

Reference manual

to the GNU C++ Iostream Facility

(Very much a work in progress)

Per Bothner (bothner@cygnus.com)

1 Introduction

The `iostream` library was written by Per Bothner.

Various people have found bugs or come with suggestions. Hongjiu Lu has worked hard to use the library as the default `stdio` implementation for Linux, and has provided much stress-testing of the library.

Some code was derived from parts of BSD 4.4, which is copyright University of California at Berkeley.

2 Using the iostream layer

2.1 C-style formatting for streams

These methods all return `*this`.

`ostream& ostream::vform(const char *format, ...)` [Method]

Similar to `fprintf(file, format, ...)`. The *format* is a `printf`-style format control string, which is used to format the (variable number of) arguments, printing the result on the `this` stream.

`ostream& ostream::vform(const char *format, va_list args)` [Method]

Similar to `vfprintf(file, format, args)`. The *format* is a `printf`-style format control string, which is used to format the argument list *args*, printing the result on the `this` stream.

`istream& istream::scan(const char *format, ...)` [Method]

Similar to `fscanf(file, format, ...)`. The *format* is a `scanf`-style format control string, which is used to read the (variable number of) arguments from the `this` stream.

`istream& istream::vscan(const char *format, va_list args)` [Method]

Like `istream::scan`, but takes a single `va_list` argument.

3 Using the streambuf layer

3.1 C-style formatting for streambufs

The GNU streambuf class supports printf-like formatting.

```
int streambuf::vform(const char *format, ...) [Method]
    Similar to fprintf(file, format, ...). The format is a printf-style format control string, which is used to format the (variable number of) arguments, printing the result on the this streambuf. The result is the number of characters printed.
```

```
int streambuf::vform(const char *format, va_list args) [Method]
    Similar to vfprintf(file, format, args). The format is a printf-style format control string, which is used to format the argument list args, printing the result on the this streambuf. The result is the number of characters printed.
```

```
int streambuf::scan(const char *format, ...) [Method]
    Similar to fscanf(file, format, ...). The format is a scanf-style format control string, which is used to read the (variable number of) arguments from the this streambuf. The result is the number of items assigned, or EOF in case of input failure before any conversion.
```

```
int streambuf::vscan(const char *format, va_list args) [Method]
    Like streambuf::scan, but takes a single va_list argument.
```

3.2 stdiobuf

A *stdiobuf* is a *streambuf* object that points to a FILE object (as defined by `stdio.h`). All *streambuf* operations on the *stdiobuf* are forwarded to the FILE. Thus the *stdiobuf* object provides a wrapper around a FILE, allowing use of *streambuf* operations on a FILE. This can be useful when mixing C code with C++ code.

The pre-defined streams `cin`, `cout`, and `cerr` are normally implemented as *stdiobufs* that point to respectively `stdin`, `stdout`, and `stderr`. This is convenient, but it does cost some extra overhead. (If you have sets things up so that you use the implementation of `stdio` provided with this library, then `cin`, `cout`, and `cerr` will be set up to use *stdiobufs*, since you get their benefits for free.)

Note that if you use `setbuf` to give a buffer to a *stdiobuf*, that buffer will provide intermediate buffering in addition that whatever is done by the FILE.

3.3 indirectbuf

An *indirectbuf* is one that forwards all of its I/O requests to another *streambuf*. All get-related requests are sent to `get_stream()`. All put-related requests are sent to `put_stream()`.

An *indirectbuf* can be used to implement Common Lisp synonym-streams and two-way-streams:

```
class synonymbuf : public indirectbuf {
    Symbol *sym;
```

```

    synonymbuf(Symbol *s) { sym = s; }
    virtual streambuf *lookup_stream(int mode) {
        return coerce_to_streambuf(lookup_value(sym)); }
};

