

```

10     return -1;
11     else if (area1 == area2)
12         return 0;
13     else
14         return 1;
15 }
16 }

```

If you create a `TreeSet` using its no-arg constructor, the `compareTo` method orders the elements in the set, assuming that the class of the elements implements the `Comparable` interface. To use a comparator, you have to use the constructor `TreeSet(Comparator)` to create a sorted set that uses the `compare` method in the comparator to order elements in the set.

Listing 22.6 gives a program that demonstrates how to sort elements in a `TreeSet` using a `Comparator` interface. The example creates a tree set of geometric objects. The elements in the set are sorted using the `compare` method in the `Comparator` interface.

LISTING 22.6 TestTreeSetWithComparator.java

```

1 import java.util.*;
2
3 public class TestTreeSetWithComparator {
4     public static void main(String[] args) {
5         // Create a tree set for geometric objects using a
tree set
6         Set<GeometricObject> set =
7         new TreeSet<GeometricObject>(new GeometricObject
8         set.add(new Rectangle(4, 5));
9         set.add(new Circle(40));
10        set.add(new Circle(40));
11        set.add(new Rectangle(4, 1));
12
13        // Display geometric objects in the tree set
display elements
14        System.out.println("A sorted set of geometric objects");
15        for (GeometricObject element : set)
16            System.out.println("area = " + element.getArea());
17    }
18 }

```



```

A sorted set of geometric objects
area = 4.0
area = 20.0
area = 5022.548245743669

```

The `Circle` and `Rectangle` classes were defined in §14.2, “Abstract `Comparable` Subclasses of `GeometricObject`.”

Two circles of the same radius are added to the set in the tree set (lines 9 and 10). Only one circle is stored, because the two circles are equal and the set does not allow duplicates.



Note

Comparable vs. Comparator

`Comparable` is used to compare the objects of the class that implements `Comparable`. `Comparator` can be used to compare the objects of the class that doesn't implement `Comparable`.

22.6 Lists

A set stores nonduplicate elements. To allow duplicate elements to be stored, you need to use a list. A list can not only store duplicate elements but also

specify where they are stored. The user can access elements by an index. The `List` interface extends `Collection` to define an ordered collection with duplicates allowed. The `List` interface adds position-oriented operations, as well as a new list iterator that enables the user to traverse the list bidirectionally. The new methods in the `List` interface are shown in Figure 22.4.

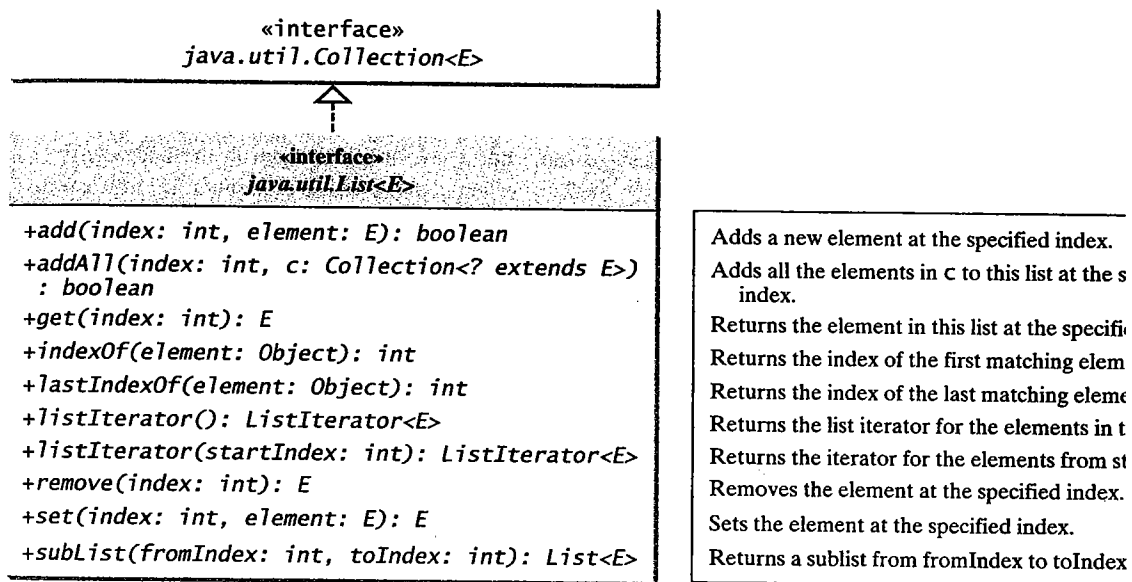


FIGURE 22.4 The `List` interface stores elements in sequence, permitting duplicates.

The `add(index, element)` method is used to insert an element at a specified index, and the `addAll(index, collection)` method to insert a collection at a specified index. The `remove(index)` method is used to remove an element at the specified index from the list. A new element can be set at the specified index using the `set(index, element)` method.

The `indexOf(element)` method is used to obtain the index of the specified element's first occurrence in the list, and the `lastIndexOf(element)` method to obtain the index of its last occurrence. A sublist can be obtained by using the `subList(fromIndex, toIndex)` method.

The `listIterator()` or `listIterator(startIndex)` method returns an instance of `ListIterator`. The `ListIterator` interface extends the `Iterator` interface to add bidirectional traversal of the list. The methods in `ListIterator` are listed in Figure 22.5.

