

Object-Oriented Programming for Scientific Computing

STL Containers and Iterators

Ole Klein

Interdisciplinary Center for Scientific Computing
Heidelberg University
`ole.klein@iwr.uni-heidelberg.de`

20. Juni 2017

The Standard Template Library (STL)

The Standard Template Library (STL) is a part of the C++ standard that provides classes with often-used functionality.

- It provides algorithms and functors to work with these classes.
- It also formulates interfaces, which must be provided by other collections of classes in order to be used as STL classes, or can be used to write algorithms that work with all STL-like container classes.
- The STL is a new level of abstraction, which frees the programmer from the necessity to write frequently used constructs such as dynamic arrays, lists, binary trees, search algorithms, and so on him/herself.
- STL algorithms are heavily optimized, i.e. if there is an STL algorithm for a problem, one should have a very good reason not to use it.
- Unfortunately, the STL is not self-explanatory.

STL Components

The main components of the STL are:

Containers are used to manage a particular type of object. The various containers have different properties and related advantages and disadvantages.

Iterators make it possible to iterate over the contents of a container. They provide a uniform interface for each STL compliant container, regardless of its internal structure.

Algorithms work with the elements of a container. They use iterators and therefore must only be written once for an arbitrary number of STL-compliant containers.

Functors can be used in STL algorithms to define local operations / comparisons that should be applied to container elements. They can be extended by user-written alternatives.

At first sight, the structure of the STL partially contradicts the original idea of object-oriented programming that algorithms and data belong together.

Containers

STL container classes, or short containers, manage a collection of elements of the same type. Depending on the type of container, the STL gives assurances on the execution speed of certain operations.

There are two fundamentally different types of container:

Sequences are ordered sets of elements with freely selectable arrangement. Each element has its place, which depends on the program execution and not on the value of the element.

Associative Containers are ordered sets sorted according to a certain sorting criterion in which the position of an element depends only on its value.

List of Available Containers

Sequences:

- `array` (C++11), fixed-size array
- `deque`, double ended queue
- `forward_list` (C++11), linked list
- `list`, doubly linked list
- `vector`, variable-size array

Sequence Adaptors:

- `stack`, LIFO (last in, first out)
- `queue`, FIFO (first in, first out)
- `priority_queue`, heap (highest priority out)

Sequence adaptors can use any sequence internally, if the interface matches (compare our `Stack` implementation).

List of Available Containers

Associative Containers:

- `set`, mathematical set
- `multiset`, entries may appear multiple times
- `map`, mathematical mapping, dictionary
- `multimap`, as above

Unordered Associative Containers (C++11):

- `unordered_set`
- `unordered_multiset`
- `unordered_map`
- `unordered_multimap`

Unordered containers sort elements using an internal hash function instead of element comparison.

List of Available Containers

Apart from the above containers, there are several container-like structures that don't fulfill the STL container interface (with some of them not being part of the STL), but are often used instead of or in conjunction with STL containers:

- `pair`, heterogeneous collection of two elements
- `tuple` (C++11), heterogeneous collection
- `valarray`, `vector` with math operations
- `bitset`, specialization of `array` for `bool`

Vector

STL sequence containers are class templates. There are two template arguments, the type of objects to be stored and a so-called allocator that can be used to change the memory management (this is useful for example when you create many small objects and don't want to pay the operating system overhead every time). The second parameter has a default value where `new()` and `delete()` are used.

Vector is a field of variable size.

- Adding and removing elements at the end of a `vector` is fast, i.e. complexity in $O(1)$.
- The element can be accessed directly via an index (random access).
- The memory used by a `vector` is guaranteed to be contiguous (unique among STL containers).
- The memory reserved by a `vector` never shrinks (!).



Abbildung: Structure of a `vector`

Amortized Complexity

- Typically, adding elements at the end of a `vector` is in $O(1)$.
- In individual cases, however, it may take much longer, especially if the allocated storage is no longer sufficient. Then new storage must be allocated, and often data has to be copied over. This is an $O(N)$ process.
- However, the standard library reserves memory blocks of increasing size for a growing vector. The overhead depends on the length of the vector. This optimizes for speed at the expense of memory usage.
- The $O(N)$ case therefore occurs very rarely. This is called “amortized complexity”.
- If it is already known that a certain amount of elements is needed, then one can reserve space with the method `reserve(size_t size)`. This doesn't change the current size of the vector, it only reserves the right amount of memory.