```

3.4 Backing up

The GNU iostream library allows you to ask streambuf to remember the current position, and then later after you've read further be able to go back to it. You're guaranteed to be able to backup arbitrary amounts, even on unbuffered files or multiple buffers worth, as long as you tell the library advance. This unbounded backup is very useful for scanning and parsing applications. This example shows a typical scenario:

```

// Read either "dog", "hound", or "hounddog".
// If "dog" is found, return 1.
// If "hound" is found, return 2.
// If "hounddog" is found, return 3.
// If non of these are found, return -1.
int my_scan(streambuf* sb)
{
    streammarker fence(sb);
    char buffer[20];
    // Try reading "hounddog":
    if (sb->sgetn(buffer, 8) == 8 && strncmp(buffer, "hounddog", 8) == 0)
        return 3;
    // No, no "hounddog": Backup to 'fence' ...
    sb->seekmark(fence); //
    // ... and try reading "dog":
    if (sb->sgetn(buffer, 3) == 3 && strncmp(buffer, "dog", 3) == 0)
        return 1;
    // No, no "dog" either: Backup to 'fence' ...
    sb->seekmark(fence); //
    // ... and try reading "hound":
    if (sb->sgetn(buffer, 5) == 5 && strncmp(buffer, "hound", 5) == 0)
        return 2;
    // No, no "hound" either: Backup to 'fence' and signal failure.
    sb->seekmark(fence); // Backup to 'fence'..
    return -1;
}

```

```

streammarker::streammarker (streambuf* sbuf) [Constructor]
    Create a streammarker associated with sbuf that remembers the current position of
    the get pointer.

```

```

int streammarker::delta (streammarker& mark2) [Method]
    Return the difference between the get positions corresponding to *this and mark2
    (which must point into the same streambuffer as this).

```

```
int streammarker::delta () [Method]
    Return the position relative to the streambuffer's current get position.

int streambuffer::seekmark (streammarker& mark) [Method]
    Move the get pointer to where it (logically) was when mark was constructed.
```

4 stdio: C input/output

Iostreams is distributed with a complete implementation of the ANSI C stdio facility. It is implemented using streambufs.

The stdio package is intended as a replacement for the whatever stdio is in your C library. It can co-exist with C libraries that have alternate implementations of stdio, but there may be some problems. Since stdio works best when you build libc to contain it, and that may be inconvenient, it is not installed by default.

Extensions beyond ANSI:

- A stdio FILE is identical to a streambuf. Hence there is no need to worry about synchronizing C and C++ input/output - they are by definition always synchronized.
- If you create a new streambuf sub-class (in C++), you can use it as a FILE from C. Thus the system is extensible using the standard streambuf protocol.
- You can arbitrarily mix reading and writing, without having to seek in between.
- Unbounded ungetc() buffer.

5 Streambuf internals

5.1 Buffer management

5.1.1 Areas

Streambuf buffer management is fairly sophisticated (this is a nice way to say "complicated"). The standard protocol has the following "areas":

- The *put area* contains characters waiting for output.
- The *get area* contains characters available for reading.
- The *reserve area* is available to virtual methods. Usually, the get and/or put areas are part of the reserve area.

The GNU streambuf design extends this by supporting two get areas:

- The *main get area* contains characters that have been read in from the character source, but not yet read by the application.
- The *backup area* contains previously read data that is being saved because of a user request, or that have been "unread" (putback).

The backup and the main get area are logically contiguous: That is, the first character of the main get area follows the last character of the backup area.

The *current get area* is whichever one of the backup or main get areas that is currently being read from. The other of the two is the *non-current get area*.

5.1.2 Pointers

The following `char*` pointers define the various areas. (Note that if a pointer points to the 'end' of an area, it means that it points to the character after the area.)

<code>char* streambuffer::base ()</code>	[Method]
The start of the reserve area.	
<code>char* streambuffer::ebuf ()</code>	[Method]
The end of the reserve area.	
<code>char* streambuffer::pbase ()</code>	[Method]
The start of the put area.	
<code>char* streambuffer::pptr ()</code>	[Method]
The current put position. If <code>pptr() < epptr()</code> , then the next write will overwrite <code>*pptr()</code> , and increment <code>pptr()</code> .	
<code>char* streambuffer::epptr ()</code>	[Method]
The end of the put area.	
<code>char* streambuffer::eback ()</code>	[Method]
The start of the current get area.	

