

# Chapter 7. Collections

---

## Topics in This Chapter

- [7.1 An Overview of the Collections Framework](#)
  - [7.2 Iterators](#)
  - [7.3 Sets](#)
  - [7.4 Maps](#)
  - [7.5 Other Collections](#)
  - [7.6 Views](#)
  - [Exercises](#)
- 

Many data structures have been developed so programmers can store and retrieve values efficiently. The Java API provides implementations of common data structures and algorithms, as well as a framework to organize them. In this chapter, you will learn how to work with lists, sets, maps, and other collections.

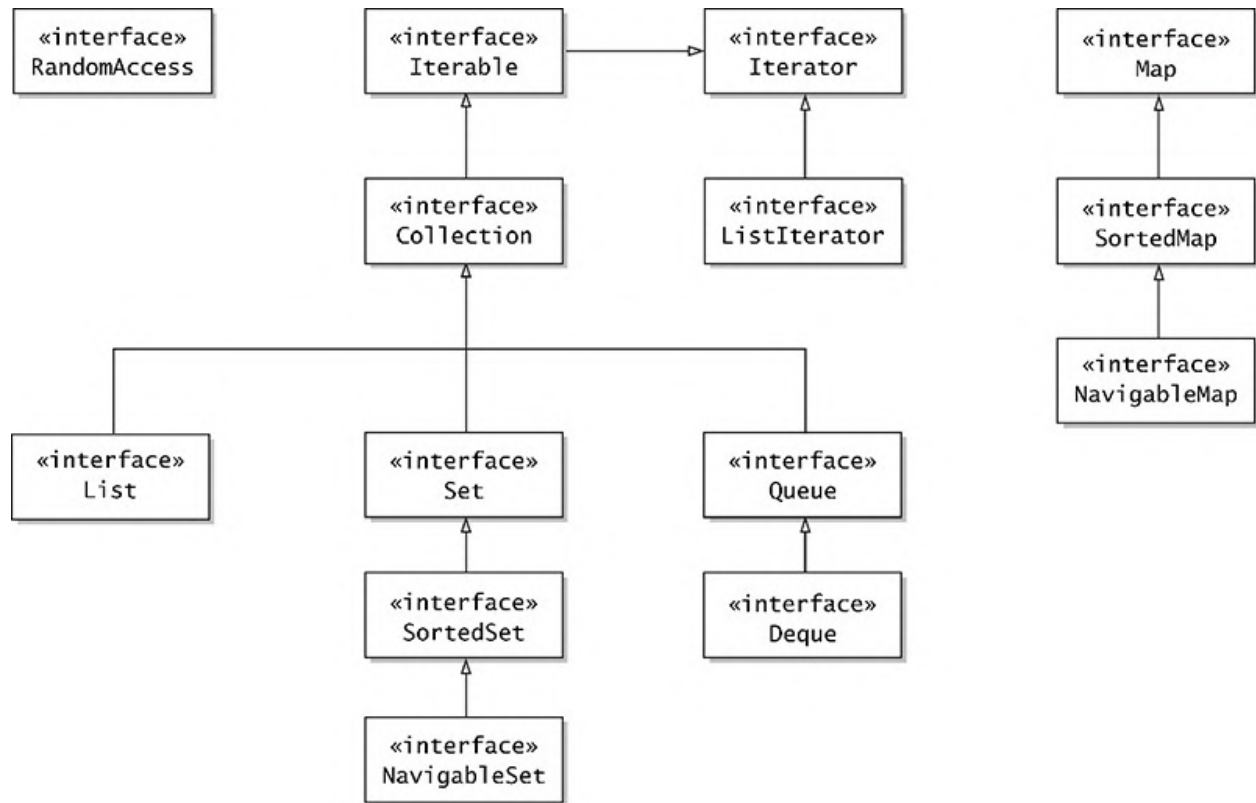
The key points of this chapter are:

1. The `Collection` interface provides common methods for all collections, except for maps which are described by the `Map` interface.
2. A list is a sequential collection in which each element has an integer index.
3. A set is optimized for efficient containment testing. Java provides `HashSet` and `TreeSet` implementations.
4. For maps, you have the choice between `HashMap` and `TreeMap` implementations. A `LinkedHashMap` retains insertion order.
5. The `Collection` interface and `Collections` class provide many useful algorithms: set operations, searching, sorting, shuffling, and more.
6. Views provide access to data stored elsewhere using the standard collection interfaces.

## 7.1 An Overview of the Collections Framework

The Java collections framework provides implementations of common data structures. To make it easy to write code that is independent of the choice of data structures, the collections framework provides a number of common interfaces, shown in [Figure 7-1](#). The fundamental interface is `Collection` whose

methods are shown in [Table 7-1](#).



