```

```

// array of arguments to thread entry routine
thread_args thr_args[thread_bufsize];
int i;
for( i = 0; i < num_threads; i++) {
    thr_args[i].filename = argv[1 + i];
// assign a tag to a thread - a value less than 256
    thr_args[i].thread_tag = total_tags++;
// create threads
    thr_create(0, 0, ThreadDuties, &thr_args[i],
        THR_SUSPENDED, &created_threads[i]);
}

for(i = 0; i < num_threads; i++) {
    thr_continue(created_threads[i]);
}
for(i = 0; i < num_threads; i++) {
    thr_join(created_threads[i], 0, 0);
}
return 0;
}

```

Associated Man Pages

The manual pages associated with the libraries described in this manual are listed in the following tables.

TABLE A-1 Man Pages for Complex Library

Man Page	Overview
<code>cplx.intro</code>	General introduction to the complex arithmetic library
<code>cartpol</code>	Cartesian and polar functions
<code>cplxerr</code>	Error-handling functions
<code>cplxexp</code>	Exponential, log, and square root functions
<code>cplxops</code>	Arithmetic operator functions
<code>cplxtrig</code>	Trigonometric functions

TABLE A-2 Man Pages for Iostream Library

Man Page	Overview
<code>ios.intro</code>	Gives an introduction to and overview of iostreams
<code>filebuf</code>	Details the public interface for the class <code>filebuf</code> , which is specialized for use with files.
<code>fstream</code>	Details specialized member functions of classes <code>ifstream</code> , <code>ofstream</code> , and <code>fstream</code> , which are specialized for use with files.
<code>ios</code>	Details parts of class <code>ios</code> , which functions as a base class for iostreams
<code>istream</code>	Describes class <code>istream</code> , which supports interpretation of characters fetched from a <code>streambuf</code>
<code>manip</code>	Describes the input and output manipulators defined in the iostream library
<code>ostream</code>	Describes class <code>ostream</code> , which supports interpretation of characters written to a <code>streambuf</code>
<code>sbufprot</code>	Describes the protected interface of class <code>streambuf</code>
<code>sbufpub</code>	Details the public interface of class <code>streambuf</code>
<code>stdiobuf</code>	Describes class <code>stdiobuf</code> , which is specialized for dealing with <code>stdio FILES</code>
<code>strstream</code>	Details the specialized member functions of <code>strstreams</code> , which are specialized for dealing with arrays of characters
<code>ssbuf</code>	Details the specialized public interface of class <code>strstreambuf</code> , which is specialized for dealing with arrays of characters.
<code>stream_MT</code>	Describes the base class that provides dynamic changing of iostream class objects to and from MT safety
<code>stream_locker</code>	Describes the class used for application-level locking of iostream class objects

TABLE A-3 Man Pages for Sun WorkShop Memory Monitor

Man Page	Overview
<code>gcmonitor</code>	Web interface for Sun WorkShop Memory Monitor
<code>gcFixPrematureFrees</code>	Enable and disable fixing of premature frees by the Sun WorkShop Memory Monitor
<code>gcInitialize</code>	Configure Sun WorkShop Memory Monitor at startup

TABLE A-4 Man Pages for C++ Standard Library

Man Page	Overview
<code>Algorithms</code>	Generic algorithms for performing various operations on containers and sequences
<code>Associative_Containers</code>	Ordered containers
<code>Bidirectional_Iterators</code>	An iterator that can both read and write and can traverse a container in both directions
<code>Containers.</code>	A standard template library (STL) collection
<code>Forward_Iterators</code>	A forward-moving iterator that can both read and write
<code>Function_Objects</code>	Object with an <code>operator()</code> defined
<code>Heap_Operations</code>	See entries for <code>make_heap</code> , <code>pop_heap</code> , <code>push_heap</code> and <code>sort_heap</code>
<code>Input_Iterators</code>	A read-only, forward moving iterator
