



28.5 — Stream states and input validation

ALEX¹

SEPTEMBER 9, 2024

Stream states

The `ios_base` class contains several state flags that are used to signal various conditions that may occur when using streams:

Flag	Meaning
<code>goodbit</code>	Everything is okay
<code>badbit</code>	Some kind of fatal error occurred (e.g. the program tried to read past the end of a file)
<code>eofbit</code>	The stream has reached the end of a file
<code>failbit</code>	A non-fatal error occurred (e.g. the user entered letters when the program was expecting an integer)

Although these flags live in `ios_base`, because `ios` is derived from `ios_base` and `ios` takes less typing than `ios_base`, they are generally accessed through `ios` (e.g. as `std::ios::failbit`).

`ios` also provides a number of member functions in order to conveniently access these states:

Member function	Meaning
good()	Returns true if the goodbit is set (the stream is ok)
bad()	Returns true if the badbit is set (a fatal error occurred)
eof()	Returns true if the eofbit is set (the stream is at the end of a file)
fail()	Returns true if the failbit is set (a non-fatal error occurred)
clear()	Clears all flags and restores the stream to the goodbit state
clear(state)	Clears all flags and sets the state flag passed in
rdstate()	Returns the currently set flags
setstate(state)	Sets the state flag passed in

The most commonly dealt with bit is the failbit, which is set when the user enters invalid input. For example, consider the following program:

```
1 | std::cout << "Enter your age: ";  
2 | int age {};  
3 | std::cin >> age;
```

Note that this program is expecting the user to enter an integer. However, if the user enters non-numeric data, such as "Alex", cin will be unable to extract anything to age, and the failbit will be set.



If an error occurs and a stream is set to anything other than goodbit, further stream operations on that stream will be ignored. This condition can be cleared by calling the clear() function.

Input validation

Input validation is the process of checking whether the user input meets some set of criteria. Input validation can generally be broken down into two types: string and numeric.

With string validation, we accept all user input as a string, and then accept or reject that string depending on whether it is formatted appropriately. For example, if we ask the user to enter a telephone number, we may want to ensure the data they enter has ten digits. In most languages (especially scripting languages like Perl and PHP), this is done via regular expressions. The C++ standard library has a [regular expression library](https://en.cppreference.com/w/cpp/regex) (<https://en.cppreference.com/w/cpp/regex>)³ as well. Because regular expressions are slow compared to manual string validation, they should only be used if performance (compile-time and run-time) is of no concern or manual validation is too cumbersome.

With numerical validation, we are typically concerned with making sure the number the user enters is within a particular range (e.g. between 0 and 20). However, unlike with string validation, it's possible for the user to enter things that aren't numbers at all -- and we need to handle these cases too.

To help us out, C++ provides a number of useful functions that we can use to determine whether specific characters are numbers or letters. The following functions live in the `cctype` header:



Function	Meaning
<code>std::isalnum(int)</code>	Returns non-zero if the parameter is a letter or a digit
<code>std::isalpha(int)</code>	Returns non-zero if the parameter is a letter
<code>std::iscntrl(int)</code>	Returns non-zero if the parameter is a control character
<code>std::isdigit(int)</code>	Returns non-zero if the parameter is a digit
<code>std::isgraph(int)</code>	Returns non-zero if the parameter is printable character that is not whitespace
<code>std::isprint(int)</code>	Returns non-zero if the parameter is printable character (including whitespace)
<code>std::ispunct(int)</code>	Returns non-zero if the parameter is neither alphanumeric nor whitespace
<code>std::isspace(int)</code>	Returns non-zero if the parameter is whitespace
<code>std::isxdigit(int)</code>	Returns non-zero if the parameter is a hexadecimal digit (0-9, a-f, A-F)

String validation

Let's do a simple case of string validation by asking the user to enter their name. Our validation criteria will be that the user enters only alphabetic characters or spaces. If anything else is encountered, the input will be rejected.