<code>char* streambuffer::gptr ()</code>	[Method]
The current get position. If <code>gptr() < egptr()</code> , then the next read will read <code>*gptr()</code> , and increment <code>gptr()</code> .	
<code>char* streambuffer::egptr ()</code>	[Method]
The end of the current get area.	
<code>char* streambuffer::Gbase ()</code>	[Method]
The start of the main get area.	
<code>char* streambuffer::eGptr ()</code>	[Method]
The end of the main get area.	
<code>char* streambuffer::Bbase ()</code>	[Method]
The start of the backup area.	
<code>char* streambuffer::Bptr ()</code>	[Method]
The start of the used part of the backup area. The area (<code>Bptr() .. eBptr()</code>) contains data that has been pushed back, while (<code>Bbase() .. eBptr()</code>) contains unused space available for future putbacks.	
<code>char* streambuffer::eBptr ()</code>	[Method]
The end of the backup area.	
<code>char* streambuffer::Nbase ()</code>	[Method]
The start of the non-current get area (either <code>main_gbase</code> or <code>backup_gbase</code>).	
<code>char* streambuffer::eNptr ()</code>	[Method]
The end of the non-current get area.	

5.2 Filebuf internals

The `filebuf` is used a lot, so it is important that it be efficient. It also supports rather complex semantics. So let us examine its implementation.

5.2.1 Tied read and write pointers

The `streambuf` model allows completely independent read and write pointers. However, a `filebuf` has only a single logical pointer used for both reads and writes. Since the `streambuf` protocol uses `gptr()` for reading and `pptr()` for writing, we map the logical file pointer into either `gptr()` or `pptr()` at different times.

- Reading is allowed when `gptr() < egptr()`, which we call get mode.
- Writing is allowed when `pptr() < epptr()`, which we call put mode.

A `filebuf` cannot be in get mode and put mode at the same time.

We have up to two buffers:

- The backup area, defined by `Bbase()`, `Bptr()`, and `eBptr()`. This can be empty.
- The reserve area, which also contains the main get area. For an unbuffered file, the (`shortbuf()..shortbuf()+1`) is used, where `shortbuf()` points to a 1-byte buffer that is part of the `filebuf`.

The file system's idea of the current position is `eGptr()`.

Character that have been written into a buffer but not yet written out (flushed) to the file systems are those between `pbase()` and `pptr()`.

The end of the valid data bytes is: `pptr() > eGptr() && pptr() < ebuf() ? pptr() : eGptr()`.

If the `filebuf` is unbuffered or line buffered, the `eptr()` is `pbase()`. This forces a call to `overflow()` on each put of a character. The logical `epptr()` is `epptr() ? ebuf() : NULL`. (If the buffer is read-only, set `pbase()`, `pptr()`, and `epptr()` to `NULL`. NOT!)

Appendix A Function and Variable Index, Concept Index,, Top

I

istream::scan(const..... 2
istream::vscan(const..... 2

O

ostream::vform(const..... 2

S

streambuf::scan(const 3
streambuf::vform(const 3
streambuf::vscan(const 3
streambuffer::base..... 7
streambuffer::Bbase..... 8
streambuffer::Bptr..... 8

streambuffer::eback..... 7
streambuffer::eBptr..... 8
streambuffer::ebuf 7
streambuffer::egptr..... 8
streambuffer::eGptr..... 8
streambuffer::eNptr..... 8
streambuffer::epptr..... 7
streambuffer::Gbase..... 8
streambuffer::gptr 8
streambuffer::Nbase..... 8
streambuffer::pbase..... 7
streambuffer::pptr 7
streambuffer::seekmark 5
streammarker::delta 4, 5
streammarker::streammarker..... 4

Appendix B Concept Index, Function and Variable Index, Top

B

backup area..... 7

G

get area..... 7

M

main get area..... 7

P

put area..... 7

R

reserve area..... 7

Table of Contents

1	Introduction	1
2	Using the iostream layer	2
2.1	C-style formatting for streams	2
3	Using the streambuf layer	3
3.1	C-style formatting for streambufs	3
3.2	stdiobuf	3
3.3	indirectbuf	3
3.4	Backing up	4
4	stdio: C input/output	6
5	Streambuf internals	7
5.1	Buffer management	7
5.1.1	Areas	7
5.1.2	Pointers	7
5.2	Filebuf internals	8
5.2.1	Tied read and write pointers	8
Appendix A Function and Variable		
	Index, Concept Index, Top	10
Appendix B Concept Index, Function and		
	Variable Index, Top	11