

# Modern C++ - STL example

---

## 19. An STL example

---

In this chapter we discuss the practical usage of STL. We will solve a simple task without STL and then we will introduce STL instruments step by step.

### Merging two sorted file

---

Suppose we have two sorted files of strings *file1* and *file2*. The task is to create *file3* as a sorted output where we merge the values from the input files.

At the beginning we create a solution without using STL at all.

```
1 #include <iostream>
2 #include <fstream>
3 #include <string>
4
5 using namespace std;
6
7 // simple merge
8 int main()
9 {
10  string s1, s2;
11  ifstream f1("file1.txt");
12  ifstream f2("file2.txt");
13  ofstream f3("file3.txt");
14  f1 >> s1;
15  f2 >> s2;
16  // the usual way:
17  while (f1 || f2)
18  {
19      if (f1 && ((s1 <= s2) || !f2))
20      {
21          f3 << s1 << endl;
22          f1 >> s1;
23      }
24      if (f2 && ((s1 >= s2) || !f1))
25      {
```

```
26     f3 << s2 << endl;
27     f2 >> s2;
28 }
29 }
30 return 0;
31 }
```

In lines 11-13 we open the files. In lines 14-15 we read one-one elements from both input streams. This will indicate *eof* state if any of the files is empty.

We are looping until *all* input files are exhausted. To reach this, in line 19 and 24 we compare the two pre-read elements, and we select the smaller one. The trick in the comparison is that we also select the element if the other file is already empty. the selected element is written to the output and a new pivot element is read from the same file.

Although this solution is working fine and effectively, it is easy to see that its complexity make it hard to maintain. For example, if the client requirement changes to use a different way of comparison, or handle *eof* in a different way, we have to modify the critical part of the loop, and then have to re-test the whole program.

## Using STL in a naïve way

---

In the following iteration we will use the **std::vector** container and the **std::merge** algorithm in a naïve way: we read both input files into two vector of **std::string** and then merge them into a third output vector.

```
1 #include <iostream>
2 #include <fstream>
3 #include <string>
4 #include <algorithm>    // for merge( b1, e1, b2, e2, b3 [,opc_comp])
5 #include <vector>
6
7 using namespace std;
8
9 int main()
10 {
11     ifstream f1("file1.txt");
12     ifstream f2("file2.txt");
13     ofstream f3("file3.txt");
14     string s;
15     vector<string> v1;
```

```

16 while ( if1 >> s ) v1.push_back(s);
17 vector<string> v2;
18 while ( if2 >> s ) v2.push_back(s);
19
20 // allocate the space for the result
21 vector<string> v3(v1.size()+v2.size()); // expensive...
22
23 merge( v1.begin(), v1.end(),
24        v2.begin(), v2.end(),
25        v3.begin()); // v3[i++] = *current++
26 for ( int i = 0; i < v3.size(); ++i)
27     f3 << v3[i] << endl;
28 return 0;
29 }

```

We read the content of input files into **v1** and **v2** in lines 16 and 18 respectively. As **merge** put the output into a fully pre-allocated container we have to create **v3** with the necessary size in line 21. This is an expensive operation as we have to create an empty string into every vector element – just to be overwritten in the next statement.

The actual merge happens in line 23-25 in the **merge** function, where the first four parameters define the input intervals, and the fifth parameter is the iterator to the beginning of the output interval. The merge algorithm supposes that the length of the output interval is long enough, that we fulfilled in line 21.

Although, the program itself is trivial now, it has both functional and performance problems. First of all: the input file sizes are limited by the memory of the program: in one point we keep all input strings twice in the memory: (ones as the input either in **v1** or **v2** and as the output in **v3**). The unnecessary creation of empty strings in line 21 is also a burden.

## Using insert adaptors

---

Although, it does not really changes the major problems with the program above, we can solve the last issue. It would be a logical choice not to preallocate (and then overwrite) all output elements, but extend the vector using the **push\_back** method one by one.

First of all, we do not allocate all elements of the output vector in advance, we just create an empty vector **v3** in line 22. Since extending a vector with  $N$  elements at the end reallocates the buffer of a vector roughly  $\log(N)$  times, and copies the elements, we pre-allocate the buffer with the **reserve** method in the next line in one step. After calling **reserve** the *logical* size of **v3** is still zero, only the *physical buffer* is allocated.