When it comes to variable length inputs, the best way to validate strings (besides using a regular expression library) is to step through each character of the string and ensure it meets the validation criteria. That's exactly what we'll do here, or better, that's what `std::all_of` does for us.

```

1  #include <algorithm> // std::all_of
2  #include <cctype> // std::isalpha, std::isspace
3  #include <iostream>
4  #include <ranges>
5  #include <string>
6  #include <string_view>
7
8  bool isValidName(std::string_view name)
9  {
10     return std::ranges::all_of(name, [](char ch) {
11         return std::isalpha(ch) || std::isspace(ch);
12     });
13
14     // Before C++20, without ranges
15     // return std::all_of(name.begin(), name.end(), [](char ch) {
16     //     return std::isalpha(ch) || std::isspace(ch);
17     // });
18 }
19
20 int main()
21 {
22     std::string name{};
23
24     do
25     {
26         std::cout << "Enter your name: ";
27         std::getline(std::cin, name); // get the entire line, including spaces
28     } while (!isValidName(name));
29
30     std::cout << "Hello " << name << "!\n";
31 }
```

Note that this code isn't perfect: the user could say their name was "asf w jweo s di we ao" or some other bit of gibberish, or even worse, just a bunch of spaces. We could address this somewhat by refining our validation criteria to only accept strings that contain at least one character and at most one space.

Author's note

Reader "Waldo" provides a C++20 solution (using `std::ranges`) that addresses these shortcomings [here \(#comment-571052\)](#)⁴

Now let's take a look at another example where we are going to ask the user to enter their phone number. Unlike a user's name, which is variable-length and where the validation criteria

are the same for every character, a phone number is a fixed length but the validation criteria differ depending on the position of the character. Consequently, we are going to take a different approach to validating our phone number input. In this case, we're going to write a function that will check the user's input against a predetermined template to see whether it matches. The template will work as follows:

A # will match any digit in the user input.

A @ will match any alphabetic character in the user input.

A _ will match any whitespace.

A ? will match anything.

Otherwise, the characters in the user input and the template must match exactly.



So, if we ask the function to match the template "(###) ###-####", that means we expect the user to enter a '(' character, three numbers, a ')' character, a space, three numbers, a dash, and four more numbers. If any of these things doesn't match, the input will be rejected.

Here is the code:

```

1  #include <algorithm> // std::equal
2  #include <cctype> // std::isdigit, std::isspace, std::isalpha
3  #include <iostream>
4  #include <map>
5  #include <ranges>
6  #include <string>
7  #include <string_view>
8
9  bool inputMatches(std::string_view input, std::string_view pattern)
10 {
11     if (input.length() != pattern.length())
12     {
13         return false;
14     }
15
16     // This table defines all special symbols that can match a range of user
17 input
18 // Each symbol is mapped to a function that determines whether the input
19 is valid for that symbol
20 static const std::map<char, int (*)(int)> validators{
21     { '#', &std::isdigit },
22     { '_', &std::isspace },
23     { '@', &std::isalpha },
24     { '?', [](int) { return 1; } }
25 };
26
27 // Before C++20, use
28 // return std::equal(input.begin(), input.end(), pattern.begin(), []
29 (char ch, char mask) -> bool {
30 // ...
31
32     return std::ranges::equal(input, pattern, [](char ch, char mask) -> bool
33 {
34     auto found{ validators.find(mask) };
35
36     if (found != validators.end())
37     {
38         // The pattern's current element was found in the validators.
39 Call the
40 // corresponding function.
41     return (*found->second)(ch);
42     }
43
44     // The pattern's current element was not found in the validators.
45 The
46 // characters have to be an exact match.
47     return ch == mask;
48 }); // end of lambda
49 }
50
51 int main()
52 {
53     std::string phoneNumber{};
54
55     do
56     {
57         std::cout << "Enter a phone number (###) ###-####: ";
58         std::getline(std::cin, phoneNumber);
59     } while (!inputMatches(phoneNumber, "(###) ###-####"));
60
61     std::cout << "You entered: " << phoneNumber << '\n';
62 }