<code>Insert_Iterators</code>	An iterator adaptor that allows an iterator to insert into a container rather than overwrite elements in the container
<code>Iterators</code>	Pointer generalizations for traversal and modification of collections
<code>Negators</code>	Function adaptors and function objects used to reverse the sense of predicate function objects

TABLE A-4 Man Pages for C++ Standard Library *(continued)*

Man Page	Overview
Operators	Operators for the C++ Standard Template Library Output
Output_Iterators	A write-only, forward moving iterator
Predicates	A function or a function object that returns a boolean (true/false) value or an integer value
Random_Access_Iterators	An iterator that reads, writes, and allows random access to a container
Sequences	A container that organizes a set of sequences
Stream_Iterators	Includes iterator capabilities for ostream and istream that allow generic algorithms to be used directly on streams
__distance_type	Determines the type of distance used by an iterator - obsolete
__iterator_category	Determines the category to which an iterator belongs - obsolete
__reverse_bi_iterator	An iterator that traverses a collection backwards
accumulate	Accumulates all elements within a range into a single value
adjacent_difference	Outputs a sequence of the differences between each adjacent pair of elements in a range
adjacent_find	Find the first adjacent pair of elements in a sequence that are equivalent
advance	Moves an iterator forward or backward (if available) by a certain distance
allocator	The default allocator object for storage management in Standard Library containers
auto_ptr	A simple, smart pointer class
back_insert_iterator	An insert iterator used to insert items at the end of a collection

TABLE A-4 Man Pages for C++ Standard Library *(continued)*

Man Page	Overview
<code>back_inserter</code>	An insert iterator used to insert items at the end of a collection
<code>basic_filebuf</code>	Class that associates the input or output sequence with a file
<code>basic_fstream</code>	Supports reading and writing of named files or devices associated with a file descriptor
<code>basic_ifstream</code>	Supports reading from named files or other devices associated with a file descriptor
<code>basic_ios</code>	A base class that includes the common functions required by all streams
<code>basic_iostream</code>	Assists in formatting and interpreting sequences of characters controlled by a stream buffer.
<code>basic_istream</code>	Assists in reading and interpreting input from sequences controlled by a stream buffer
<code>basic_istringstream</code>	Supports reading objects of class <code>basic_string<charT,traits,Allocator></code> from an array in memory
<code>basic_ofstream</code>	Supports writing into named files or other devices associated with a file descriptor
<code>basic_ostream</code>	Assists in formatting and writing output to sequences controlled by a stream buffer
<code>basic_ostringstream</code>	Supports writing objects of class <code>basic_string<charT,traits,Allocator></code>
<code>basic_streambuf</code>	Abstract base class for deriving various stream buffers to facilitate control of character sequences
<code>basic_string</code>	A templated class for handling sequences of character-like entities
<code>basic_stringbuf</code>	Associates the input or output sequence with a sequence of arbitrary characters
<code>basic_stringstream</code>	Supports writing and reading objects of class <code>basic_string<charT,traits,Alocator></code> to/from an array in memory

TABLE A-4 Man Pages for C++ Standard Library *(continued)*

Man Page	Overview
<code>binary_function</code>	Base class for creating binary function objects
<code>binary_negate</code>	A function object that returns the complement of the result of its binary predicate
<code>binary_search</code>	Performs a binary search for a value on a container
<code>bind1st</code>	Templatized utilities to bind values to function objects.
<code>bind2nd</code>	Templatized utilities to bind values to function objects.