**Figure 7-1** Interfaces in the Java collections framework

**Table 7-1** The Methods of the `Collection<E>` Interface

Method	Description
<code>boolean add(E e)</code>	
<code>boolean addAll(Collection&lt;? extends E&gt; c)</code>	Adds <code>e</code> , or the elements in <code>c</code> . Returns <code>true</code> if the collection changed.
<code>boolean remove(Object o)</code>	
<code>boolean removeAll(Collection&lt;?&gt; c)</code>	
<code>boolean retainAll(Collection&lt;?&gt; c)</code>	Removes <code>o</code> , or the elements in <code>c</code> , or the elements not in <code>C</code> , or matching elements, or all elements. The first four methods return <code>true</code> if the collection changed.
<code>boolean removeIf(Predicate&lt;? super E&gt; filter)</code>	

<code>void clear()</code>	
<code>int size()</code>	Returns the number of elements in this collection.
<code>boolean isEmpty()</code> <code>boolean contains(Object o)</code> <code>boolean containsAll(Collection&lt;?&gt; c)</code>	Returns <code>true</code> if this collection is empty, or contains <code>o</code> , or contains all elements in <code>c</code> .
<code>Iterator&lt;E&gt; iterator()</code> <code>Stream&lt;E&gt; stream()</code> <code>Stream&lt;E&gt; parallelStream()</code> <code>Splitter&lt;E&gt; splitter()</code>	Yields an iterator, or a stream, or a possibly parallel stream, or a splitter for visiting the elements of this collection. See <a href="#">Section 7.2</a> for iterators and <a href="#">Chapter 8</a> for streams. Splitters are only of interest to implementors of streams.
<code>Object[] toArray()</code> <code>T[] toArray(T[] a)</code>	Returns an array with the elements of this collection. The second method returns <code>a</code> if it has sufficient length.

A `List` is a sequential collection: Elements have position 0, 1, 2, and so on. [Table 7-2](#) shows the methods of that interface.

**Table 7-2** The `List` Interface

Method	Description
<code>boolean add(int index, E e)</code> <code>boolean addAll(int index, Collection&lt;? extends E&gt; c)</code> <code>boolean add(E e)</code> <code>boolean addAll(Collection&lt;? extends E&gt; c)</code>	Adds <code>e</code> , or the elements in <code>c</code> , before <code>index</code> or to the end. Returns <code>true</code> if the list changed.
<code>E get(int index)</code> <code>E set(int index, E element)</code> <code>E remove(int index)</code>	Gets, sets, or removes the element at the given index. The last two methods return the element at the index before the call.
<code>int indexOf(Object o)</code> <code>int lastIndexOf(Object o)</code>	Returns the index of the first or last element equal to <code>o</code> , or <code>-1</code> if there is no

match.

---

<code>ListIterator&lt;E&gt;</code> <code>listIterator()</code> <code>ListIterator&lt;E&gt;</code> <code>listIterator(int index)</code>	Yields a list iterator for all elements or the elements starting at <code>index</code> .
<code>void</code> <code>replaceAll(UnaryOperator&lt;E&gt;</code> <code>operator)</code>	Replaces each element with the result of applying the operator to it.
<code>void sort(Comparator&lt;?</code> <code>super E&gt; c)</code>	Sorts this list, using the ordering given by <code>c</code> .
<code>static List&lt;E&gt; of(E...</code> <code>elements)</code>	Yields an unmodifiable list containing the given elements.
<code>List&lt;E&gt; subList(int</code> <code>fromIndex, int toIndex)</code>	Yields a view ( <a href="#">Section 7.6</a> ) of the sublist starting at <code>fromIndex</code> and ending before <code>toIndex</code> .

---

The `List` interface is implemented both by the `ArrayList` class, which you have seen throughout this book, and the `LinkedList` class. If you took a course on data structures, you probably remember a linked list—a sequence of linked nodes, each carrying an element. Insertion in the middle of a linked list is speedy—you just splice in a node. But to get to the middle, you have to follow all the links from the beginning, which is slow. There are applications for linked lists, but most application programmers will probably stick with array lists when they need a sequential collection. Still, the `List` interface is useful. For example, the method `Collections.nCopies(n, o)` returns a `List` object with `n` copies of the object `o`. That object “cheats” in that it doesn't actually store `n` copies but, when you ask about any one of them, returns `o`.



### Caution

The `List` interface provides methods to access the  $n$ th element of a list, even though such an access may not be efficient. To indicate that it is, a collection class should implement the `RandomAccess` interface. This is a tagging interface without methods. For example, `ArrayList` implements `List` and `RandomAccess`, but `LinkedList` implements only the `List` interface.

---

In a `Set`, elements are not inserted at a particular position, and duplicate elements are not allowed. A `SortedSet` allows iteration in sort order, and a `NavigableSet` has methods for finding neighbors of elements. You will learn more about sets in [Section 7.3](#), “Sets” (page 242).

A `Queue` retains insertion order, but you can only insert elements at the tail and remove them from the head (just like a queue of people). A `Deque` is a double-ended queue with insertion and removal at both ends.

All collection interfaces are generic, with a type parameter for the element type (`Collection<E>`, `List<E>`, and so on). The `Map<K, V>` interface has a type parameter `K` for the key type and `V` for the value type.

You are encouraged to use the interfaces as much as possible in your code. For example, after constructing an `ArrayList`, store the reference in a variable of type `List`:

[Click here to view code image](#)

```
List<String> words = new ArrayList<>();
```

Whenever you implement a method that processes a collection, use the least restrictive interface as parameter type. Usually, a `Collection`, `List`, or `Map` will suffice.