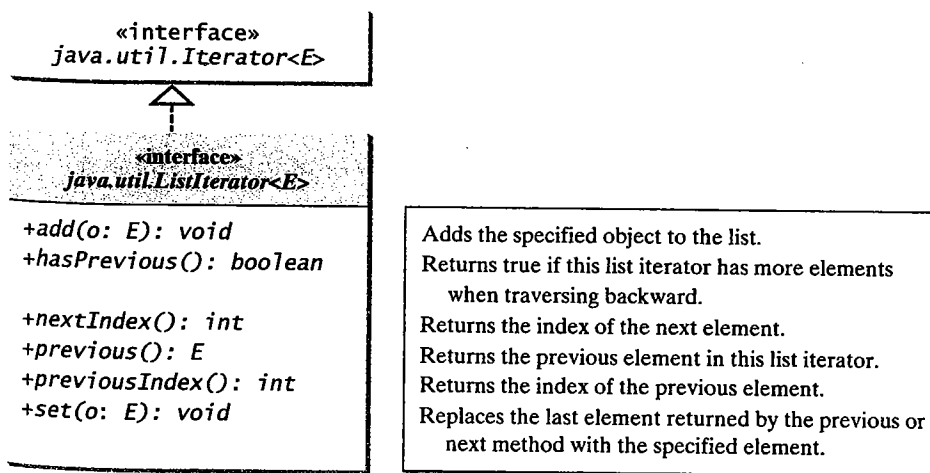


FIGURE 22.5 `ListIterator` enables traversal of a list bidirectionally.

The `add(element)` method inserts the specified element into the list, inserted immediately before the next element that would be returned by the `next()` method defined in the `Iterator` interface, if any, and after the element that would be returned by the `previous()` method, if any. If the list contains no elements, the new element is the sole element on the list. The `set(element)` method can be used to replace the element returned by the `next()` method or the `previous()` method with the specified element.

The `hasNext()` method defined in the `Iterator` interface is used to check whether the iterator has more elements when traversed in the forward direction, and the `hasPrevious()` method is used to check whether the iterator has more elements when traversed in the reverse direction.

The `next()` method defined in the `Iterator` interface returns the next element in the list, and the `previous()` method returns the previous element. The `nextIndex()` method returns the index of the next element in the list, and the `previousIndex()` method returns the index of the previous element in the list.

The `AbstractList` class provides a partial implementation for the `List` interface. The `AbstractSequentialList` class extends `AbstractList` to provide a partial implementation for the `SequentialList` interface.

22.6.1 The ArrayList and LinkedList Classes

The `ArrayList` class (introduced in §11.11) and the `LinkedList` class are two implementations of the `List` interface. `ArrayList` stores elements in an array that is dynamically created. If the capacity of the array is exceeded, a larger array is created, and all the elements from the current array are copied to the new array. `LinkedList` stores elements in a linked list. Which of the two classes you use depends on your application. If you need to support random access through an index without inserting or deleting elements except at the end, `ArrayList` offers the most efficient collection. If you need to support the insertion or deletion of elements anywhere in the list, `LinkedList` is more efficient. A list can grow or shrink dynamically. Once it is created, your application does not require the insertion or deletion of elements. `ArrayList` is an efficient data structure.

`ArrayList` is a resizable-array implementation of the `List` interface. It provides methods for manipulating the size of the array used internally to store the elements in the list. Figure 22.6. Each `ArrayList` instance has a capacity, which is the number of elements that can be stored in the list. It is always at least as large as the number of elements currently stored in the list. When an element is added to an `ArrayList`, its capacity grows automatically. An `ArrayList` can be constructed using the `ArrayList(Collection)` constructor, the `ArrayList(int)` constructor, or the `ArrayList()` constructor. You can use the `trimToSize()` method to reduce the capacity of the list to its current size.

`trimToSize()`

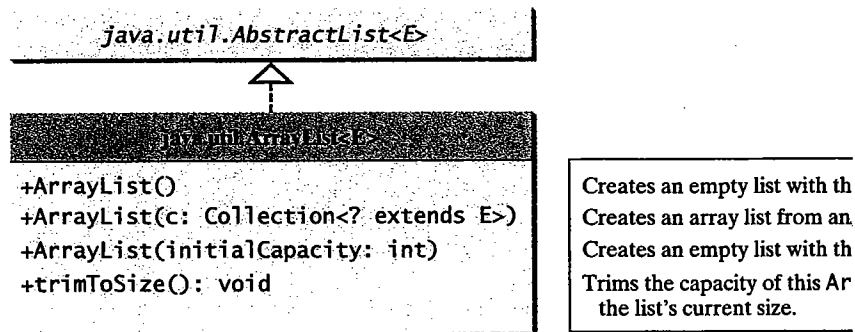


FIGURE 22.6 ArrayList implements List using an array.

`LinkedList` is a linked list implementation of the `List` interface. In addition to implementing the `List` interface, this class provides the methods for retrieving, inserting, and removing elements from both ends of the list, as shown in Figure 22.7. A `LinkedList` can be constructed using its no-arg constructor or `LinkedList(Collection)`.

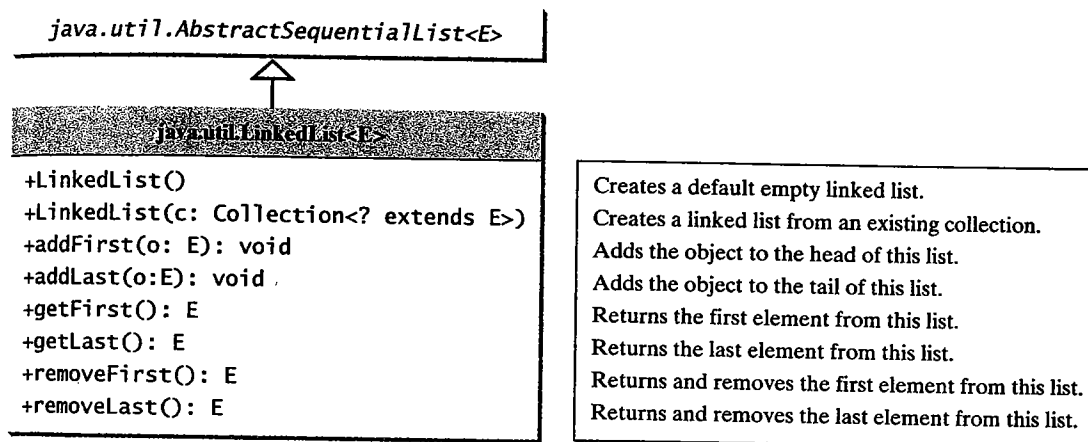


FIGURE 22.7 `LinkedList` provides methods for adding and inserting elements at both ends of the list.