Example: STL Vector

```
#include <iostream>
#include <vector>
#include <string>

int main()
{
    std::vector<double> a(7);
    std::cout << a.size() << std::endl;
    for (int i = 0; i < 7; ++i)
        a[i] = i * 0.1;
    double d = 4 * a[2];
    std::vector<double> c(a);
    std::cout << a.back() << "\u2193" << c.back() << std::endl;
    std::vector<std::string> b;
    b.resize(3);
    for (int i = 2; i >= 0; --i)
        std::cin >> b[i];
    b.resize(4);
    b[3] = "foo";
    b.push_back("bar");
    for (int i = 0; i < b.size(); ++i)
        std::cout << b[i] << std::endl;
}
```

Deque

`Deque` is a “double-ended” queue, it is also a field of dynamic size, but:

- The addition and removal of elements is quick also at the beginning of `deque`, that is $O(1)$, in addition to its end.
- Element access can again be achieved using an index, but the index of an element may change when elements are added to the beginning of the container.
- Internally, `deques` can be implemented using e.g. arrays of arrays, but this is implementation-dependent.




Abbildung: Structure of a deque

List

`List` is a doubly linked list of elements.

- There is no direct access to list elements.
- To reach the 10th element, one must start at the beginning of the `list` and traverse the first nine elements, access to a specific element is therefore $O(N)$.
- Adding and removing elements is fast in any location in the `list`, i.e. $O(1)$.




Abbildung: Structure of a `list`

Example: STL List

```
#include <iostream>
#include <list>
#include <string>

int main()
{
    std::list<double> vals;
    for (int i = 0; i < 7; ++i)
        vals.push_back(i * 0.1);
    vals.push_front(-1);
    std::list<double> copy(vals);
    std::cout << vals.back() << "\n" << copy.back() << std::endl;
    std::cout << vals.front() << "\n" << copy.front() << std::endl;
    for (int i = 0; i < vals.size(); ++i)
    {
        std::cout << i << ":\n" << vals.front() << "\n" << vals.size() <<
            std::endl;
        vals.pop_front();
    }
    std::cout << std::endl;
    for (int i = 0; i < copy.size(); ++i)
    {
        std::cout << i << ":\n" << copy.back() << "\n" << copy.size() <<
            std::endl;
        copy.pop_back();
    }
}
```

C++11: Array

Array is a C++11 replacement for the classical C arrays, i.e. a field of fixed size.

- `array` has two template parameters, the type of the stored objects and the number of elements of the container
- Adding and removing elements isn't possible.
- Element access can be achieved directly via index.
- In contrast to C arrays, an STL `array` knows its own size and can be used like the other STL containers.



Abbildung: Structure of an array

Example: STL Array

```
#include <iostream>
#include <array>
#include <string>

int main()
{
    std::array<double,7> a;
    std::cout << a.size() << std::endl;
    for (int i = 0; i < 7; ++i)
        a[i] = i * 0.1;
    double d = 4 * a[2];
    std::array<double,7> c(a);
    std::cout << a.back() << "␣" << c.back() << std::endl;
    std::array<std::string,4> b;
    for (int i = 2; i >= 0; --i)
        std::cin >> b[i];
    b[3] = "foo";
    for (int i = 0; i < b.size(); ++i)
        std::cout << b[i] << std::endl;
}
```

Set/Multiset

- The containers `set` and `multiset` are sorted sets of elements. Internally, these elements are stored in a tree structure.
- While in a `set` every element may only appear once, a `multiset` may contain elements several times.
- In a `set`, it is particularly important to be able to quickly determine whether an element is in the set or not (and in a `multiset`, how often).
- The search for an element is of optimal complexity $O(\log(N))$.
- `set` and `multiset` have three template parameters: the type of objects, a comparison operator and an allocator. For the last two, there are default values (the predicate `less` and the standard allocator).

Map/Multimap

- The containers `map` and `multimap` consist of sorted pairs of two variables, a key and a value. The entries in the `map` are sorted by the key.
- While each key can only appear once in a `map`, it may exist several times in a `multimap` (independent of the associated value).
- A `map` can be quickly searched for a key and then gives access to the appropriate value.
- The search for a key is of optimal complexity $O(\log(N))$.
- `map` and `multimap` have four template parameters: the type of the keys, the type of the values, a comparison operator and an allocator. For the last two, there are again default values (`less` and `new/delete`).