The issue comes with the **merge** function. Merge wants to *overwrite* the output vector and not to *extend* it. Here we apply an *adaptor*.

```
1 #include <iostream>
2 #include <fstream>
3 #include <string>
4 #include <algorithm>    // for merge( b1, e1, b2, e2, b3 [,opc_comp])
5 #include <vector>
6 #include <iterator>    // for back_insert_iterator
7
8 using namespace std;
9
10 int main()
11 {
12     ifstream f1("file1.txt");
13     ifstream f2("file2.txt");
14     ofstream f3("file3.txt");
15     string s;
16     vector<string> v1;
17     while ( f1 >> s ) v1.push_back(s);
18     vector<string> v2;
19     while ( f2 >> s ) v2.push_back(s);
20
21     // do not allocate the space for the result
22     vector<string> v3;           // cheap...
23     v3.reserve(v1.size()+v2.size()); // allocate all space in once
24
25     merge( v1.begin(), v1.end(),
26           v2.begin(), v2.end(),
27           back_inserter(v3));    // v3.push_back(*current++)
28     for ( int i = 0; i < v3.size(); ++i)
29         f3 << v3[i] << endl;
30     return 0;
31 }
```

Adaptors modify the interface of a container (or an iterator). That is here we will replace **v3.begin()** with a **back\_inserter\_iterator** object in line 27. The `back_insert_iterator` class is defined in header `<iterator>` included in line 6 and the object itself is created by the **back\_inserter** factory function.

A `back_inserter_iterator` adaptor is *acting* like an output iterator, but its **operator=**

instead of overwriting the referred element on the container calls the **push\_back** method on the original container (here on **v3**). Thus merge *thinks* that overwrites the output range, but in reality it *extends* it.

There are also inserters to insert elements into front calling **push\_front** or just **insert** for associative containers.

## Input and output iterator adaptors

---

The major issue with the previous program version was to store all elements in the memory (twice) in temporary vectors.

To solve this we will apply adaptors for the input and output files. Such adaptors are called *input* and *output iterator adaptors*.

```
1 #include <iostream>
2 #include <fstream>
3 #include <string>
4 #include <algorithm>
5 #include <iterator>      // input- and output-iterators
6
7 using namespace std;
8
9 int main()
10 {
11     ifstream f1("file1.txt");
12     ifstream f2("file2.txt");
13     ofstream f3("file3.txt");
14     // istream_iterator(if1) -> f1 >> *current
15     // istream_iterator() -> EOF
16     // ostream_iterator(f3,x) -> of << *current << x
17     merge( istream_iterator<string>(f1), istream_iterator<string>(),
18           istream_iterator<string>(f2), istream_iterator<string>(),
19           ostream_iterator<string>(f3, "\n") );
20     return 0;
21 }
```

The iterator adaptors have the same idea as the inserters. They are acting like iterators but reading from and write to the wrapped stream objects. To express the *end* of the input range (here the *eof* event on read), default constructed `istream_iterator` objects are given as the first and third parameter of `merge`.

As the same stream can be read as sequence of characters, numbers or strings, the **input\_iterator** class is a template to define the unit to read (here string). For output iterator we can define a *separator* to be written after every output action (here “\n”).

Notice that in this version there is no vectors to buffer input and output, therefore our program has a low memory footprint and can also work on infinite input streams.

## Functors

---

Suppose, later the customer of our program wants a change: in the sorting of the strings we have to neglect the difference between upper case and lower case letters. The problem is, that the **merge** algorithm by default uses the **operator<** between the input elements, here **string** objects.

