Chapter 15

Collection Considerations

Goals

- Distinguish Collection classes, data structures, and ADTs
- Consider three data structures used in this textbook: array, singly linked, and tree
- Observe that a Java interface can be used to specify a type to be implemented with different classes using different data structures

15.1 ADTs, Collection Classes, and Data Structures

A **collection class** is a Java class whose main purpose is to store, retrieve, and manage a collection of objects in some meaningful way. Collection classes have the following characteristics:

- The main responsibility of a collection class is to store a collection of values.
- Elements may be added and removed from the collection.
- A collection class allows clients to access the individual elements.

- A collection class may have search-and-sort operations for locating a particular element.
- Some collections allow duplicate elements while other collections do not.
- Some collections are naturally ordered while other collections are not.

Some collection classes are designed to store large amounts of information where any element may need to be found quickly—to find a phone listing, for example. Other collections are designed to have elements that are added, removed, and changed frequently—an inventory system, for example. Some collections store elements in a first in first out basis—such as a queue of incoming packets on an Internet router. Other collections are used to help build computer systems—a stack that manages method calls, for example.

Collection classes support many operations. The following is a short list of operations that are common to many collection classes:

- **add** Place an object into a collection (insertion point varies).
- **find** Get a reference to an object in a collection so you can send messages to it.
- **remove** Take the object out of the collection (extraction point varies).

The Java class is a convenient way to encapsulate algorithms and store data in one module. In addition to writing the class and method headings, decisions have to be made about what data structures to use.

Data Structures

A **data structure** is a way of storing data on a computer so it can be used efficiently. There are several structures available to programmers for storing data. Some are more appropriate than others, depending on how you need to manage your information. Although not a complete list, here are some of the storage structures you have or will see in this book:

1. arrays (Chapters 7, 8, and 10)
2. linked structure (Chapter 16)

3. hash tables (beyond the scope of this book)

Data can be stored in contiguous memory with arrays or in non-contiguous memory with linked structures. Arrays allow you to reserve memory where each element can be physically located next to its predecessor and successor. Any element can be directly changed and accessed through an index.

```
String[] data = new String[5];
data[0] = "Al";
data[1] = "Di";
data[2] = "Mo";
data[3] = "Xu";
```

data (where data.length == 5):

| [0] | [1] | [2] | [3] | [4] |
|-----|-----|-----|-----|-----|
| "Al" | "Di" | "Mo" | "Xu" | null |

A **singly linked structure** contains nodes that store a reference to an element and a reference to another node. The reference to another node may be null to indicate there are no more elements stored.



You will see that both of these storage mechanisms — arrays and linked structures — implement the same List ADT in subsequent chapters.

You will also see how binary trees store data that can be added, removed, and found more quickly. The following picture indicates a linked structure where each node contains a reference to data (integers are used here to save space) and references to two other nodes, or to nothing (shown as diagonal lines).
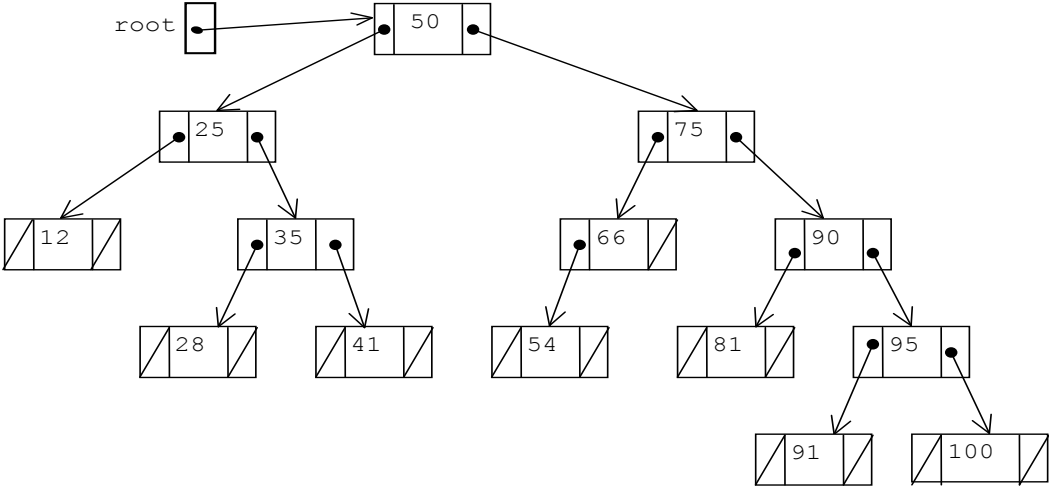


The hash table is another data structure uses a key—such as a student ID, employee number, or Social Security Number—to find the value very quickly.

| Array Index | Key | Value mapped to the key (null or five instance variables shown) |
|---|---|---|
| [0] | "1023" | "Devon" 40.0 10.50 1 'S' |
| [1] | null | null |
| [2] | "5462" | "Al" 42.5 12.00 2 'M' |
| [3] | null | null |
| [4] | "3343" | "Ali" 20.0 9.50 0 'S' |
| ... | | |
| [753] | "0930" | "Chris" 0.0 13.50 1 'S' |

Although hash tables are beyond the scope of this book, you will see collection classes using the other three data structures to implement abstract data types with a java class.