Container Concepts

- The properties of STL containers are divided into specific categories, called concepts.
- They are, for example, `Assignable`, `EqualityComparable`, `Comparable`, `DefaultConstructible`...
- The objects of a class that are to be stored in a container must be `Assignable` (there is an assignment operator), `Copyable` (there is a copy constructor) , `Destroyable` (there is a public destructor), `EqualityComparable` (there is an `operator==`) and `Comparable` (there is an `operator<`).
- Concepts (constraint sets) like the above can help in writing templates, and have been considered for inclusion in the standard (not just in the context of the STL). However, they haven't been included so far and remain an experimental feature / a way to talk about template design.

Container

A Container itself is `Assignable` (there is an assignment operator), `EqualityComparable` (there is an `operator==`) and `Comparable` (there is an `operator<`).

Associated types:

<code>value_type</code>	The type of object stored. Needs to be <code>Assignable</code> , but not <code>DefaultConstructible</code> .
<code>iterator</code>	The type of iterator. Must be an <code>InputIterator</code> and a conversion to <code>const_iterator</code> must exist.
<code>const_iterator</code>	An iterator through which the elements may be read but not changed.
<code>reference</code>	The type of a reference to <code>value_type</code> .
<code>const_reference</code>	As above, but constant reference.
<code>pointer</code>	As above, but pointer.
<code>const_pointer</code>	As above, but pointer to constant.
<code>difference_type</code>	A type suitable for storing the difference between two iterators.
<code>size_type</code>	An unsigned integer type that can store the distance between two elements.

Container

In addition to the methods of `Assignable`, `EqualityComparable` and `Comparable`, a container always has the following methods:

<code>begin()</code>	Returns an iterator to the first element. If the container is <code>const</code> this is a <code>const_iterator</code> .
<code>end()</code>	As <code>begin()</code> , but points to the location <i>after</i> the last element.
<code>size()</code>	Returns the size of the container, i.e. number of elements, return type is <code>size_type</code> .
<code>max_size()</code>	Returns the maximum size allowed at the moment (capacity), return type is <code>size_type</code> .
<code>empty()</code>	True if the container is empty.
<code>swap(b)</code>	Swaps elements with container <code>b</code> .

Specializations of the Container Concept

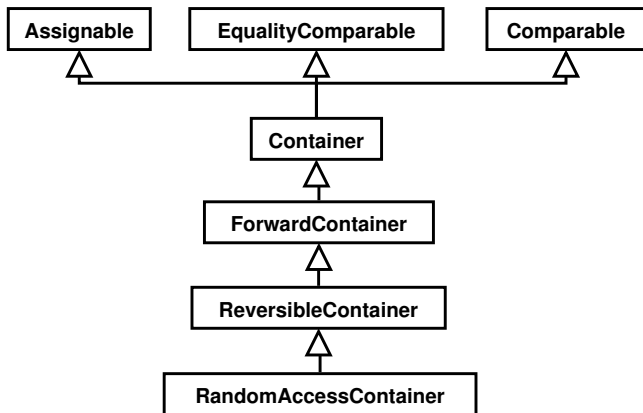


Abbildung: Container concepts

ForwardContainer

`ForwardContainer` is a specialization of the `Container` concept:

- There is an `iterator` with which one can pass through the container in forward direction (`ForwardIterator`).
- Main addition to vanilla `Container`: elements are arranged in definite order.
- This allows for:
 - element-by-element equality (requires `EqualityComparable` elements)
 - lexicographical ordering (requires `Comparable` elements)

The `Container` concept itself does *not* require an iterator to always produce the same sequence of elements. This is an addition, and formalized in the concept `ForwardContainer`.

Other restrictions and constraints are introduced through similar definitions.

ReversibleContainer

Concept ReversibleContainer:

- There is an iterator which allows passing back and forth through the container (`BidirectionalIterator`).

- Additional associated types:

`reverse_iterator`

Iterator in which the `operator++` moves to the previous item instead of the next item.

`const_reverse_iterator`

As above, but `const` version.

- Additional methods:

`rbegin()`

Returns an iterator to the first element of a reverse pass (last element of the container).

`rend()`

As `rbegin()`, but points to the location *before* the first element.

Implementations:

- `std::list`
- `std::set`
- `std::map`

RandomAccessContainer

Concept RandomAccessContainer:

- Is a specialization of ReversibleContainer.
- There is an iterator with which one can gain access to arbitrary elements of the container (RandomAccessIterator, uses an index).
- Additional methods: `operator[]` (size_type) (and `const` version), access operators for random access.