In fact, merge has a second version, where the sixth parameter is used to compare the input elements. Such *functors* are classes acting like functions, defining the *function call operator* (**operator()**). Functors can appear in various ways, but *unary predicates*, taking one parameter and returning **bool**, and *binary predicates* taking two parameters and returning **bool** has a major role for STD algorithms. Unary predicates are used in algorithms, like **find\_if** or **remove\_if**, while binary ones mainly as comparisons.

```
1 #include <iostream>
2 #include <fstream>
3 #include <string>
4 #include <cctype>
5 #include <algorithm>
6 #include <iterator>
7
8 using namespace std;
9
10 struct my_less // function object: "functor"
11 {
12     bool operator()(const string& s1, const string& s2) const
13     {
14         string us1 = s1;
15         string us2 = s2;
16         // TODO: use locale object
17         transform( s1.begin(), s1.end(), us1.begin(), toupper);
18         transform( s2.begin(), s2.end(), us2.begin(), toupper);
19         return us1 < us2;
20     }
```

```
21 };
22 int main()
23 {
24     ifstream f1("file1.txt");
25     ifstream f2("file2.txt");
26     ifstream f3("file3.txt");
27
28     merge( istream_iterator<string>(f1), istream_iterator<string>(),
29           istream_iterator<string>(f2), istream_iterator<string>(),
30           ostream_iterator<string>(f3, "\n"), my_less() );
31     return 0;
32 }
```

The **my\_less** class has a function call operator method to compare two strings and returning **true** when the case insensitive comparison recognizes that the first parameter is *less* than the second. This class is *callable*. Thus it acts as a *binary predicates*.

A temporary object is created as the sixth parameter of merge algorithm and its function call operator is called for every time merge compares two strings.

The implementation of the function call operator is only demonstration purposes, a real solution would utilize the **local** library.

## Functor with state

---

More complex functors can be defined using the full feature set of object-oriented programming.

For example, the customer may not simply want to merge the input files in a sorted way, but taking one element from file1 and than one from file2, “zipping” the files. We can modify the previous solution to returning the predicates alternating **true** and **false**. This requires to store the previous state: we will create an attribute in the functor for this purpose.

```
1 #include <iostream>
2 #include <fstream>
3 #include <string>
4 #include <cctype>
5 #include <algorithm>
6 #include <iterator>
7
```

```
8 using namespace std;
9
10 struct zipper
11 {
12     zipper() : _flag(false) {}
13     bool operator()(const string&, const string&) const
14     {
15         _flag = ! _flag;
16         return _flag;
17     }
18     bool _flag;
19 };
20 int main()
21 {
22     ifstream f1("file1.txt");
23     ifstream f2("file2.txt");
24     ifstream f3("file3.txt");
25
26     merge( istream_iterator<string>(f1), istream_iterator<string>(),
27           istream_iterator<string>(f2), istream_iterator<string>(),
28           ostream_iterator<string>(f3, "\n"), zipper() );
29     return 0;
30 }
```

Usually, it is not suggested to define functors with *state*. Some STL algorithm may copy or assign the functor and therefore its behavior may hard to predict. In such simple cases, however, we can use predicates with state.

Notice, that since the function call operator does not depend on its parameters, we declare them without formal parameter name.

## Generalization

---

In the last step we generalize the zipper class. We make it template, so it can work for any type we can read in and write out to files. (We do not use any comparison anymore, so this is not a restriction now.)

We can apply further (run-time) parameters to tell teh zipper how many elements to read from “left” (file1) and from “right” (file2). Also we can specify whether to start from left or right.



```
1 #include <iostream>
2 #include <fstream>
3 #include <string>
4 #include <cctype>
5 #include <algorithm>
6 #include <iterator>
7
8 using namespace std;
9
10 template <typename T>
11 class zipper
12 {
13 public:
14     zipper(int l, int r, bool fl = true) :
15         left(l), right(r), from_left(fl), cnt(0) { }
16     bool operator()( const T&, const T&) const
17     {
18         bool ret = from_left;
19         const int max = from_left ? left : right;
20         if ( ++cnt == max )
21         {
22             cnt = 0;
23             from_left = ! from_left;
24         }
25         return ret;
26     }
27 private:
28     const int left;
29     const int right;
30     const bool from_left;
31     int cnt;
32 };
33 int main()
34 {
35     ifstream f1("file1.txt");
36     ifstream f2("file2.txt");
37     ifstream f3("file3.txt");
38
39     merge( istream_iterator<string>(f1), istream_iterator<string>(),
```

```
40     istream_iterator<string>(f2), istream_iterator<string>(),
41     ostream_iterator<string>(f3, "\n"), zipper() );
42     return 0;
43 }
```

A further advantage of such STL-style programming is that we separated the functor logic: a different individual can code (and test) it, while somebody else can write the other part of the program.