One advantage of a collections framework is that you don't have to reinvent the wheel when it comes to common algorithms. Some basic algorithms (such as `addAll` and `removeIf`) are methods of the `Collection` interface. The `Collections` utility class contains many additional algorithms that operate on various kinds of collections. You can sort, shuffle, rotate, and reverse lists, find the maximum or minimum, or the position of an arbitrary element in a collection, and generate collections with no elements, one element, or `n` copies of the same element. [Table 7-3](#) provides a summary.

**Table 7-3** Useful Methods of the `Collections` Class

<b>Method (all are static)</b>	<b>Description</b>
<code>boolean disjoint(Collection&lt;?&gt; c1, Collection&lt;?&gt; c2)</code>	Returns <code>true</code> if the collections have no elements in common.
<code>boolean addAll(Collection&lt;? super</code>	Adds all elements to <code>c</code> .

<code>T&gt; c, T... elements)</code>	
<code>void copy(List&lt;? super T&gt; dest, List&lt;? extends T&gt; src)</code>	Copies all elements from <code>src</code> to the same indexes in <code>dest</code> (which must be at least as long as <code>src</code> ).
<code>boolean replaceAll(List&lt;T&gt; list, T oldVal, T newVal)</code>	Replaces all <code>oldVal</code> elements with <code>newVal</code> , either of which may be <code>null</code> . Returns <code>true</code> if at least one match was found.
<code>void fill(List&lt;? super T&gt; list, T obj)</code>	Sets all elements of the list to <code>obj</code> .
<code>List&lt;T&gt; nCopies(int n, T o)</code>	Yields an immutable list with <code>n</code> copies of <code>o</code> .
<code>int frequency(Collection&lt;?&gt; c, Object o)</code>	Returns the number of elements in <code>c</code> equal to <code>o</code> .
<code>int indexOfSubList(List&lt;?&gt; source, List&lt;?&gt; target)</code> <code>int lastIndexOfSubList(List&lt;?&gt; source, List&lt;?&gt; target)</code>	Returns the start of the first or last occurrence of the target list within the source list, or <code>-1</code> if there is none.
<code>int binarySearch(List&lt;? extends Comparable&lt;? super T&gt;&gt; list, T key)</code> <code>int binarySearch(List&lt;? extends T&gt; list, T key, Comparator&lt;? super T&gt; c)</code>	Returns the position of the key, assuming that the list is sorted by the natural element order or <code>c</code> . If the key is not present, returns <code>-i - 1</code> where <code>i</code> is the location at which the key should be inserted.
<code>sort(List&lt;T&gt; list)</code> <code>sort(List&lt;T&gt; list, Comparator&lt;? super T&gt; c)</code>	Sorts the list, using the natural element order or <code>c</code> .
<code>void swap(List&lt;?&gt; list, int i, int j)</code>	Swaps the elements at the given position.
	Rotates the list, moving the

<code>void rotate(List&lt;?&gt; list, int distance)</code>	element with index <code>i</code> to <code>(i + distance) % list.size()</code> .
<code>void reverse(List&lt;?&gt; list)</code> <code>void shuffle(List&lt;?&gt; list)</code> <code>void shuffle(List&lt;?&gt; list, Random rnd)</code>	Reverses or randomly shuffles the list.
<code>synchronized(Collection List Set SortedSet NavigableSet Map SortedMap NavigableMap)()</code>	Yields a synchronized view (see <a href="#">Section 7.6</a> ).
<code>unmodifiable(Collection List Set SortedSet NavigableSet Map SortedMap NavigableMap)()</code>	Yields an unmodifiable view (see <a href="#">Section 7.6</a> ).
<code>checked(Collection List Set SortedSet NavigableSet Map SortedMap NavigableMap Queue)()</code>	Yields a checked view (see <a href="#">Section 7.6</a> ).

## 7.2 Iterators

Each collection provides a way to iterate through its elements in some order. The `Iterable<T>` superinterface of `Collection` defines a method

```
Iterator<T> iterator()
```

It yields an iterator that you can use to visit all elements.

[Click here to view code image](#)

```
Collection<String> coll = ...;
Iterator<String> iter = coll.iterator();
while (iter.hasNext()) {
    String element = iter.next();
    Process element
}
```

In this case, you can simply use the enhanced for loop:

```
for (String element : coll) {
    Process element
}
```



### Note

For *any* object `c` of a class that implements the `Iterable<E>`

interface, the enhanced `for` loop is translated to the preceding form.

---

The `Iterator` interface also has a `remove` method which removes the previously visited element. This loop removes all elements that fulfill a condition:

[Click here to view code image](#)

```
while (iter.hasNext()) {
    String element = iter.next();
    if (element fulfills the condition)
        iter.remove();
}
```

However, it is easier to use the `removeIf` method:

[Click here to view code image](#)

```
coll.removeIf(e -> e fulfills the condition);
```

---



### Caution

The `remove` method removes the last element that the iterator has returned, not the element to which the iterator points. You can't call `remove` twice without an intervening call to `next` or `previous`.

---

The `ListIterator` interface is a subinterface of `Iterator` with methods for adding an element before the iterator, setting the visited element to a different value, and for navigating backwards. It is mainly useful for working with linked lists.