Listing 22.7 gives a program that creates an array list filled with numbers and inserts new elements into specified locations in the list. The example also creates a linked list from the array list and inserts and removes elements from the list. Finally, the example traverses the list forward and backward.

LISTING 22.7 `TestArrayAndLinkedList.java`

```

1 import java.util.*;
2
3 public class TestArrayAndLinkedList {
4     public static void main(String[] args) {
5         List<Integer> arrayList = new ArrayList<Integer>();           array list
6         arrayList.add(1); // 1 is autoboxed to new Integer(1)
7         arrayList.add(2);
8         arrayList.add(3);
9         arrayList.add(1);
10        arrayList.add(4);
11        arrayList.add(0, 10);
12        arrayList.add(3, 30);
13
14        System.out.println("A list of integers in the array list:");
15        System.out.println(arrayList);
16
17        LinkedList<Object> linkedList = new LinkedList<Object>(arrayList);   linked list
18        linkedList.add(1, "red");
19        linkedList.removeLast();
20        linkedList.addFirst("green");
21
22        System.out.println("Display the linked list forward:");
23        ListIterator<Object> listIterator = linkedList.listIterator();       list iterator
24        while (listIterator.hasNext()) {
25            System.out.print(listIterator.next() + " ");
26        }
27        System.out.println();
  
```

```

list iterator
28
29     System.out.println("Display the linked list back
30     listIterator = linkedList.listIterator(linkedLi
31     while (listIterator.hasPrevious()) {
32         System.out.print(listIterator.previous() + "
33     }
34 }
35 }

```



A list of integers in the array list:
 [10, 1, 2, 30, 3, 1, 4]
 Display the linked list forward:
 green 10 red 1 2 30 3 1
 Display the linked list backward:
 1 3 30 2 1 red 10 green

A list can hold identical elements. Integer 1 is stored twice in the list (`ArrayList` and `LinkedList` are operated similarly). The critical difference between `ArrayList` and `LinkedList` is their internal implementation, which affects their performance. `ArrayList` is efficient for inserting and removing elements from the end of the list, while `LinkedList` is efficient for inserting and removing elements anywhere in the list.



Tip

Java provides the static `asList` method for creating a list from a variable-length array of a generic type. Thus you can use the following code to create a list of strings:

```

List<String> list1 = Arrays.asList("red", "green",
List<Integer> list2 = Arrays.asList(10, 20, 30, 40

```

`Arrays.asList(T... a)`
 method

22.7 Static Methods for Lists and Collections

You can use `TreeSet` to store sorted elements in a set. But there is more to the Java Collections Framework. The `Collections` class provides static methods in the `Collection` interface to be used to sort a list. The `Collections` class also contains the binary search methods `binarySearch`, `shuffle`, `copy`, and `fill` methods on lists, and `max`, `min`, `disjoint` methods on collections, as shown in Figure 22.8.

You can sort the comparable elements in a list in its natural order through the `Comparable` interface. You may also specify a comparator. For example, the following code sorts strings in a list.

```

List<String> list = Arrays.asList("red", "green", "blue",
Collections.sort(list);
System.out.println(list);

```

The output is [blue, green, red].

The preceding code sorts a list in ascending order. To sort it in descending order, simply use the `Collections.reverseOrder()` method to return a comparator that orders the elements in reverse order. For example, the following code sorts strings in descending order.

```

List<String> list = Arrays.asList("yellow", "red", "green", "blue",
Collections.sort(list, Collections.reverseOrder());
System.out.println(list);

```

The output is [yellow, red, green, blue].

sort list

ascending order
 descending order

22.7 Static Methods for Lists and Coll

java.util.Collections		
	<p>+sort(list: List): void +sort(list: List, c: Comparator): void +binarySearch(list: List, key: Object): int +binarySearch(list: List, key: Object, c: Comparator): int</p>	<p>Sorts the specified list. Sorts the specified list with the comparat Searches the key in the sorted list using Searches the key in the sorted list using with the comparator.</p>
List	<p>+reverse(list: List): void +reverseOrder(): Comparator +shuffle(list: List): void +shuffle(list: List, rnd: Random): void +copy(des: List, src: List): void +nCopies(n: int, o: Object): List +fill(list: List, o: Object): void</p>	<p>Reverses the specified list. Returns a comparator with the reverse o Shuffles the specified list randomly. Shuffles the specified list with a random Copies from the source list to the destin Returns a list consisting of <i>n</i> copies of th Fills the list with the object.</p>
Collection	<p>+max(c: Collection): Object +max(c: Collection, c: Comparator): Object +min(c: Collection): Object +min(c: Collection, c: Comparator): Object +disjoint(c1: Collection, c2: Collection): boolean +frequency(c: Collection, o: Object): int</p>	<p>Returns the max object in the collection Returns the max object using the compa Returns the min object in the collection Returns the min object using the compa Returns true if c1 and c2 have no elem Returns the number of occurrences of th element in the collection.</p>

FIGURE 22.8 The Collections class contains static methods for manipulating lists and collections.