## Abstract Data Types

An **abstract data type** (ADT) describes a set of data values and associated operations that are precisely specified independent of any particular implementation. An abstract data type can be specified using axiomatic semantics. For example, here is the Bag ADT as described by the National Institute of Standards and Technology (NIST)[1].

---

[1]  http://www.nist.gov/dads/HTML/bag.html

**Bag**

**Definition:** An unordered collection of values that may have duplicates.

**Formal Definition:** A bag has a single query function, occurencesOf(v, B), which tells how many copies of an element are in the bag, and two modifier functions, add(v, B) and remove(v, B). These may be defined with axiomatic semantics as follows.

1. new() returns a bag
2. occurencesOf(v, new()) = 0
3. occurencesOf(v, add(v, B)) = 1 + occurencesOf(v, B) [2]
4. occurencesOf(v, add(u, B)) = occurencesOf(v, B) if $v \neq u$
5. remove(v, new()) = new()
6. remove(v, add(v, B)) = B
7. remove(v, add(u, B)) = add(u, remove(v, B)) if $v \neq u$

where B is a bag and u and v are elements.

The predicate isEmpty(B) may be defined with the following additional axioms:

8. isEmpty(new()) = true
9. isEmpty(add(v, B)) = false

**Also known as** multi-set.
Note: A bag, or multi-set, is a set where values may be repeated. Inserting 2, 1, 2 into an empty set gives the set {1, 2}. Inserting those values into an empty bag gives {1, 2, 2}.

---

[2]    This definition shows a bag B is passed as an argument. In an object-oriented language you send a message to an object of type B as in `aBag.add("New Element);` rather than `add("New Element", aBag);`

Although an abstract data type describes how each operation affects data, an ADT

- does not specify how data will be stored
- does not specify the algorithms needed to accomplish the listed operations
- is not dependent on a particular programming language or computer system.

An ADT can also be described as a Java interface that specifies the operations and JUnit assertions to specify behavior.

## ADTs Described with Java Interfaces and JUnit Assertions

The Java `interface` can also be used to specify an abstract data type. For example, the following Java interface specifies the operations for a Bag ADT as method headings. The `new` operation from NIST's Bag ADT is not included here simply because the Java interface does not allow constructor to be specified in an interface. We will use `new` later when there is some class that implements the interface. The syntactic element <E> requires the class to specify the type of element to be stored.

```java
/**
 * This interface specifies the methods for a Bag ADT. It is designed to be
 * generic so any type of element can be stored.
 */
public interface Bag<E> {

  /**
   * Add element to this Bag.
   * @param element: The element to insert.
   */
  public void add(E element);
```

```java
  /**
   * Return true if no elements have been added to this bag.
   * @return False if there is one or more elements in this bag.
   */
  public boolean isEmpty();

  /**
   * Return how many elements match element according to the equals method of
   * the type specified at construction.
   * @param element The element to count in this Bag.
   */
  public int occurencesOf(E element);

  /**
   * Remove the first occurrence of element and return true if found. Otherwise
   * leave this Bag unchanged.
   * @param element: The element to remove
   * @return true if element was removed, false if element was not equal to any
   *          in this bag.
   */
  public boolean remove(E element);
}
```

Like an ADT specified with axiomatic expressions, an `interface` does not specify the data structure and the algorithms to add, remove, or find elements. Comments and well-named identifiers imply the behavior of the operations; however the behavior can be made more explicit with assertions. For example, the assertions like those shown in the following test method help describe the behavior of `add` and `occurencesOf`. This code assumes that a class named `ArrayBag` implements `interface Bag<E>`.

```java
@Test
public void testAddWithOccurencesOf() {
  Bag<String> names = new ArrayBag<String>();

  // A new bag has no occurrences of any string
  assertEquals(0, names.occurencesOf("NOT here"));
  names.add("Sam");
  names.add("Devon");
  names.add("Sam");
  names.add("Sam");
  assertEquals(0, names.occurencesOf("NOT Here"));
  assertEquals(1, names.occurencesOf("Devon"));
  assertEquals(3, names.occurencesOf("Sam"));
}
```

If any assertion fails, either that assertion is stated incorrectly or a method such as isEmpty or add is incorrect. Or the constructor (new) may be wrong. In any case, you will know something is wrong with the concrete implementation or testing of the new type (a Java class).