[Click here to view code image](#)

```
List<String> friends = new LinkedList<>();
ListIterator<String> iter = friends.listIterator();
iter.add("Fred"); // Fred |
iter.add("Wilma"); // Fred Wilma |
iter.previous(); // Fred | Wilma
iter.set("Barney"); // Fred | Barney
```

---



### Caution

If you have multiple iterators visiting a data structure and one of them mutates it, the other ones can become invalid. An invalid iterator may throw a `ConcurrentModificationException` if you continue



using it.

---

## 7.3 Sets

A set can efficiently test whether a value is an element, but it gives up something in return: It doesn't remember in which order elements were added. Sets are useful whenever the order doesn't matter. For example, if you want to disallow a set of bad words as usernames, their order doesn't matter. You just want to know whether a proposed username is in the set or not.

The `Set` interface is implemented by the `HashSet` and `TreeSet` classes. Internally, these classes use very different implementations. If you have taken a course in data structures, you may know how to implement hash tables and binary trees—but you can use these classes without knowing their internals.

Generally, hash sets are a bit more efficient, provided you have a good *hash function* for your elements. Library classes such as `String` or `Path` have good hash functions. You learned how to write hash function for your own classes in [Chapter 4](#).

For example, that set of bad words can be implemented simply as

[Click here to view code image](#)

```
Set<String> badWords = new HashSet<>();
badWords.add("sex");
badWords.add("drugs");
badWords.add("c++");
if (badWords.contains(username.toLowerCase()))
    System.out.println("Please choose a different user name");
```

You use a `TreeSet` if you want to traverse the set in sorted order. One reason you might want to do this is to present users a sorted list of choices.

The element type of the set must implement the `Comparable` interface, or you need to supply a `Comparator` in the constructor.

[Click here to view code image](#)

```
TreeSet<String> countries = new TreeSet<>(); // Visits added countries in sorted
order
countries = new TreeSet<>((u, v) ->
    u.equals(v) ? 0
    : u.equals("USA") ? -1
    : v.equals("USA") ? 1
    : u.compareTo(v));
// USA always comes first
```

The `TreeSet` class implements the `SortedSet` and `NavigableSet`

interfaces, whose methods are shown in [Tables 7-4](#) and [7-5](#).

**Table 7-4** SortedSet<E> Methods

Method	Description
E first() E last()	The first and last element in this set.
SortedSet<E> headSet(E toElement)	
SortedSet<E> subSet(E fromElement, E toElement)	Returns a view of the elements starting at fromElement and ending before toElement.
SortedSet<E> tailSet(E fromElement)	

**Table 7-5** NavigableSet<E> Methods

Method	Description
E higher(E e) E ceiling(E e) E floor(E e) E lower(E e)	Returns the closest element $> \geq < e$ .
E pollFirst() E pollLast()	Removes and returns the first or last element, or returns null if the set is empty.
NavigableSet<E> headSet(E toElement, boolean inclusive) NavigableSet<E> subSet(E fromElement, boolean fromInclusive, E toElement, boolean toExclusive) NavigableSet<E> tailSet(E fromElement, boolean inclusive)	Returns a view of the elements from fromElement to toElement (inclusive or exclusive).

## 7.4 Maps

Maps store associations between keys and values. Call `put` to add a new association, or change the value of an existing key:

[Click here to view code image](#)

```
Map<String, Integer> counts = new HashMap<>();
counts.put("Alice", 1); // Adds the key/value pair to the map
counts.put("Alice", 2); // Updates the value for the key
```

This example uses a hash map which, as for sets, is usually the better choice if you don't need to visit the keys in sorted order. If you do, use a `TreeMap` instead.

Here is how you can get the value associated with a key:

[Click here to view code image](#)

```
int count = counts.get("Alice");
```

If the key isn't present, the `get` method returns `null`. In this example, that would cause a `NullPointerException` when the value is unboxed. A better alternative is

[Click here to view code image](#)

```
int count = counts.getDefault("Alice", 0);
```

Then a count of `0` is returned if the key isn't present.

When you update a counter in a map, you first need to check whether the counter is present, and if so, add 1 to the existing value. The `merge` method simplifies that common operation. The call

[Click here to view code image](#)

```
counts.merge(word, 1, Integer::sum);
```

associates `word` with 1 if the key wasn't previously present, and otherwise combines the previous value and 1, using the `Integer::sum` function.

[Table 7-6](#) summarizes the map operations.

**Table 7-6** Map<K, V> Methods

Method	Description
<pre>V get(Object key) V getDefault(Object key, V defaultValue)</pre>	If <code>key</code> is associated with a non- <code>null</code> value <code>v</code> , returns <code>v</code> . Otherwise, returns <code>null</code> or <code>defaultValue</code> .
<pre>V put(K key, V value)</pre>	If <code>key</code> is associated with a non- <code>null</code> value <code>v</code> , associates <code>key</code> with <code>value</code> and returns <code>v</code> .