You can use the `binarySearch` method to search for a key in a list. The list must be presorted in increasing order. If the key is not in the list, the method returns $-(insertion\ point + 1)$. Recall that the insertion point is where the item would fall in the list if it were present. For example, the following code searches the keys in a list of integers and a list of strings.

```

List<Integer> list1 =
    Arrays.asList(2, 4, 7, 10, 11, 45, 50, 59, 60, 66);
System.out.println("(1) Index: " + Collections.binarySearch(list1, 7));
System.out.println("(2) Index: " + Collections.binarySearch(list1, 9));

List<String> list2 = Arrays.asList("blue", "green", "red");
System.out.println("(3) Index: " +
    Collections.binarySearch(list2, "red"));
System.out.println("(4) Index: " +
    Collections.binarySearch(list2, "cyan"));
    
```

The output of the preceding code is

```

(1) Index: 2
(2) Index: -4
(3) Index: 2
(4) Index: -2
    
```

You can use the `reverse` method to reverse the elements in a list. For example, the following code displays `[blue, green, red, yellow]`.

```

List<String> list = Arrays.asList("yellow", "red", "green", "blue");
Collections.reverse(list);
System.out.println(list);
    
```

shuffle

You can use the `shuffle(List)` method to randomly reorder the elements in a list. For example, the following code shuffles the elements in `list`.

```
List<String> list = Arrays.asList("yellow", "red", "green");
Collections.shuffle(list);
System.out.println(list);
```

You can also use the `shuffle(List, Random)` method to randomly reorder a list with a specified `Random` object. Using a specified `Random` object results in a list with identical sequences of elements for the same original list. For example, the following code shuffles the elements in `list`.

```
List<String> list1 = Arrays.asList("yellow", "red", "green");
List<String> list2 = Arrays.asList("yellow", "red", "green");
Collections.shuffle(list1, new Random(20));
Collections.shuffle(list2, new Random(20));
System.out.println(list1);
System.out.println(list2);
```

You will see that `list1` and `list2` have the same sequence of elements after shuffling.

copy

You can use the `copy(dest, src)` method to copy all the elements from the source list to the destination list on the same index. The destination must be as long as the source list; otherwise, the longer, the remaining elements in the source list are not affected. For example, the following code copies `list2` to `list1`.

```
List<String> list1 = Arrays.asList("yellow", "red", "green");
List<String> list2 = Arrays.asList("white", "black");
Collections.copy(list1, list2);
System.out.println(list1);
```

The output for `list1` is `[white, black, green, blue]`. The `copy` method creates a shallow copy. Only the references of the elements from the source list are copied.

nCopies

You can use the `nCopies(int n, Object o)` method to create an immutable list consisting of `n` copies of the specified object. For example, the following code creates five `Calendar` objects.

```
List<GregorianCalendar> list1 = Collections.nCopies(5, new GregorianCalendar(2005, 0, 1));
```

The list created from the `nCopies` method is immutable, so you cannot update elements in the list. All the elements have the same references.

fill

You can use the `fill(List list, Object o)` method to replace all the elements in the list with the specified element. For example, the following code displays `[black, black, black]`.

```
List<String> list = Arrays.asList("red", "green", "blue");
Collections.fill(list, "black");
System.out.println(list);
```

max and min methods

You can use the `max` and `min` methods for finding the maximum and minimum elements in a collection. The elements must be comparable using the `Comparable` interface. For example, the following code displays the largest and smallest strings in a collection.

```
Collection<String> collection = Arrays.asList("red", "green", "blue");
System.out.println(Collections.max(collection));
System.out.println(Collections.min(collection));
```

22.8 Performance of Sets and Lists

The `disjoint(collection1, collection2)` method returns `true` if the two collections have no elements in common. For example, in the following code, `disjoint(collection1, collection2)` returns `false`, but `disjoint(collection1, collection3)` returns `true`. disjoint

```
Collection<String> collection1 = Arrays.asList("red", "cyan");
Collection<String> collection2 = Arrays.asList("red", "blue");
Collection<String> collection3 = Arrays.asList("pink", "tan");
System.out.println(Collections.disjoint(collection1, collection2));
System.out.println(Collections.disjoint(collection1, collection3));
```

The `frequency(collection, element)` method finds the number of occurrences of the element in the collection. For example, `frequency(collection, "red")` returns 2 in the following code. frequency

```
Collection<String> collection = Arrays.asList("red", "cyan", "red");
System.out.println(Collections.frequency(collection, "red"));
```

22.8 Performance of Sets and Lists

We now conduct an interesting experiment to test the performance of sets and lists. Listing 22.8 gives a program that shows the execution time of adding and removing elements in a hash set, linked hash set, tree set, array list, and linked list.