<code>binder1st</code>	Templatized utilities to bind values to function objects
<code>binder2nd</code>	Templatized utilities to bind values to function objects
<code>bitset</code>	A template class and related functions for storing and manipulating fixed-size sequences of bits
<code>cerr</code>	Controls output to an unbuffered stream buffer associated with the object <code>stderr</code> declared in <code><cstdio></code>
<code>char_traits</code>	A traits class with types and operations for the <code>basic_string</code> container and <code>iostream</code> classes
<code>cin</code>	Controls input from a stream buffer associated with the object <code>stdin</code> declared in <code><cstdio></code>
<code>clog</code>	Controls output to a stream buffer associated with the object <code>stderr</code> declared in <code><cstdio></code>
<code>codecvt</code>	A code conversion facet
<code>codecvt_byname</code>	A facet that includes code set conversion classification facilities based on the named locales
<code>collate</code>	A string collation, comparison, and hashing facet
<code>collate_byname</code>	A string collation, comparison, and hashing facet
<code>compare</code>	A binary function or a function object that returns true or false
<code>complex</code>	C++ complex number library
<code>copy</code>	Copies a range of elements

TABLE A-4 Man Pages for C++ Standard Library *(continued)*

Man Page	Overview
<code>copy_backward</code>	Copies a range of elements
<code>count</code>	Count the number of elements in a container that satisfy a given condition
<code>count_if</code>	Count the number of elements in a container that satisfy a given condition
<code>cout</code>	Controls output to a stream buffer associated with the object <code>stdout</code> declared in <code><cstdio></code>
<code>ctype</code>	A facet that includes character classification facilities
<code>ctype_byname</code>	A facet that includes character classification facilities based on the named locales
<code>deque</code>	A sequence that supports random access iterators and efficient insertion/deletion at both beginning and end
<code>distance</code>	Computes the distance between two iterators
<code>divides</code>	Returns the result of dividing its first argument by its second
<code>equal</code>	Compares two ranges for equality
<code>equal_range</code>	Finds the largest subrange in a collection into which a given value can be inserted without violating the ordering of the collection
<code>equal_to</code>	A binary function object that returns true if its first argument equals its second
<code>exception</code>	A class that supports logic and runtime errors
<code>facets</code>	A family of classes used to encapsulate categories of locale functionality
<code>filebuf</code>	Class that associates the input or output sequence with a file
<code>fill</code>	Initializes a range with a given value
<code>fill_n</code>	Initializes a range with a given value
<code>find</code>	Finds an occurrence of value in a sequence

TABLE A-4 Man Pages for C++ Standard Library *(continued)*

Man Page	Overview
<code>find_end</code>	Finds the last occurrence of a sub-sequence in a sequence
<code>find_first_of</code>	Finds the first occurrence of any value from one sequence in another sequence
<code>find_if</code>	Finds an occurrence of a value in a sequence that satisfies a specified predicate
<code>for_each</code>	Applies a function to each element in a range
<code>fpos</code>	Maintains position information for the iostream classes
<code>front_insert_iterator</code>	An insert iterator used to insert items at the beginning of a collection
<code>front_inserter</code>	An insert iterator used to insert items at the beginning of a collection
<code>fstream</code>	Supports reading and writing of named files or devices associated with a file descriptor
<code>generate</code>	Initialize a container with values produced by a value-generator class
<code>generate_n</code>	Initialize a container with values produced by a value-generator class
<code>get_temporary_buffer</code>	Pointer based primitive for handling memory
<code>greater</code>	A binary function object that returns true if its first argument is greater than its second
<code>greater_equal</code>	A binary function object that returns true if its first argument is greater than or equal to its second
<code>gslice</code>	A numeric array class used to represent a generalized slice from an array
<code>gslice_array</code>	A numeric array class used to represent a BLAS-like slice from a valarray
<code>has_facet</code>	A function template used to determine if a locale has a given facet

TABLE A-4 Man Pages for C++ Standard Library *(continued)*

Man Page	Overview
<code>ifstream</code>	Supports reading from named files or other devices associated with a file descriptor
<code>includes</code>	A basic set of operation for sorted sequences
<code>indirect_array</code>	A numeric array class used to represent elements selected from a valarray
<code>inner_product</code>	Computes the inner product $A \times B$ of two ranges A and B
<code>inplace_merge</code>	Merges two sorted sequences into one
<code>insert_iterator</code>	An insert iterator used to insert items into a collection rather than overwrite the collection
<code>inserter</code>	An insert iterator used to insert items into a collection rather than overwrite the collection