	Otherwise, adds entry and returns <code>null</code> .
<code>V putIfAbsent(K key, V value)</code>	If <code>key</code> is associated with a non- <code>null</code> value <code>v</code> , ignores <code>value</code> and returns <code>v</code> . Otherwise, adds entry and returns <code>null</code> .
<code>V merge(K key, V value, BiFunction&lt; ? super V, ? super V, ? extends V&gt; remappingFunction)</code>	If <code>key</code> is associated with a non- <code>null</code> value <code>v</code> , applies the function to <code>v</code> and <code>value</code> and either associates <code>key</code> with the result or, if the result is <code>null</code> , removes the key. Otherwise, associates <code>key</code> with <code>value</code> . Returns <code>get(key)</code> .
<code>V compute(K key, BiFunction&lt; ? super K, ? super V, ? extends V&gt; remappingFunction)</code>	Applies the function to <code>key</code> and <code>get(key)</code> . Either associates <code>key</code> with the result or, if the result is <code>null</code> , removes the key. Returns <code>get(key)</code> .
<code>V computeIfPresent(K key, BiFunction&lt; ? super K, ? super V, ? extends V&gt; remappingFunction)</code>	If <code>key</code> is associated with a non- <code>null</code> value <code>v</code> , applies the function to <code>key</code> and <code>v</code> and either associates <code>key</code> with the result or, if the result is <code>null</code> , removes the key. Returns <code>get(key)</code> .
<code>V computeIfAbsent(K key, Function&lt; ? super K, ? extends V&gt; mappingFunction)</code>	Applies the function to <code>key</code> unless <code>key</code> is associated with a non- <code>null</code> value. Either associates <code>key</code> with the result or, if the result is <code>null</code> , removes the key. Returns <code>get(key)</code> .
<code>void putAll(Map&lt;? extends K, ? extends V&gt; m)</code>	Adds all entries from <code>m</code> .
<code>V remove(Object key)</code>	Removes the key and its associated value, or
<code>V replace(K key, V newValue)</code>	replaces the old value. Returns the old value, or <code>null</code> if none existed.
<code>boolean remove(Object key, Object value)</code> <code>boolean replace(K key, V value, V newValue)</code>	Provided that <code>key</code> was associated with <code>value</code> , removes the entry or replaces the old value and returns <code>true</code> . Otherwise, does nothing and returns <code>false</code> . These methods are mainly of interest when the map is accessed concurrently.
<code>int size()</code>	Returns the number of entries.

<code>boolean isEmpty()</code>	Checks if this map is empty.
<code>void clear()</code>	Removes all entries.
<code>void forEach(BiConsumer&lt;? super K, ? super V&gt; action)</code>	Applies the action to all entries.
<code>void replaceAll(BiFunction&lt;? super K, ? super V, ? extends V&gt; function)</code>	Calls the function on all entries. Associates keys with non- <code>null</code> results and removes keys with <code>null</code> results.
<code>boolean containsKey(Object key)</code>	Checks whether the map contains the given key or value.
<code>boolean containsValue(Object value)</code>	
<code>Set&lt;K&gt; keySet()</code>	Returns views of the keys, values, and entries.
<code>Collection&lt;V&gt; values()</code>	
<code>Set&lt;Map.Entry&lt;K, V&gt;&gt; entrySet()</code>	
<code>static Map&lt;K, V&gt; of()</code> <code>static Map&lt;K, V&gt; of(K k1, V v1)</code> <code>static Map&lt;K, V&gt; of(K k1, V v1, K k2, V v2)</code> <code>...</code>	
Yields an unmodifiable map containing up to ten keys and values.	

You can get *views* of the keys, values, and entries of a map by calling these methods:

```
Set<K> keySet()
Set<Map.Entry<K, V>> entrySet()
Collection<V> values()
```

The collections that are returned are not copies of the map data, but they are connected to the map. If you remove a key or entry from the view, then the entry is also removed from the underlying map.

To iterate through all keys and values of a map, you can iterate over the set returned by the `entrySet` method:

[Click here to view code image](#)

```
for (Map.Entry<String, Integer> entry : counts.entrySet()) {  
    String k = entry.getKey();  
    Integer v = entry.getValue();  
    Process k, v  
}
```

Or simply use the `forEach` method:

```
counts.forEach((k, v) -> {  
    Process k, v  
});
```



### Caution

Some map implementations (for example, `ConcurrentHashMap`) disallow `null` for keys or values. And with those that allow it (such as `HashMap`), you need to be very careful if you do use `null` values. A number of map methods interpret a `null` value as an indication that an entry is absent, or should be removed.

---

---



### Tip

Sometimes, you need to present map keys in an order that is different from the sort order. For example, in the JavaServer Faces framework, you specify labels and values of a selection box with a map. Users would be surprised if the choices were sorted alphabetically (Friday, Monday, Saturday, Sunday, Thursday, Tuesday, Wednesday) or in the hash code order. In that case, use a `LinkedHashMap` that remembers the order in which entries were added and iterates through them in that order.

---

---

## 7.5 Other Collections

In the following sections, I briefly discuss some collection classes that you may find useful in practice.