LISTING 22.8 SetListPerformanceTest.java

```
1 import java.util.*;
2
3 public class SetListPerformanceTest {
4     public static void main(String[] args) {
5         // Create a hash set, and test its performance
6         Collection<Integer> set1 = new HashSet<Integer>();           a hash set
7         System.out.println("Time for hash set is " +
8             getTestTime(set1, 500000) + " milliseconds");
9
10        // Create a linked hash set, and test its performance
11        Collection<Integer> set2 = new LinkedHashSet<Integer>();    a linked has
12        System.out.println("Time for linked hash set is " +
13            getTestTime(set2, 500000) + " milliseconds");
14
15        // Create a tree set, and test its performance
16        Collection<Integer> set3 = new TreeSet<Integer>();         a tree set
17        System.out.println("Time for tree set is " +
18            getTestTime(set3, 500000) + " milliseconds");
19
20        // Create an array list, and test its performance
21        Collection<Integer> list1 = new ArrayList<Integer>();      an array list
22        System.out.println("Time for array list is " +
23            getTestTime(list1, 60000) + " milliseconds");
24
25        // Create a linked list, and test its performance
26        Collection<Integer> list2 = new LinkedList<Integer>();    a linked list
27        System.out.println("Time for linked list is " +
28            getTestTime(list2, 60000) + " milliseconds");
29    }
30
31    public static long getTestTime(Collection<Integer> c, int size) {
32        long startTime = System.currentTimeMillis();              start time
33    }
```



```

34 // Add numbers 0, 1, 2, ..., size - 1 to the array
35 List<Integer> list = new ArrayList<Integer>();
36 for (int i = 0; i < size; i++)
37     list.add(i);
38
shuffle 39 Collections.shuffle(list); // Shuffle the array 1-
40
41 // Add the elements to the container
add to container 42 for (int element: list)
43     c.add(element);
44
shuffle 45 Collections.shuffle(list); // Shuffle the array 1-
46
47 // Remove the element from the container
remove from container 48 for (int element: list)
49     c.remove(element);
50
end time 51 long endTime = System.currentTimeMillis();
return elapsed time 52 return endTime - startTime; // Return the executio
53 }
54 }

```



```

Time for hash set is 1437 milliseconds
Time for linked hash set is 1891 milliseconds
Time for tree set is 2891 milliseconds
Time for array list is 13797 milliseconds
Time for linked list is 15344 milliseconds

```

The `getTestTime` method creates a list of distinct integers from 0 to `size`; shuffles the list (line 39), adds the elements from the list to a container `c`; shuffles the list again (line 45), removes the elements from the container (line 48); and returns the execution time (line 52).

The program creates a hash set (line 6), a linked hash set (line 11), an array list (line 21), and a linked list (line 26). The program obtains the execution times for adding and removing 500,000 elements in the three sets and adding and removing 500,000 elements in the two lists.

As you see, sets are much more efficient than lists. If sets are sufficient for your application, use sets. Furthermore, if no particular order is needed for your application, use sets.

The program tested general remove operations for array lists and linked lists. The complexity is about the same. Please note that linked lists are more efficient for insertion and deletion anywhere in the list except at the end.

22.9 The Vector and Stack Classes

The Java Collections Framework was introduced with Java 2. Several classes were supported earlier, among them the `Vector` and `Stack` classes. These classes do not fit into the Java Collections Framework, but all their old-style methods are supported for compatibility.

`Vector` is the same as `ArrayList`, except that it contains synchronized methods for adding and modifying the vector. Synchronized methods can prevent data corruption when the vector is accessed and modified by two or more threads concurrently. For the majority of applications that do not require synchronization, using `ArrayList` is more efficient than using `Vector`.

The `Vector` class implements the `List` interface. It also has the methods of the original `Vector` class defined prior to Java 2, as shown in Figure 22.9.

sets are better

22.11 Maps

Suppose your program stores a million students and frequently searches for a student using the social security number. An efficient data structure for this task is the *map*. A map is a container that stores the elements along with the keys. The keys are like indexes. In *List*, the indexes are integers. In *Map*, the keys can be any objects. A map cannot contain duplicate keys. Each key maps to one value. A key and its corresponding value form an entry, which is actually stored in a map, as shown in Figure 22.14.

why map?

There are three types of maps: *HashMap*, *LinkedHashMap*, and *TreeMap*. The common features of these maps are defined in the *Map* interface. Their relationship is shown in Figure 22.15.

The *Map* interface provides the methods for querying, updating, and obtaining a collection of values and a set of keys, as shown in Figure 22.16.

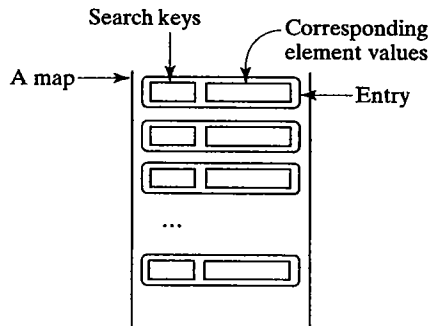


FIGURE 22.14 The entries consisting of key/value pairs are stored in a map.

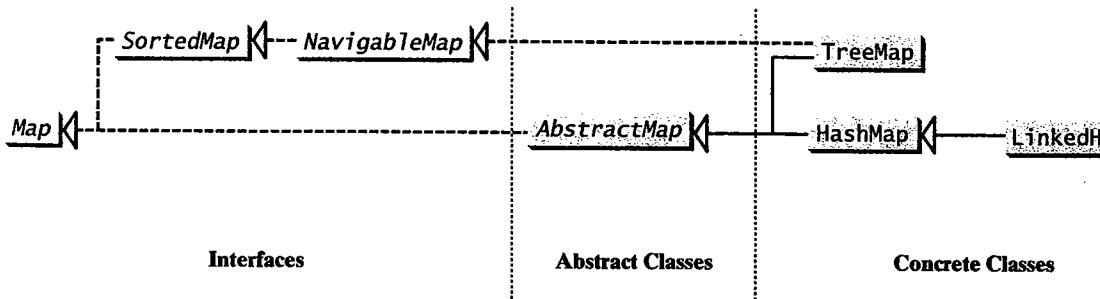


FIGURE 22.15 A map stores key/value pairs.