<code>ios</code>	A base class that includes the common functions required by all streams
<code>ios_base</code>	Defines member types and maintains data for classes that inherit from it
<code>iosfwd</code>	Declares the input/output library template classes and specializes them for wide and tiny characters
<code>isalnum</code>	Determines if a character is alphabetic or numeric
<code>isalpha</code>	Determines if a character is alphabetic
<code>iscntrl</code>	Determines if a character is a control character
<code>isdigit</code>	Determines if a character is a decimal digit
<code>isgraph</code>	Determines if a character is a graphic character
<code>islower</code>	Determines whether a character is lower case
<code>isprint</code>	Determines if a character is printable
<code>ispunct</code>	Determines if a character is punctuation
<code>isspace</code>	Determines if a character is a space

TABLE A-4 Man Pages for C++ Standard Library *(continued)*

Man Page	Overview
<code>istream</code>	Assists in reading and interpreting input from sequences controlled by a stream buffer
<code>istream_iterator</code>	A stream iterator that has iterator capabilities for istreams
<code>istreambuf_iterator</code>	Reads successive characters from the stream buffer for which it was constructed
<code>istringstream</code>	Supports reading objects of class <code>basic_string<charT, traits, Allocator></code> from an array in memory
<code>istrstream</code>	Reads characters from an array in memory
<code>isupper</code>	Determines whether a character is upper case
<code>isxdigit</code>	Determines whether a character is a hexadecimal digit
<code>iter_swap</code>	Exchanges values in two locations
<code>iterator</code>	A base iterator class
<code>iterator_traits</code>	Returns basic information about an iterator
<code>less</code>	A binary function object that returns true if its first argument is less than its second
<code>less_equal</code>	A binary function object that returns true if its first argument is less than or equal to its second
<code>lexicographical_compare</code>	Compares two ranges lexicographically
<code>limits</code>	Refer to the <code>numeric_limits</code> section of this reference guide
<code>list</code>	A sequence that supports bidirectional iterators
<code>locale</code>	A localization class containing a polymorphic set of facets
<code>logical_and</code>	A binary function object that returns true if both of its arguments are true
<code>logical_not</code>	A unary function object that returns true if its argument is false

TABLE A-4 Man Pages for C++ Standard Library *(continued)*

Man Page	Overview
<code>logical_or</code>	A binary function object that returns true if either of its arguments are true
<code>lower_bound</code>	Determines the first valid position for an element in a sorted container
<code>make_heap</code>	Creates a heap
<code>map</code>	An associative container with access to non-key values using unique keys
<code>mask_array</code>	A numeric array class that gives a masked view of a valarray
<code>max</code>	Finds and returns the maximum of a pair of values
<code>max_element</code>	Finds the maximum value in a range
<code>mem_fun</code>	Function objects that adapt a pointer to a member function, to take the place of a global function
<code>mem_fun1</code>	Function objects that adapt a pointer to a member function, to take the place of a global function
<code>mem_fun_ref</code>	Function objects that adapt a pointer to a member function, to take the place of a global function
<code>mem_fun_ref1</code>	Function objects that adapt a pointer to a member function, to take the place of a global function
<code>merge</code>	Merges two sorted sequences into a third sequence
<code>messages</code>	Messaging facets
<code>messages_byname</code>	Messaging facets
<code>min</code>	Finds and returns the minimum of a pair of values
<code>min_element</code>	Finds the minimum value in a range
<code>minus</code>	Returns the result of subtracting its second argument from its first
<code>mismatch</code>	Compares elements from two sequences and returns the first two elements that don't match each other

TABLE A-4 Man Pages for C++ Standard Library *(continued)*

Man Page	Overview
<code>modulus</code>	Returns the remainder obtained by dividing the first argument by the second argument
<code>money_get</code>	Monetary formatting facet for input
<code>money_put</code>	Monetary formatting facet for output
<code>moneypunct</code>	Monetary punctuation facets
<code>moneypunct_byname</code>	Monetary punctuation facets
<code>multimap</code>	An associative container that gives access to non-key values using keys
<code>multiplies</code>	A binary function object that returns the result of multiplying its first and second arguments
<code>multiset</code>	An associative container that allows fast access to stored key values
<code>negate</code>	Unary function object that returns the negation of its argument
<code>next_permutation</code>	Generates successive permutations of a sequence based on an ordering function
<code>not1</code>	A function adaptor used to reverse the sense of a unary predicate function object
<code>not2</code>	A function adaptor used to reverse the sense of a binary predicate function object