## 7.5.1 Properties

The `Properties` class implements a map that can be easily saved and loaded using a plain text format. Such maps are commonly used for storing configuration options for programs. For example:

[Click here to view code image](#)

```
Properties settings = new Properties();
settings.put("width", "200");
settings.put("title", "Hello, World!");
try (OutputStream out = Files.newOutputStream(path)) {
    settings.store(out, "Program Properties");
}
```

The result is the following file:

```
#Program Properties
#Mon Nov 03 20:52:33 CET 2014
width=200
title=Hello, World\!
```



### Note

As of Java 9, property files are encoded in UTF-8. (Previously, they were encoded in ASCII, with characters greater than `'\u007e'` written as Unicode escapes `\unnnn`.) Comments start with `#` or `!`. A newline in a key or value is written as `\n`. The characters `\`, `#`, `!` are escaped as `\\`, `\#`, `\!`.

---

To load properties from a file, call

[Click here to view code image](#)

```
try (InputStream in = Files.newInputStream(path)) {
    settings.load(in);
}
```

Then use the `getProperty` method to get a value for a key. You can specify a default value used when the key isn't present:

[Click here to view code image](#)

```
String title = settings.getProperty("title", "New Document");
```

---



### Note

For historical reasons, the `Properties` class implements `Map<Object, Object>` even though the values are always strings. Therefore, don't use the `get` method—it returns the value as an `Object`.

---

The `System.getProperties` method yields a `Properties` object with system properties. [Table 7-7](#) describes the most useful ones.

**Table 7-7** Useful System Properties

<b>Property Key</b>	<b>Description</b>
<code>user.dir</code>	The “current working directory” of this virtual machine
<code>user.home</code>	The user's home directory
<code>user.name</code>	The user's account name
<code>java.version</code>	The Java runtime version of this virtual machine
<code>java.home</code>	The home directory of the Java installation
<code>java.class.path</code>	The class path with which this VM was launched
<code>java.io.tmpdir</code>	A directory suitable for temporary files (such as <code>/tmp</code> )
<code>os.name</code>	The name of the operating system (such as <code>Linux</code> )
<code>os.arch</code>	The architecture of the operating system (such as <code>amd64</code> )
<code>os.version</code>	The version of the operating system (such as <code>3.13.0-34-generic</code> )
<code>file.separator</code>	The file separator ( <code>/</code> on Unix, <code>\</code> on Windows)
<code>path.separator</code>	The path separator ( <code>:</code> on Unix, <code>;</code> on Windows)
<code>line.separator</code>	The newline separator ( <code>\n</code> on Unix, <code>\r\n</code> on Windows)

## 7.5.2 Bit Sets

The `BitSet` class stores a sequence of bits. A bit set packs bits into an array of `long` values, so it is more efficient to use a bit set than an array of `boolean` values. Bit sets are useful for sequences of flag bits or to represent sets of non-negative integers, where the `i`th bit is `1` to indicate that `i` is contained in the set.

The `BitSet` class gives you convenient methods for getting and setting individual bits. This is much simpler than the bit-fiddling necessary to store bits



in `int` or `long` variables. There are also methods that operate on all bits together for set operations, such as union and intersection. See [Table 7-8](#) for a complete list. Note that the `BitSet` class is not a collection class—it does not implement `Collection<Integer>`.

**Table 7-8** Methods of the `BitSet` Class

<b>Method</b>	<b>Description</b>
<code>BitSet()</code> <code>BitSet(int nbits)</code>	Constructs a bit set that can initially hold 64, or <code>nbits</code> , bits.
<code>void set(int bitIndex)</code> <code>void set(int fromIndex, int toIndex)</code> <code>void set(int bitIndex, boolean value)</code> <code>void set(int fromIndex, int toIndex, boolean value)</code>	Sets the bit at the given index, or from <code>fromIndex</code> (inclusive) to <code>toIndex</code> (exclusive), to 1 or to the given value.
<code>void clear(int bitIndex)</code> <code>void clear(int fromIndex, int toIndex)</code> <code>void clear()</code>	Sets the bit at the given index, or from <code>fromIndex</code> (inclusive) to <code>toIndex</code> (exclusive), or all bits to 0.
<code>void flip(int bitIndex)</code> <code>void flip(int fromIndex, int toIndex)</code>	Flips the bit at the given index, or from <code>fromIndex</code> (inclusive) to <code>toIndex</code> (exclusive).
<code>boolean get(int bitIndex)</code> <code>BitSet get(int fromIndex, int toIndex)</code>	Gets the bit at the given index, or from <code>fromIndex</code> (inclusive) to <code>toIndex</code> (exclusive).
<code>int nextSetBit(int fromIndex)</code> <code>int previousSetBit(int</code>	

fromIndex) Returns the index of the next/previous 1/0 bit, or  
int nextClearBit(int -1 if none exists.  
fromIndex)  
int  
previousClearBit(int  
fromIndex)

---

void and(BitSet set)  
void andNot(BitSet set) Forms the intersection|difference|union|symmetric  
difference with set.  
void or(BitSet set)  
void xor(BitSet set)

---

int cardinality() Returns the number of 1 bits in this bit set.  
**Caution:** The size method returns the current  
size of the bit vector, not the size of the set.

---

byte[] toByteArray[] Packs the bits of this bit set into an array.  
long[] toLongArray[]

---

IntStream stream() Returns a stream or string of the integers (that is,  
String toString() indexes of 1 bits) in this bit set.