interface <i>Map</i> <i>extends Map<K, V></i>	
<code>+clear(): void</code>	Removes all entries from this map.
<code>+containsKey(key: Object): boolean</code>	Returns true if this map contains entries for the specified key.
<code>+containsValue(value: Object): boolean</code>	Returns true if this map maps one or more keys to the specified value.
<code>+entrySet(): Set<Map.Entry<K, V>></code>	Returns a set consisting of the entries in this map.
<code>+get(key: Object): V</code>	Returns the value for the specified key in this map.
<code>+isEmpty(): boolean</code>	Returns true if this map contains no entries.
<code>+keySet(): Set<K></code>	Returns a set consisting of the keys in this map.
<code>+put(key: K, value: V): V</code>	Puts a mapping in this map.
<code>+putAll(m: Map<? extends K, ? extends V>): void</code>	Adds all the entries from <i>m</i> to this map.
<code>+remove(key: Object): V</code>	Removes the entries for the specified key.
<code>+size(): int</code>	Returns the number of entries in this map.
<code>+values(): Collection<V></code>	Returns a collection consisting of the values in this map.

FIGURE 22.16 The Map interface maps keys to values.

update methods

The *update methods* include `clear`, `put`, `putAll`, and `remove`. The `remove` method removes all entries from the map. The `put(K key, V value)` method adds a key in the map. If the map formerly contained a mapping for this key, the previous value associated with the key is returned. The `putAll(Map m)` method adds the mappings from the map `m` to the current map. The `remove(Object key)` method removes the map elements for the specified key from the map.

query methods

The *query methods* include `containsKey`, `containsValue`, `isEmpty`, `size`, `keySet`, `values`, and `entrySet`. The `containsKey(Object key)` method checks whether the map contains the specified key. The `containsValue(Object value)` method checks whether the map contains a mapping for this value. The `isEmpty()` method checks whether the map is empty. The `size()` method returns the number of mappings in the map.

`keySet()`

`values()`

`entrySet()`

You can obtain a set of the keys in the map using the `keySet()` method. The `values()` method returns a collection of the values in the map using the `values()` method. The `entrySet()` method returns a set of objects that implement the `Map.Entry<K, V>` interface, where `Entry` is an interface for the `Map` interface, as shown in Figure 22.17. Each object in the set is a pair in the underlying map.

```

interface
java.util.Map.Entry<K,V>
+getKey(): K
+getValue(): V
+setValue(value: V): void
    
```

Returns the key corresponding to this entry.
Returns the value corresponding to this entry.
Replaces the value in this entry with the specified value.

FIGURE 22.17 The `Map.Entry` interface operates on an entry in the map.

`AbstractMap`

The `AbstractMap` class is a convenience class that implements all the methods of the `Map` interface except the `entrySet()` method.

The `SortedMap` interface extends the `Map` interface to maintain the natural order of keys with additional methods `firstKey()` and `lastKey()` for returning the lowest and highest key, `headMap(toKey)` for returning the portion of the map with keys less than `toKey`, and `tailMap(fromKey)` for returning the portion of the map with keys greater than or equal to `fromKey`.

concrete implementation

The `HashMap`, `LinkedHashMap`, and `TreeMap` classes are three concrete implementations of the `Map` interface, as shown in Figure 22.18.

`HashMap`

The `HashMap` class is efficient for locating a value, inserting a mapping, and deleting a mapping.

`LinkedHashMap`

`LinkedHashMap` extends `HashMap` with a linked-list implementation that maintains the order of the entries in the map. The entries in a `HashMap` are not ordered, but the entries in a `LinkedHashMap` can be retrieved either in the order in which they were inserted into the map (known as the *insertion order*) or in the order in which they were last accessed (known as the *access order*). The no-arg constructor of `LinkedHashMap` uses the insertion order. To construct a `LinkedHashMap` with the access order, use the `LinkedHashMap(int initialCapacity, float loadFactor, boolean accessOrder)` constructor.

insertion order
access order

`TreeMap`

The `TreeMap` class is efficient for traversing the keys in a sorted order. The keys in a `TreeMap` are sorted using the `Comparable` interface or the `Comparator` interface. If you use the no-arg constructor, the `compareTo` method in the `Comparable` interface is used to compare the elements in the map, assuming that the class of the elements implements the `Comparable` interface. To use a comparator, you have to use the `TreeMap(Comparator c)` constructor to create a sorted map that uses the comparator to order the elements in the map based on the keys.

`SortedMap`

`SortedMap` is a subinterface of `Map`, which guarantees that the entries in the map are sorted. Additionally, it provides the methods `firstKey()` and `lastKey()` for returning the first and last keys in the map, and `headMap(toKey)` and `tailMap(fromKey)` for returning the portion of the map with keys less than `toKey` and the portion of the map with keys greater than or equal to `fromKey`.