One Class, Any Type Element

The Bag ADT from NIST does not specify what type of elements can be stored. The Bag interface uses type parameters top specify the type of elements that can be added, removed and gotten. Any class that implements the `Bag interface` above will be able to store the type of element with the type passed during the construction of the object. This is specified with the parameter `<E>` in add, occurencesOf, and remove. The compiler replaces the type specified in a construction wherever E is found. So, when a `Bag` is constructed to store `BankAccount` objects, you can add `BankAccount` objects. At compiletime, the parameter E in **public void** add(E element) is considered to be `BankAccount`.

```
Bag<BankAccount> accounts = new ArrayBag<BankAccount>();
accounts.add(new BankAccount("Chris", 100.00));
accounts.add(new BankAccount("Skyler", 2802.67));
accounts.add(new BankAccount("Sam", 31.13));
```

When a `Bag` is constructed to store `Point` objects, you can add `Point` objects. At compile time, the parameter `E` in **public void** add(E element) is considered to be `Point`.

```
Bag<Point> points = new ArrayBag<Point>();
points.add(new Point(2, 3));
points.add(new Point(0, 0));
points.add(new Point(-2, -1));
```

Since Java 5.0, you can add primitive values. This is possible because integer literals such as `100` are autoboxed as new `Integer(100)`. At compiletime, the parameter `E` in **public void** add(E element) is considered to be `Integer` for the Bag `ints`.

```
Bag<Integer> ints = new ArrayBag<Integer>();
ints.add(100);
ints.add(95);
ints.add(new Integer(95));
```

One of the advantages of using generics is that you cannot accidentally add the wrong type of element. Each of the following attempts to add the element that is not of type E for that instance results in a compiletime error.

```
ints.add(4.5);
ints.add("Not a double");
points.add("A String is not a Point");
accounts.add("A string is not a BankAccount");
```

A JUnit test describes the behavior of methods in a very real sense by sending messages and making assertions about the state of a `Bag`. The assertions shown describe operations of the Bag ADT in a concrete fashion.

Not Implementing the `Bag<E>` interface

This chapter does not have a class to implement the methods of the Bag<E> interface because it would have algorithms very similar to the non-generic `StringBag` from an earlier chapter. Additionally, the next chapter will implement a new interface (`OurList`) using a type parameter to give us a generic collection using an array instance variable.

Self-Check

15-1 Write a test method for `remove`.

15-2 Implement the `remove` method as if you were writing it inside the `ArrayBag` class when the class begins as

```java
// An ArrayBag object can store elements as a multi-set where elements can
// "equals" each other.  There is no specific order to the elements.
public class ArrayBag<E> implements Bag<E> {

  // --Instance variables
  private Object[] data;
  private int n;

  // Construct an empty bag that can store any type of element.
  public ArrayBag() {
    data = new Object[20];
    n = 0;
  }
```

```java
  // Return true if there are 0 elements in this bag.
   public boolean isEmpty() {
    return n == 0;
  }

  // Add element to this bag
  public void add(E element) {
    data[n] = element;
    n++;
  }

   // Add remove here
}
```

Answers to Self-Check Questions

15-1 One possible test method for remove:

```java
@Test
public void testRemove() {
  Bag<String> names = new ArrayBag<String>();
  names.add("Sam");
  names.add("Chris");
  names.add("Devon");
  names.add("Sandeep");
  assertFalse(names.remove("Not here"));

  assertTrue(names.remove("Sam"));
  assertEquals(0, names.occurencesOf("Sam"));
  // Attempt to remove after remove
  assertFalse(names.remove("Sam"));
  assertEquals(0, names.occurencesOf("Sam"));

  assertEquals(1, names.occurencesOf("Chris"));
  assertTrue(names.remove("Chris"));
  assertEquals(0, names.occurencesOf("Chris"));

  assertEquals(1, names.occurencesOf("Sandeep"));
  assertTrue(names.remove("Sandeep"));
  assertEquals(0, names.occurencesOf("SanDeep"));

  // Only 1 left
  assertEquals(1, names.occurencesOf("Devon"));
  assertTrue(names.remove("Devon"));
  assertEquals(0, names.occurencesOf("Devon"));
}
```

15-2 remove method in the ArrayBag class

```java
// Remove element if it is in this bag, otherwise return false.
public boolean remove(E element) {

  for (int i = 0; i < n; i++) {
    if (element.equals(data[i])) {
      // Replace with the element at the send
      data[i] = data[n - 1];
      // Reduce the size
      n--;
      return true; // Found--end the search
    }
  }
  // Did not find element "equals" anything in this bag.
  return false;
}
```