Implementations:

- `std::vector`
- `std::deque`

Sequence




Abbildung: Sequence concepts

Sequence Methods

A Sequence is a specialization of the concept of `ForwardContainer` (so one can at least in one direction iterate over the container) and is `DefaultConstructible` (there is a constructor without argument / an empty container).

<code>X(n, t)</code>	Generates a sequence with $n \geq 0$ elements initialized with <code>t</code> .
<code>X(n)</code>	As above, but initialized with the default constructor.
<code>X(i, j)</code>	Generates a sequence which is a copy of the range <code>[i, j)</code> . Here <code>i</code> and <code>j</code> are <code>InputIterators</code> .
<code>insert(p, t)</code>	Inserts the element <code>t</code> in front of the one the iterator <code>p</code> points to, and returns an iterator pointing to the inserted element.
<code>insert(p, i, j)</code>	As above, but for the range <code>[i, j)</code> .
<code>insert(p, n, t)</code>	As above, but inserts <code>n</code> copies of <code>t</code> and returns an iterator to the last of them.
<code>erase(p)</code>	Invokes the destructor for the element the iterator <code>p</code> points to and deletes it from the container.
<code>erase(p, q)</code>	As above, but for the range <code>[p, q)</code> .
<code>erase()</code>	Deletes all elements.
<code>resize(n, t)</code>	Shrinks or enlarges the container to size <code>n</code> and initializes new elements with <code>t</code> .
<code>resize(n)</code>	The same as <code>resize(n, T())</code> .

C++11: `emplace`

Creating an element and then inserting it creates a copy of the stored element. This is somewhat alleviated with the introduction of C++11 move semantics, but copies are still necessary if the element has to be stored in a specific location (contiguous memory layout of `vector`).

C++11 introduces `emplace` as a substitute for `insert`:

<code>emplace(p, args)</code>		Inserts an element <code>t</code> , constructed using <code>args</code> , in front of the one <code>p</code> points to, and returns an iterator pointing to the emplaced element.
-------------------------------	--	---

The same holds for `emplace_back(args)` (instead of `push_back`) and `emplace_front(args)` (instead of `push_front`) for containers that have that functionality.

Emplacement is possible whenever insertion is possible, and will not be repeated in the following container descriptions.

Complexity Guarantees for Sequences

The following complexity guarantees are given for sequence containers of the STL:

- The constructors, $x(n, t)$ $x(n)$ and $x(i, j)$ have linear complexity.
- Inserting elements with `insert(p, n, t)` and `insert(p, i, j)` and deleting them with `erase(p, q)` has linear complexity.
- The complexity of inserting and removing single elements depends on the sequence implementation.

BackInsertionSequence

Methods in addition to those from the `Sequence` concept:

<code>back()</code>	Returns a reference to the last element.
<code>push_back(t)</code>	Inserts a copy of <code>t</code> after the last element.
<code>pop_back()</code>	Deletes the last element of the sequence.

Complexity Guarantees

`back()`, `push_back(t)`, and `pop_back()` have amortized constant complexity, i.e. in individual cases it may take longer but the average time is independent of the number of elements.

Implementations

- `std::vector`
- `std::list`
- `std::deque`

FrontInsertionSequence

Methods in addition to those from the `Sequence` concept:

<code>front()</code>	Returns a reference to the first element.
<code>push_front(t)</code>	Inserts a copy of <code>t</code> before the first element.
<code>pop_front()</code>	Removes the first element of the sequence.

Complexity Guarantees

`front()`, `push_front(t)`, and `pop_front()` have amortized constant complexity.