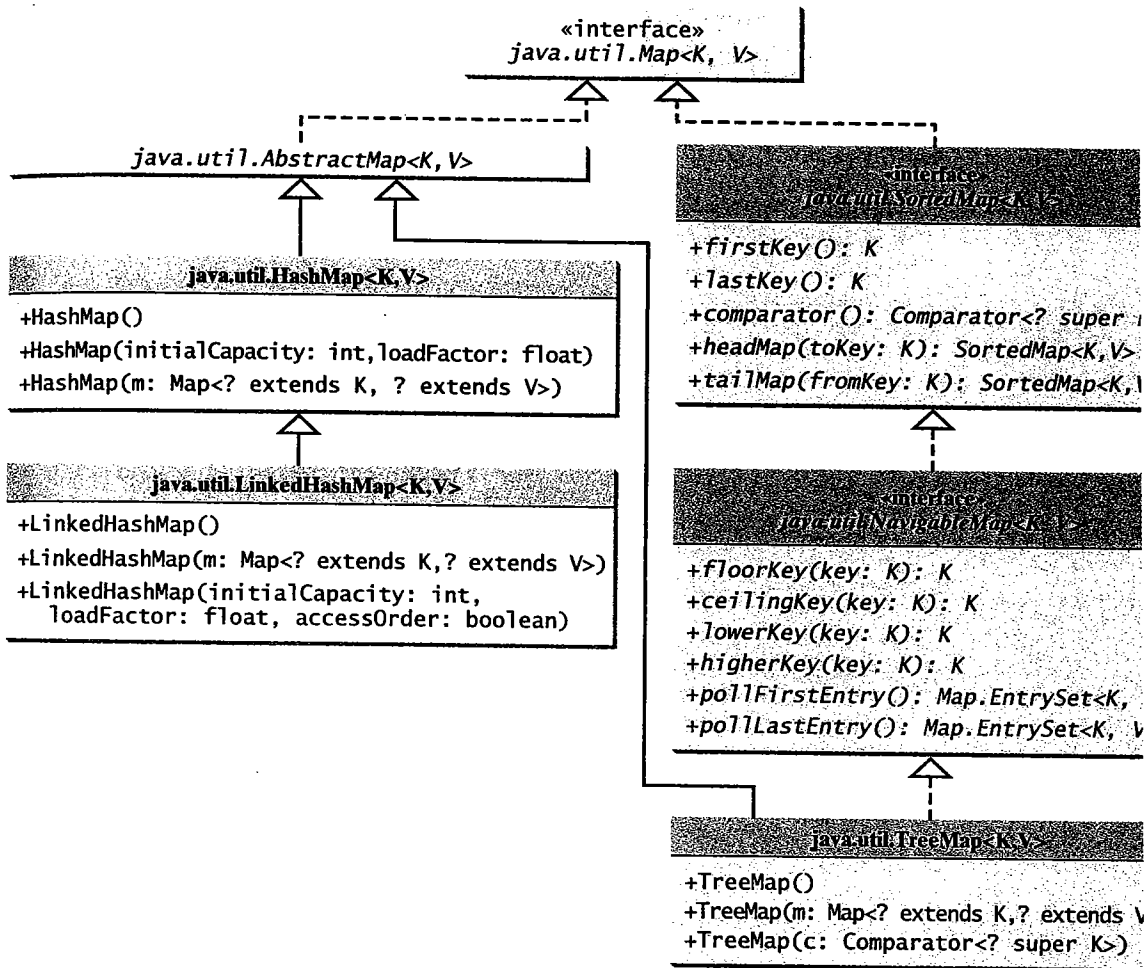


FIGURE 22.18 The Java Collections Framework provides three concrete map classes.

returning a portion of the map whose keys are less than `toKey` and greater than or equal to `fromKey`.

`NavigableMap` extends `SortedMap` to provide navigation methods `lowerKey(key)`, `floorKey(key)`, `ceilingKey(key)`, and `higherKey(key)` that return keys respectively less than, less than or equal, greater than or equal, and greater than a given key and return null if there is no such key. The `pollFirstEntry()` and `pollLastEntry()` methods remove and return the first and last entry in the tree map, respectively.



Note

Prior to Java 2, `java.util.Hashtable` was used for mapping keys with elements. `Hashtable` was redesigned to fit into the Java Collections Framework with all its methods retained for compatibility. `Hashtable` implements the `Map` interface and is used in the same way as `HashMap`, except that `Hashtable` is synchronized.

Navigable

Hashtable

Listing 22.11 gives an example that creates a hash map, a linked hash map, and a tree map that map students to ages. The program first creates a hash map with the student's name as its key and the age as its value. The program then creates a tree map from the hash map and displays the mappings in ascending order of the keys. Finally, the program creates a linked hash map, adds the same entries to the map, and displays the entries.

LISTING 22.11 TestMap.java

```

1 import java.util.*;
2
3 public class TestMap {
4     public static void main(String[] args) {
5         // Create a HashMap
6         Map<String, Integer> hashMap = new HashMap<String, Integer>(16);
7         hashMap.put("Smith", 30);
8         hashMap.put("Anderson", 31);
9         hashMap.put("Lewis", 29);
10        hashMap.put("Cook", 29);
11
12        System.out.println("Display entries in HashMap");
13        System.out.println(hashMap + "\n");
14
15        // Create a TreeMap from the previous HashMap
16        Map<String, Integer> treeMap =
17            new TreeMap<String, Integer>(hashMap);
18        System.out.println("Display entries in ascending order of key");
19        System.out.println(treeMap);
20
21        // Create a LinkedHashMap
22        Map<String, Integer> linkedHashMap =
23            new LinkedHashMap<String, Integer>(16, 0.75f, true);
24        linkedHashMap.put("Smith", 30);
25        linkedHashMap.put("Anderson", 31);
26        linkedHashMap.put("Lewis", 29);
27        linkedHashMap.put("Cook", 29);
28
29        // Display the age for Lewis
30        System.out.println("The age for " + "Lewis is " +
31            linkedHashMap.get("Lewis").intValue());
32
33        System.out.println("\nDisplay entries in LinkedHashMap");
34        System.out.println(linkedHashMap);
35    }
36 }

```



```

Display entries in HashMap
{Cook=29, Smith=30, Lewis=29, Anderson=31}

Display entries in ascending order of key
{Anderson=31, Cook=29, Lewis=29, Smith=30}
The age for Lewis is 29

Display entries in LinkedHashMap
{Smith=30, Anderson=31, Cook=29, Lewis=29}

```

As shown in the output, the entries in the `HashMap` are in random order. `TreeMap` are in increasing order of the keys. The entries in the `LinkedHashMap` are in order of their access, from least recently accessed to most recently.