<code>not_equal_to</code>	A binary function object that returns true if its first argument is not equal to its second
<code>nth_element</code>	Rearranges a collection so that all elements lower in sorted order than the <i>n</i> th element come before it and all elements higher in sorted order than the <i>n</i> th element come after it
<code>num_get</code>	A numeric formatting facet for input
<code>num_put</code>	A numeric formatting facet for output
<code>numeric_limits</code>	A class for representing information about scalar types

TABLE A-4 Man Pages for C++ Standard Library *(continued)*

Man Page	Overview
<code>numpunct</code>	A numeric punctuation facet
<code>numpunct_byname</code>	A numeric punctuation facet
<code>ofstream</code>	Supports writing into named files or other devices associated with a file descriptor
<code>ostream</code>	Assists in formatting and writing output to sequences controlled by a stream buffer
<code>ostream_iterator</code>	Stream iterators allow for use of iterators with <code>ostreams</code> and <code>istreams</code>
<code>ostreambuf_iterator</code>	Writes successive characters onto the stream buffer object from which it was constructed
<code>ostringstream</code>	Supports writing objects of class <code>basic_string<charT, traits, Allocator></code>
<code>ostrstream</code>	Writes to an array in memory
<code>pair</code>	A template for heterogeneous pairs of values
<code>partial_sort</code>	Templatized algorithm for sorting collections of entities
<code>partial_sort_copy</code>	Templatized algorithm for sorting collections of entities
<code>partial_sum</code>	Calculates successive partial sums of a range of values
<code>partition</code>	Places all of the entities that satisfy the given predicate before all of the entities that do not
<code>permutation</code>	Generates successive permutations of a sequence based on an ordering function
<code>plus</code>	A binary function object that returns the result of adding its first and second arguments
<code>pointer_to_binary_function</code>	A function object that adapts a pointer to a binary function, to take the place of a <code>binary_function</code>
<code>pointer_to_unary_function</code>	A function object class that adapts a pointer to a function, to take the place of a <code>unary_function</code>
<code>pop_heap</code>	Moves the largest element off the heap

TABLE A-4 Man Pages for C++ Standard Library *(continued)*

Man Page	Overview
<code>prev_permutation</code>	Generates successive permutations of a sequence based on an ordering function
<code>priority_queue</code>	A container adapter that behaves like a priority queue
<code>ptr_fun</code>	A function that is overloaded to adapt a pointer to a function, to take the place of a function
<code>push_heap</code>	Places a new element into a heap
<code>queue</code>	A container adaptor that behaves like a queue (first in, first out)
<code>random_shuffle</code>	Randomly shuffles elements of a collection
<code>raw_storage_iterator</code>	Enables iterator-based algorithms to store results into uninitialized memory
<code>remove</code>	Moves desired elements to the front of a container, and returns an iterator that describes where the sequence of desired elements ends
<code>remove_copy</code>	Moves desired elements to the front of a container, and returns an iterator that describes where the sequence of desired elements ends
<code>remove_copy_if</code>	Moves desired elements to the front of a container, and returns an iterator that describes where the sequence of desired elements ends
<code>remove_if</code>	Moves desired elements to the front of a container, and returns an iterator that describes where the sequence of desired elements ends
<code>replace</code>	Substitutes elements in a collection with new values
<code>replace_copy</code>	Substitutes elements in a collection with new values, and moves the revised sequence into result
<code>replace_copy_if</code>	Substitutes elements in a collection with new values, and moves the revised sequence into result
<code>replace_if</code>	Substitutes elements in a collection with new values
<code>return_temporary_buffer</code>	A pointer-based primitive for handling memory

TABLE A-4 Man Pages for C++ Standard Library *(continued)*

Man Page	Overview
<code>reverse</code>	Reverses the order of elements in a collection
<code>reverse_copy</code>	Reverses the order of elements in a collection while copying them to a new collection
<code>reverse_iterator</code>	An iterator that traverses a collection backwards
<code>rotate</code>	Swaps the segment that contains elements from first through middle-1 with the segment that contains the elements from middle through last
<code>rotate_copy</code>	Swaps the segment that contains elements from first through middle-1 with the segment that contains the elements from middle through last
<code>search</code>	Finds a sub-sequence within a sequence of values that is element-wise equal to the values in an indicated range
<code>search_n</code>	Finds a sub-sequence within a sequence of values that is element-wise equal to the values in an indicated range
<code>set</code>	An associative container that supports unique keys
<code>set_difference</code>	A basic set operation for constructing a sorted difference
<code>set_intersection</code>	A basic set operation for constructing a sorted intersection