Implementations

- `std::list`
- `std::deque`

STL Sequence Containers

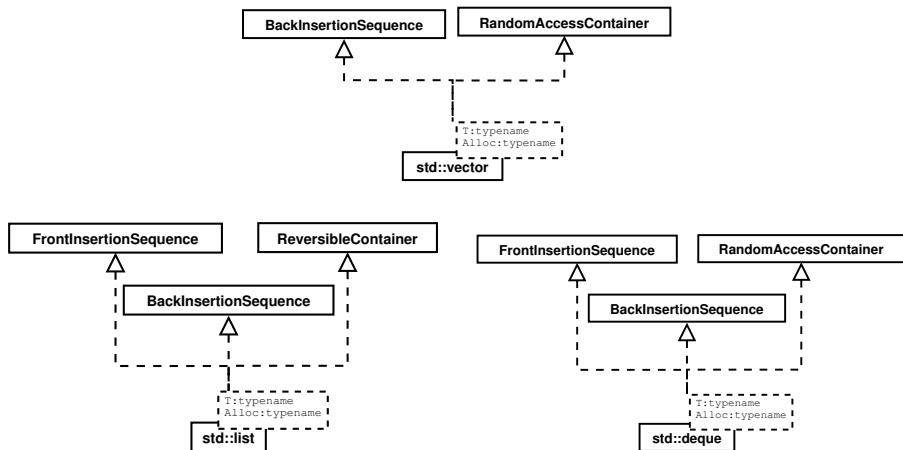


Abbildung: STL sequence containers

Sequence Test: Vector vs. List

Task by J. Bentley and B. Stroustrup (Element Shuffle):

- For a fixed number N , generate N random integers and insert them into a sorted sequence.

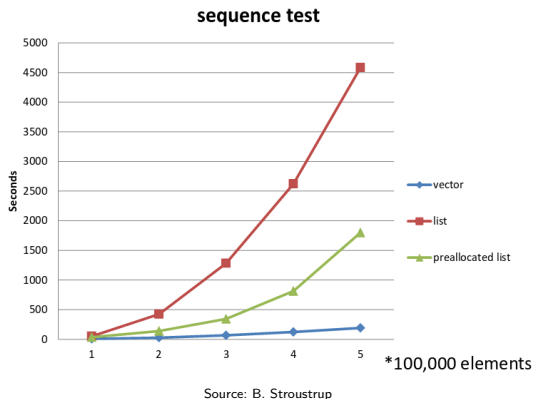
Example:

- 5
 - 1, 5
 - 1, 4, 5
 - 1, 2, 4, 5
- Remove elements at random while keeping the sequence sorted.

Example:

- 1, 2, 4, 5
 - 1, 4, 5
 - 1, 4
 - 4
- For which N should a `list` be used, and in which cases a `vector`?

Sequence Test: Vector vs. List



- Despite random insertion / deletion, `vector` is faster by an order of magnitude
- Linear search for both containers, despite bisection being available for `vector` (!)
- Search completely dominates move required by `vector`
- Non-optimized `list` performs one allocation / deallocation per element (!)

Associative Containers




Abbildung: Associative container concepts

AssociativeContainer

An `AssociativeContainer` is a specialisation of `ForwardContainer` and `DefaultConstructible`.

- Additional associated type: `key_type` is the type of a key.

- Additional Methods:

<code>erase(k)</code>	Deletes all entries with key <code>k</code> .
<code>erase(p)</code>	Deletes the element that iterator <code>p</code> points to.
<code>erase(p,q)</code>	As above, but for the range <code>[p,q)</code> .
<code>clear()</code>	Deletes all elements.
<code>find(k)</code>	Returns an iterator pointing at the item (or one of the items) with key <code>k</code> or <code>end()</code> if this key doesn't exist.
<code>count(k)</code>	Returns the number of elements with key <code>k</code> .
<code>equal_range(k)</code>	Returns a pair <code>p</code> of iterators so that <code>[p.first,p.second)</code> consists of all elements that have key <code>k</code> .

AssociativeContainer

Assurances

Continuous memory: all elements with the same key directly follow one another.

Immutability of the key: The key of each element of an associative container is unchangeable.

Complexity Guarantees

<code>erase(k)</code>	Average complexity at most $O(\log(\text{size}()) + \text{count}(k))$
<code>erase(p)</code>	Average complexity constant
<code>erase(p, q)</code>	Average complexity at most $O(\log(\text{size}()) + N)$
<code>count(k)</code>	Average complexity at most $O(\log(\text{size}()) + \text{count}(k))$
<code>find(k)</code>	Average complexity at most logarithmic
<code>equal_range(k)</code>	Average complexity at most logarithmic

These are just average complexities, and the worst case can be significantly more expensive!

SimpleAssociativeContainer and PairAssociativeContainer

These are specializations of the `AssociativeContainer`.

SimpleAssociativeContainer

has the following restrictions:

- `key_type` and `value_type` must be the same.
- `iterator` and `const_iterator` must have the same type.

PairAssociativeContainer

- introduces the associated data type `mapped_type`. The container maps `key_type` to `mapped_type`.
- The `value_type` is `std::pair<key_type,mapped_type>`.

SortedAssociativeContainer

This specialization uses a sorting criterion for the key. Two keys are equivalent if none is smaller than the other.