All the concrete classes that implement the `Map` interface have at least two constructors: the no-arg constructor that constructs an empty map, and the other constructor that constructs a map from an existing map. Thus `new TreeMap<String, Integer>(hashMap)` constructs a tree map from a hash map.

You can create an insertion-ordered or access-ordered linked hash map. An access-ordered linked hash map is created in lines 22–23. The most recently accessed entry is placed at the end of the map. The entry with the key Lewis is last accessed in line 31, so it is displayed last in line 34.

**Tip**

If you don't need to maintain an order in a map when updating it, use a `HashMap`. When you need to maintain the insertion order or access order in the map, use a `LinkedHashMap`. When you need the map to be sorted on keys, use a `TreeMap`.

22.11.1 Case Study: Occurrences of Words

This case study writes a program that counts the occurrences of words in a text and displays the words and their occurrences in alphabetical order of words. The program uses a `TreeMap` to store an entry consisting of a word and its count. For each word, check whether it is already a key in the map. If not, add to the map an entry with the word as the key and value 1. Otherwise, increase the value for the word (key) by 1 in the map. Assume the words are case insensitive; e.g., Good is treated the same as good.

Listing 22.12 gives the solution to the problem.

LISTING 22.12 CountOccurrenceOfWords.java

```

1 import java.util.*;
2
3 public class CountOccurrenceOfWords {
4     public static void main(String[] args) {
5         // Set text in a string
6         String text = "Good morning. Have a good class. " +
7             "Have a good visit. Have fun!";
8
9         // Create a TreeMap to hold words as key and count as value
10        TreeMap<String, Integer> map = new TreeMap<String, Integer>();
11
12        String[] words = text.split("[ \\n\\t\\r.,;:!?@{}\"'");
13        for (int i = 0; i < words.length; i++) {
14            String key = words[i].toLowerCase();
15
16            if (key.length() > 0) {
17                if (map.get(key) == null) {
18                    map.put(key, 1);
19                }
20                else {
21                    int value = map.get(key).intValue();
22                    value++;
23                    map.put(key, value);
24                }
25            }
26        }
27
28        // Get all entries into a set
29        Set<Map.Entry<String, Integer>> entrySet = map.entrySet();
30
31        // Get key and value from each entry
32        for (Map.Entry<String, Integer> entry: entrySet)
33            System.out.println(entry.getValue() + "\\t" + entry.getKey());
34    }
35 }

```



```

2      a
1      class
1      fun
3      good
3      have
1      morning
1      visit

```

The program creates a `TreeMap` (line 10) to store pairs of words and their counts. The words serve as the keys. Since all elements in the map must be stored as objects, the count is wrapped in an `Integer` object.

The program extracts a word from a text using the `split` method (line 11) of the `String` class (see §9.2.7). For each word extracted, the program checks whether it is a key in the map (line 17). If not, a new pair consisting of the word and its count is stored to the map (line 18). Otherwise, the count for the word is incremented (lines 21–23).

The program obtains the entries of the map in a set (line 29), and traverses the set to display the count and the key in each entry (lines 32–33).

Since the map is a tree map, the entries are displayed in increasing order of the occurrence counts, see Exercise 22.8.

Now sit back and think how you would write this program without using `map`. The program will be longer and more complex. You will find that `map` is a very useful data structure for solving problems such as this.

22.12 Singleton and Unmodifiable Collection and Maps

The `Collections` class contains the static methods for lists and collections. It also provides the methods for creating singleton sets, lists, and maps, and for creating unmodifiable lists, and maps, as shown in Figure 22.19.

The `Collections` class defines three constants: one for an empty set, one for an empty list, and one for an empty map (`EMPTY_SET`, `EMPTY_LIST`, and `EMPTY_MAP`). It provides the `singleton(Object o)` method for creating an immutable set containing one element, the `singletonList(Object o)` method for creating an immutable list containing one element, the `singletonMap(Object key, Object value)` method for creating an immutable map with one entry, the `unmodifiableCollection(Collection c)` method for creating an unmodifiable collection, the `unmodifiableList(List l)` method for creating an unmodifiable list, the `unmodifiableMap(Map m)` method for creating an unmodifiable map, the `unmodifiableSet(Set s)` method for creating an unmodifiable set, the `unmodifiableSortedMap(SortedMap s)` method for creating an unmodifiable sorted map, and the `unmodifiableSortedSet(SortedSet s)` method for creating an unmodifiable sorted set.

Method Signature	Description
<code>+singleton(o: Object): Set</code>	Returns a singleton set containing the element o.
<code>+singletonList(o: Object): List</code>	Returns a singleton list containing the element o.
<code>+singletonMap(key: Object, value: Object): Map</code>	Returns a singleton map with the key and value.
<code>+unmodifiableCollection(c: Collection): Collection</code>	Returns an unmodifiable collection.
<code>+unmodifiableList(list: List): List</code>	Returns an unmodifiable list.
<code>+unmodifiableMap(m: Map): Map</code>	Returns an unmodifiable map.
<code>+unmodifiableSet(s: Set): Set</code>	Returns an unmodifiable set.
<code>+unmodifiableSortedMap(s: SortedMap): SortedMap</code>	Returns an unmodifiable sorted map.
<code>+unmodifiableSortedSet(s: SortedSet): SortedSet</code>	Returns an unmodifiable sorted set.

FIGURE 22.19 The `Collections` class contains the static methods for creating singleton and unmodifiable collections and maps.