```

Using this function, we can force the user to match our specific format exactly. However, this function is still subject to several constraints: if #, @, _, and ? are valid characters in the user input, this function won't work, because those symbols have been given special meanings. Also, unlike with regular expressions, there is no template symbol that means "a variable number of characters can be entered". Thus, such a template could not be used to ensure the user enters two words separated by a whitespace, because it can not handle the fact that the words are of variable lengths. For such problems, the non-template approach is generally more appropriate.

Numeric validation

When dealing with numeric input, the obvious way to proceed is to use the extraction operator to extract input to a numeric type. By checking the failbit, we can then tell whether the user entered a number or not.

Let's try this approach:

```
1  #include <iostream>
2  #include <limits>
3
4  int main()
5  {
6      int age{};
7
8      while (true)
9      {
10         std::cout << "Enter your age: ";
11         std::cin >> age;
12
13         if (std::cin.fail()) // no extraction took place
14         {
15             std::cin.clear(); // reset the state bits back to goodbit so we
16             can use ignore()
17             std::cin.ignore(std::numeric_limits<std::streamsize>::max(),
18             '\n'); // clear out the bad input from the stream
19             continue; // try again
20         }
21
22         if (age <= 0) // make sure age is positive
23             continue;
24
25         break;
26     }
27     std::cout << "You entered: " << age << '\n';
}
```

If the user enters an integer, the extraction will succeed. `std::cin.fail()` will evaluate to false, skipping the conditional, and (assuming the user entered a positive number), we will hit the break statement, exiting the loop.

If the user instead enters input starting with a letter, the extraction will fail. `std::cin.fail()` will evaluate to true, and we will go into the conditional. At the end of the conditional block, we will hit the continue statement, which will jump back to the top of the while loop, and the user will be asked to enter input again.



However, there's one more case we haven't tested for, and that's when the user enters a string that starts with numbers but then contains letters (e.g. "34abcd56"). In this case, the starting numbers (34) will be extracted into age, the remainder of the string ("abcd56") will be left in the input stream, and the failbit will NOT be set. This causes two potential problems:

1. If you want this to be valid input, you now have garbage in your stream.
2. If you don't want this to be valid input, it is not rejected (and you have garbage in your stream).

Let's fix the first problem. This is easy:

```
1  #include <iostream>
2  #include <limits>
3
4  int main()
5  {
6      int age{};
7
8      while (true)
9      {
10         std::cout << "Enter your age: ";
11         std::cin >> age;
12
13         if (std::cin.fail()) // no extraction took place
14         {
15             std::cin.clear(); // reset the state bits back to goodbit so we
16             can use ignore()
17             std::cin.ignore(std::numeric_limits<std::streamsize>::max(),
18             '\n'); // clear out the bad input from the stream
19             continue; // try again
20         }
21
22         std::cin.ignore(std::numeric_limits<std::streamsize>::max(), '\n');
23         // clear out any additional input from the stream
24
25         if (age <= 0) // make sure age is positive
26             continue;
27
28         break;
29     }
30
31     std::cout << "You entered: " << age << '\n';
32 }
```

If you don't want such input to be valid, we'll have to do a little extra work. Fortunately, the previous solution gets us half way there. We can use the `gcount()` function to determine how many characters were ignored. If our input was valid, `gcount()` should return 1 (the newline character that was discarded). If it returns more than 1, the user entered something that wasn't extracted properly, and we should ask them for new input. Here's an example of this:


```
1  #include <iostream>
2  #include <limits>
3
4  int main()
5  {
6      int age{};
7
8      while (true)
9      {
10         std::cout << "Enter your age: ";
11         std::cin >> age;
12
13         if (std::cin.fail()) // no extraction took place
14             {
15                 std::cin.clear(); // reset the state bits back to goodbit so we
16 can use ignore()
17                 std::cin.ignore(std::numeric_limits<std::streamsize>::max(),
18 '\n'); // clear out the bad input from the stream
19                 continue; // try again
20             }
21
22         std::cin.ignore(std::numeric_limits<std::streamsize>::max(), '\n');
23 // clear out any additional input from the stream
24         if (std::cin.gcount() > 1) // if we cleared out more than one
25 additional character
26             {
27                 continue; // we'll consider this input to be invalid
28             }
29
30         if (age <= 0) // make sure age is positive
31             {
32                 continue;
33             }
34
35         break;
36     }
37
38     std::cout << "You entered: " << age << '\n';
39 }
```