<code>set_symmetric_difference</code>	A basic set operation for constructing a sorted symmetric difference
<code>set_union</code>	A basic set operation for constructing a sorted union
<code>slice</code>	A numeric array class for representing a BLAS-like slice from an array
<code>slice_array</code>	A numeric array class for representing a BLAS-like slice from a valarray
<code>smanip</code>	Helper classes used to implement parameterized manipulators

TABLE A-4 Man Pages for C++ Standard Library *(continued)*

Man Page	Overview
<code>smanip_fill</code>	Helper classes used to implement parameterized manipulators
<code>sort</code>	A templated algorithm for sorting collections of entities
<code>sort_heap</code>	Converts a heap into a sorted collection
<code>stable_partition</code>	Places all of the entities that satisfy the given predicate before all of the entities that do not, while maintaining the relative order of elements in each group
<code>stable_sort</code>	A templated algorithm for sorting collections of entities
<code>stack</code>	A container adapter that behaves like a stack (last in, first out)
<code>streambuf</code>	Abstract base class for deriving various stream buffers to facilitate control of character sequences
<code>string</code>	A typedef for <code>basic_string<char, char_traits<char>, allocator<char>></code>
<code>stringbuf</code>	Associates the input or output sequence with a sequence of arbitrary characters
<code>stringstream</code>	Supports writing and reading objects of class <code>basic_string<charT, traits, Allocator></code> to/from an array in memory
<code>strstream</code>	Reads and writes to an array in memory
<code>strstreambuf</code>	Associates either the input sequence or the output sequence with a tiny character array whose elements store arbitrary values
<code>swap</code>	Exchanges values
<code>swap_ranges</code>	Exchanges a range of values in one location with those in another
<code>time_get</code>	A time formatting facet for input
<code>time_get_byname</code>	A time formatting facet for input, based on the named locales

TABLE A-4 Man Pages for C++ Standard Library *(continued)*

Man Page	Overview
<code>time_put</code>	A time formatting facet for output
<code>time_put_byname</code>	A time formatting facet for output, based on the named locales
<code>tolower</code>	Converts a character to lower case.
<code>toupper</code>	Converts a character to upper case
<code>transform</code>	Applies an operation to a range of values in a collection and stores the result
<code>unary_function</code>	A base class for creating unary function objects
<code>unary_negate</code>	A function object that returns the complement of the result of its unary predicate
<code>uninitialized_copy</code>	An algorithm that uses construct to copy values from one range to another location
<code>uninitialized_fill</code>	An algorithm that uses the construct algorithm for setting values in a collection
<code>uninitialized_fill_n</code>	An algorithm that uses the construct algorithm for setting values in a collection
<code>unique</code>	Removes consecutive duplicates from a range of values and places the resulting unique values into the result
<code>unique_copy</code>	Removes consecutive duplicates from a range of values and places the resulting unique values into the result
<code>upper_bound</code>	Determines the last valid position for a value in a sorted container
<code>use_facet</code>	A template function used to obtain a facet
<code>valarray</code>	An optimized array class for numeric operations
<code>vector</code>	A sequence that supports random access iterators
<code>wcerr</code>	Controls output to an unbuffered stream buffer associated with the object <code>stderr</code> declared in <code><cstdio></code>
<code>wcin</code>	Controls input from a stream buffer associated with the object <code>stdin</code> declared in <code><cstdio></code>

TABLE A-4 Man Pages for C++ Standard Library *(continued)*

Man Page	Overview
wclog	Controls output to a stream buffer associated with the object stderr declared in <cstdio>
wcout	Controls output to a stream buffer associated with the object stdout declared in <cstdio>
wfilebuf	Class that associates the input or output sequence with a file
wfstream	Supports reading and writing of named files or devices associated with a file descriptor
wifstream	Supports reading from named files or other devices associated with a file descriptor
wios	A base class that includes the common functions required by all streams
wistream	Assists in reading and interpreting input from sequences controlled by a stream buffer
wistringstream	Supports reading objects of class <code>basic_string<charT, traits, Allocator></code> from an array in memory
wofstream	Supports writing into named files or other devices associated with a file descriptor
wostream	Assists in formatting and writing output to sequences controlled by a stream buffer
wostreambuf	Supports writing objects of class <code>basic_string<charT, traits, Allocator></code>
wstreambuf	Abstract base class for deriving various stream buffers to facilitate control of character sequences
wstring	A typedef for <code>basic_string<wchar_t, char_traits<wchar_t>, allocator<wchar_t>></code>
wstringbuf	Associates the input or output sequence with a sequence of arbitrary characters

Index