---

static BitSet  
valueOf(byte[]  
bytes)  
static BitSet  
valueOf(long[]  
longs) Yields a bit set containing the supplied bits.  
static BitSet  
valueOf(ByteBuffer  
bb)  
static BitSet  
valueOf(LongBuffer  
lb)

---

boolean isEmpty()  
boolean  
intersects(BitSet set) Checks whether this bit set is empty, or has an  
element in common with set.

---

### 7.5.3 Enumeration Sets and Maps

If you collect sets of enumerated values, use the `EnumSet` class instead of `BitSet`. The `EnumSet` class has no public constructors. Use a static factory method to construct the set:

[Click here to view code image](#)

```
enum Weekday { MONDAY, TUESDAY, WEDNESDAY, THURSDAY, FRIDAY, SATURDAY,
SUNDAY };
Set<Weekday> always = EnumSet.allOf(Weekday.class);
Set<Weekday> never = EnumSet.noneOf(Weekday.class);
Set<Weekday> workday = EnumSet.range(Weekday.MONDAY, Weekday.FRIDAY);
Set<Weekday> mwf = EnumSet.of(Weekday.MONDAY, Weekday.WEDNESDAY,
Weekday.FRIDAY);
```

You can use the methods of the `Set` interface to work with an `EnumSet`.

An `EnumMap` is a map with keys that belong to an enumerated type. It is implemented as an array of values. You specify the key type in the constructor:

[Click here to view code image](#)

```
EnumMap<Weekday, String> personInCharge = new EnumMap<>
(Weekday.class);
personInCharge.put(Weekday.MONDAY, "Fred");
```

## 7.5.4 Stacks, Queues, Deques, and Priority Queues

A stack is a data structure for adding and removing elements at one end (the “top” of the stack). A queue lets you efficiently add elements at one end (the “tail”) and remove them from the other end (the “head”). A double-ended queue, or deque, supports insertion and removal at both ends. With all these data structures, adding elements in the middle is not supported.

The `Queue` and `Deque` interfaces define the methods for these data structures. There is no `Stack` interface in the Java collections framework, just a legacy `Stack` class from the earliest days of Java that you should avoid. If you need a stack, queue, or deque and are not concerned about thread safety, use an `ArrayDeque`.

With a stack, use the `push` and `pop` methods.

[Click here to view code image](#)

```
ArrayDeque<String> stack = new ArrayDeque<>();
stack.push("Peter");
stack.push("Paul");
stack.push("Mary");
while (!stack.isEmpty())
System.out.println(stack.pop());
```

With a queue, use `add` and `remove`.

[Click here to view code image](#)

```
Queue<String> queue = new ArrayDeque<>();
queue.add("Peter");
queue.add("Paul");
queue.add("Mary");
while (!queue.isEmpty())
    System.out.println(queue.remove());
```

Threadsafe queues are commonly used in concurrent programs. You will find more information about them in [Chapter 10](#).

A *priority queue* retrieves elements in sorted order after they were inserted in arbitrary order. That is, whenever you call the `remove` method, you get the smallest element currently in the priority queue.

A typical use for a priority queue is job scheduling. Each job has a priority. Jobs are added in random order. Whenever a new job can be started, the highest priority job is removed from the queue. (Since it is traditional for priority 1 to be the “highest” priority, the `remove` operation yields the minimum element.)

[Click here to view code image](#)

```
public class Job implements Comparable<Job> { ... }
...
PriorityQueue<Job> jobs = new PriorityQueue<>();
jobs.add(new Job(4, "Collect garbage"));
jobs.add(new Job(9, "Match braces"));
jobs.add(new Job(1, "Fix memory leak"));
...
while (jobs.size() > 0) {
    Job job = jobs.remove(); // The most urgent jobs are removed first
    execute(job);
}
```

Just like a `TreeSet`, a priority queue can hold elements of a class that implements the `Comparable` interface, or you can supply a `Comparator` in the constructor. However, unlike a `TreeSet`, iterating over the elements does not necessarily yield them in sorted order. The priority queue uses algorithms for adding and removing elements that cause the smallest element to gravitate to the root, without wasting time on sorting all elements.

### 7.5.5 Weak Hash Maps

The `WeakHashMap` class was designed to solve an interesting problem. What happens with a value whose key is no longer used anywhere in your program? If

the last reference to a key has gone away, there is no longer any way to refer to the value object so it should be removed by the garbage collector.

It isn't quite so simple. The garbage collector traces live objects. As long as the map object is live, all entries in it are live and won't be reclaimed—and neither will be the values that are referenced by the entries.

This is the problem that the `WeakHashMap` class solves. This data structure cooperates with the garbage collector to remove key/value pairs when the only reference to the key is the one from the hash table entry.

Technically, the `WeakHashMap` uses *weak references* to hold keys. A `WeakReference` object holds a reference to another object—in our case, a hash table key. Objects of this type are treated in a special way by the garbage collector. If an object is reachable *only* by a weak reference, the garbage collector reclaims the object and places the weak reference into a queue associated with the `WeakReference` object. Whenever a method is invoked on it, a `WeakHashMap` checks its queue of weak references for new arrivals and removes the associated entries.

## 7.6 Views

A collection *view* is a lightweight object that implements a collection interface, but doesn't store elements. For example, the `keySet` and `values` methods of a map yield views into the map.