Numeric validation as a string

The above example was quite a bit of work simply to get a simple value! Another way to process numeric input is to read it in as a string, then try to convert it to a numeric type. The following program makes use of that methodology:

```
1 #include <charconv> // std::from_chars
2 #include <iostream>
3 #include <limits>
4 #include <optional>
5 #include <string>
6 #include <string_view>
7
8 // std::optional<int> returns either an int or nothing
9 std::optional<int> extractAge(std::string_view age)
10 {
11     int result{};
12     const auto end{ age.data() + age.length() }; // get end iterator of
13     underlying C-style string
14
15     // Try to parse an int from age
16     // If we got an error of some kind...
17     if (std::from_chars(age.data(), end, result).ec != std::errc{})
18     {
19         return {}; // return nothing
20     }
21
22     if (result <= 0) // make sure age is positive
23     {
24         return {}; // return nothing
25     }
26
27     return result; // return an int value
28 }
29
30 int main()
31 {
32     int age{};
33
34     while (true)
35     {
36         std::cout << "Enter your age: ";
37         std::string strAge{};
38
39         // Try to get a line of input
40         if (!std::getline(std::cin >> std::ws, strAge))
41         {
42             // If we failed, clean up and try again
43             std::cin.clear();
44             std::cin.ignore(std::numeric_limits<std::streamsize>::max(),
45 '\n');
46             continue;
47         }
48
49         // Try to extract the age
50         auto extracted{ extractAge(strAge) };
51
52         // If we failed, try again
53         if (!extracted)
54             continue;
55
56         age = *extracted; // get the value
57         break;
58     }
59
60     std::cout << "You entered: " << age << '\n';
61 }
```

Whether this approach is more or less work than straight numeric extraction depends on your validation parameters and restrictions.

As you can see, doing input validation in C++ is a lot of work. Fortunately, many such tasks (e.g. doing numeric validation as a string) can be easily turned into functions that can be reused in a wide variety of situations.



[Next lesson](#)

28.6 [Basic file I/O](#)

5



[Back to table of contents](#)

6



[Previous lesson](#)

28.4 [Stream classes for strings](#)

7





8


[B](#) [U](#) [URL](#) [INLINE CODE](#) [C++ CODE BLOCK](#) [HELP!](#)

Leave a comment...

Notify me about replies:



 Find a mistake? Leave a comment above! 

 Avatars from <https://gravatar.com/>¹⁰ are connected to your provided email address.

112 COMMENTS

Newest ▼

Links

1. <https://www.learncpp.com/author/Alex/>
2. <https://www.humix.com/video/-PJCQh4jXO2>
3. <https://en.cppreference.com/w/cpp/regex>
4. <https://www.learncpp.com/cpp-tutorial/stream-states-and-input-validation/#comment-571052>
5. <https://www.learncpp.com/cpp-tutorial/basic-file-io/>
6. <https://www.learncpp.com/>
7. <https://www.learncpp.com/cpp-tutorial/stream-classes-for-strings/>
8. <https://www.learncpp.com/stream-states-and-input-validation/>
9. <https://www.learncpp.com/wordpress/tiga-and-wordpress-25/>
10. <https://gravatar.com/>
11. <https://g.ezoic.net/privacy/learncpp.com>