Additional associated types

<code>key_compare</code>	The type implementing <code>StrictWeakOrdering</code> to compare two keys.
<code>value_compare</code>	The type implementing <code>StrictWeakOrdering</code> to compare two values. Compares two objects of type <code>value_type</code> by handing their keys over to <code>key_compare</code> .

Additional Methods

<code>key_compare()</code>	Returns the key comparison object.
<code>value_compare()</code>	Returns the value comparison object.
<code>lower_bound(k)</code>	Returns an iterator to the first element whose key is not less than <code>k</code> , or <code>end()</code> if there is no such element.
<code>upper_bound(k)</code>	Returns an iterator to the first element whose key is greater than <code>k</code> , or <code>end()</code> if there is no such element.

SortedAssociativeContainer

Assurances

`value_compare`: if `t1` and `t2` have keys `k1` and `k2`, then it is guaranteed that

`value_compare()(t1,t2) == key_compare()(k1,k2)` is **true**.

Increasing order of the elements is guaranteed.

Complexity Guarantees

- `key_compare()(k1,k2)` and `value_compare()(t1,t2)` have constant complexity.
- `lower_bound(k)` and `upper_bound(k)` are logarithmic.

UniqueAssociativeContainer and MultipleAssociativeContainer

A `UniqueAssociativeContainer` is an `AssociativeContainer` with the additional property that each key occurs at most once.

A `MultipleAssociativeContainer` is an `AssociativeContainer` in which each key can appear several times.

Additional Methods

<code>x(i,j)</code>	Creates an associative container from the items in the range <code>[i,j)</code> .
<code>insert(t)</code>	Inserts the <code>value_type</code> <code>t</code> and returns a <code>std::pair</code> containing an iterator to the copy of <code>t</code> and a <code>bool</code> (<code>true</code> if the copy has just been inserted)
<code>insert(i,j)</code>	Inserts all elements in the range <code>[i,j)</code> .

UniqueAssociativeContainer and MultipleAssociativeContainer

Complexity Guarantees

- The average complexity of `insert(t)` is at most logarithmic.
- The average complexity of `insert(i,j)` is at most $O(N * \log(\text{size}()) + N)$, where $N=j-i$

Associative Container Classes

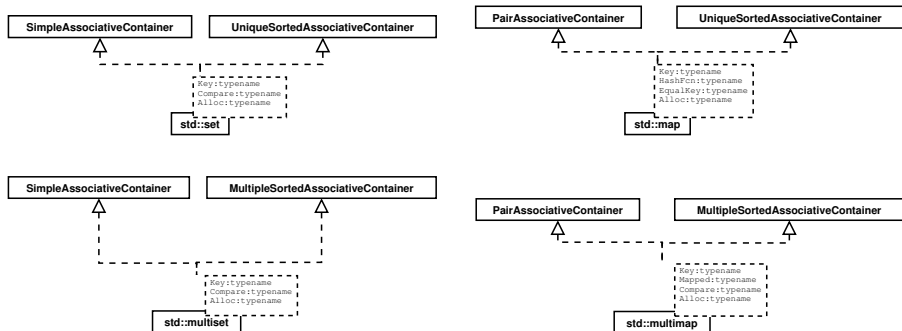


Abbildung: Associative container classes

Properties of the Different Container Classes

	vector	deque	list	set	map
Typical internal data structure	Dynamic array	Array of arrays	Doubly linked list	Binary tree	Binary tree
Elements	values	values	values	values	keys/values
Search	slow	slow	very slow	fast	fast (key)
Insert/delete fast	end	beginning and end	everywhere	—	—
Frees memory of removed elements	never	sometimes	always	always	always
Allows preallocation	yes	no	—	—	—

Table: Properties of the different container classes

Which Container Should Be Used?

- If there is no reason to use a specific container, then `vector` should be used, because it is the simplest data structure and allows random access.
- If elements often have to be inserted/removed at the beginning or at the end, then a `deque` should be used. This container will shrink again when items are removed.
- If elements have to be inserted/removed/moved at arbitrary locations, then a `list` is the container of choice. Even moving all elements from one `list` into another can be done in constant time. But there is no random access.
- If it should be possible to repeatedly search for items in a fast way, one should use a `set` or `multiset`.
- If it is necessary to manage pairs of keys and values (as in a dictionary or phone book) then one uses a `map` or `multimap`.

These are only general recommendations, the “correct” choice depends on specifics! If in doubt, use `vector` if at all possible.