In the following sections, you will see some views that are provided by the Java collections framework.

### 7.6.1 Small Collections

The `List`, `Set`, and `Map` interfaces provide static methods yielding a set or list with given elements, and a map with given key/value pairs.

For example,

[Click here to view code image](#)

```
List<String> names = List.of("Peter", "Paul", "Mary");  
Set<Integer> numbers = Set.of(2, 3, 5);
```

yield a list and a set with three elements. For a map, specify the keys and values like this:

[Click here to view code image](#)

```
Map<String, Integer> scores = Map.of("Peter", 2, "Paul", 3, "Mary",  
5);
```

The elements, keys, or values may not be `null`.

The `List` and `Set` interfaces have 11 `of` methods with zero to ten arguments, and an `of` method with a variable number of arguments. The specializations are provided for efficiency.

For the `Map` interface, it is not possible to provide a version with variable arguments since the argument types alternate between the key and value types. There is a static method `ofEntries` that accepts an arbitrary number of `Map.Entry<K, V>` objects, which you can create with the static `entry` method. For example,

[Click here to view code image](#)

```
import static java.util.Map.*;
...
Map<String, Integer> scores = ofEntries(
    entry("Peter", 2),
    entry("Paul", 3),
    entry("Mary", 5));
```

The `of` and `ofEntries` methods produce objects of classes that have an instance variable for each element, or that are backed by an array.

These collection objects are *unmodifiable*. Any attempt to change their contents results in an `UnsupportedOperationException`.

If you want a mutable collection, you can pass the unmodifiable collection to the constructor:

[Click here to view code image](#)

```
List<String> names = new ArrayList<>(List.of("Peter", "Paul",
"Mary"));
```



### Note

There is also a static `Arrays.asList` method that is similar to `List.of`. It returns a mutable list that is not resizable. That is, you can call `set` but not `add` or `remove` on the list.

---

## 7.6.2 Ranges

You can form a sublist view of a list. For example,

[Click here to view code image](#)

```
List<String> sentence = ...;
```

```
List<String> nextFive = sentence.subList(5, 10);
```

This view accesses the elements with index 5 through 9. Any mutations of the sublist (such as setting, adding, or removing elements) affect the original.

For sorted sets and maps, you specify a range by the lower and upper bound:

[Click here to view code image](#)

```
TreeSet<String> words = ...;  
SortedSet<String> asOnly = words.subSet("a", "b");
```

As with `subList`, the first bound is inclusive, and the second exclusive.

The `headSet` and `tailSet` methods yield a subrange with no lower or upper bound.

[Click here to view code image](#)

```
NavigableSet<String> nAndBeyond = words.tailSet("n");
```

With the `NavigableSet` interface, you can choose for each bound whether it should be inclusive or exclusive—see [Table 7-5](#).

For a sorted map, there are equivalent methods `subMap`, `headMap`, and `tailMap`.

### 7.6.3 Unmodifiable Views

Sometimes, you want to share the contents of a collection but you don't want it to be modified. Of course, you could copy the values into a new collection, but that is potentially expensive. An unmodifiable view is a better choice. Here is a typical situation. A `Person` object maintains a list of friends. If the `getFriends` gave out a reference to that list, a caller could mutate it. But it is safe to provide an unmodifiable list view:

[Click here to view code image](#)

```
public class Person {  
    private ArrayList<Person> friends;  
  
    public List<Person> getFriends() {  
        return Collections.unmodifiableList(friends);  
    }  
    ...  
}
```

All mutator methods throw an exception when they are invoked on an unmodifiable view.

As you can see from [Table 7-3](#), you can get unmodifiable views as collections,

lists, sets, sorted sets, navigable sets, maps, sorted maps, and navigable maps.

---



In [Chapter 6](#), you saw how it is possible to smuggle the wrong kind of elements into a generic collection (a phenomenon called “heap pollution”), and that a runtime error is reported when the inappropriate element is retrieved, not when it is inserted. If you need to debug such a problem, use a *checked view*. Where you constructed, say, an `ArrayList<String>`, instead use

[Click here to view code image](#)

```
List<String> strings  
= Collections.checkedList(new ArrayList<>(), String.class);
```

The view monitors all insertions into the list and throws an exception when an object of the wrong type is added.

---

---



The `Collections` class produces *synchronized* views that ensure safe concurrent access to data structures. In practice, these views are not as useful as the data structures in the `java.util.concurrent` package that were explicitly designed for concurrent access. I suggest you use those classes and stay away from synchronized views.

---

---

## Exercises

1. Implement the “Sieve of Erathostenes” algorithm to determine all prime numbers  $\leq n$ . Add all numbers from 2 to  $n$  to a set. Then repeatedly find the smallest element  $s$  in the set, and remove  $s^2$ ,  $s \cdot (s + 1)$ ,  $s \cdot (s + 2)$ , and so on. You are done when  $s^2 > n$ . Do this with both a `HashSet<Integer>` and a `BitSet`.
2. In an array list of strings, make each string uppercase. Do this with (a) an iterator, (b) a loop over the index values, and (c) the `replaceAll` method.
3. How do you compute the union, intersection, and difference of two sets, using just the methods of the `Set` interface and without using loops?
4. Produce a situation that yields a