

# Data Abstraction in Java

H. Conrad Cunningham

15 June 2022

## Contents

<b>2</b>	<b>Data Abstraction in Java</b>	<b>2</b>
2.1	Chapter Introduction . . . . .	2
2.2	Java as an Object-Oriented Language . . . . .	2
2.3	Java Classes . . . . .	4
2.3.1	Class and Instance Methods . . . . .	4
2.3.2	Class and Instance Variables . . . . .	5
2.3.3	Public and Private Accessibility . . . . .	6
2.3.4	Primitive and Reference Variables . . . . .	6
2.4	Implementing ADTs as Java Classes . . . . .	7
2.4.1	Java Implementation of Bounded Stack . . . . .	12
2.5	Better Approach to Implementing ADTs in Java . . . . .	14
2.5.1	Java Class Implementation for Day . . . . .	16
2.6	What Next? . . . . .	22
2.7	Exercises . . . . .	22
2.8	Acknowledgments . . . . .	22
2.9	Concepts . . . . .	23
2.10	References . . . . .	23

Copyright (C) 1996-2022, H. Conrad Cunningham  
Professor of Computer and Information Science  
University of Mississippi  
214 Weir Hall  
P.O. Box 1848  
University, MS 38677  
(662) 915-7396 (dept. office)

**Browser Advisory:** The HTML version of this textbook requires a browser that supports the display of MathML. A good choice as of June 2022 is a recent version of Firefox from Mozilla.

## 2 Data Abstraction in Java

### 2.1 Chapter Introduction

This chapter is a Java-specific supplement to the chapter *Data Abstraction Concepts* [2]. However, most of the concepts apply to other object-oriented languages (e.g., Scala [9,10]). The discussion here is meant to be read in conjunction with the *Data Abstraction Concepts* chapter.

This chapter discusses use of Java classes to implement abstract data types (ADTs). It seeks to use good object-oriented programming practices, but it does not cover the principles and practices of object-oriented programming fully. For more information on object orientation, see my the chapters Object-Oriented Software Development [5] and Object-Based Paradigms [4] (a chapter in the ELIFP textbook [3, Ch. 3]).

Caveat: I wrote this chapter originally in the late 1990s. It should be updated to use Java generics (and possibly other newer Java features such as default methods in interfaces and lambda expressions). However, the basic principles used here are still relevant to contemporary Java programming.

TODO: Update chapter to use generics and possibly other newer Java features as needed.

### 2.2 Java as an Object-Oriented Language

TODO: Make sure this new section and the older sections that follow use consistent terminology and flow smoothly.

According to the concepts and terminology used in the Object-Based Paradigms [4] chapter of the ELIFP textbook [3, Ch. 3], Java is an object-oriented language. Its object model includes:

- objects
  - Java objects are instances of classes.
- classes
  - Java classes are declared with the keyword **class**.
- inheritance
  - A Java **class** inherits from (i.e., **extends**) one other Java class—the builtin, top-level class `Object` if no class is explicitly specified.
  - In addition to the single inheritance from a class, a Java class **implements** zero or more interfaces. A Java **interface** is similar to an abstract class; it cannot be instantiated to create objects.
- subtype polymorphism

Java's the system treats child classes as subtypes of the parent class. It dynamically binds the invocation of a method to the appropriate method implementation by searching up the inheritance hierarchy.

As we use Java in this chapter, a Java object exhibits all three “essential characteristics” of objects [4]:

a. state

The state of a Java object is the mapping of the object's attributes (i.e., instance variables) to their values.

b. operations

The operations of a Java object are the methods defined for the object. A method takes the object and zero or more other arguments and either returns information about the object's state or changes the its state.

c. identity

Each Java object (i.e., instance of a class) has an identity that is distinct from all other Java objects (e.g., the address of the object in memory).

As we use Java in this chapter, a Java object also exhibits both “important but non-essential characteristics” of objects [4]:

d. encapsulation

A Java class defines sets of instance variables and methods for instances of the class. If these instance variables or methods are declared as **private**, then they can only be accessed from within the class and, hence, are encapsulated within the class. If they are marked as **public**, then they can be accessed from outside the class. As we see, we normally explicitly declare all instance variables **private** and all operations **public**.

Java's encapsulation is thus at the class-level not the object level.

e. independent lifecycle

Java objects exist independently of the program units (e.g., methods or classes) that create them. An object can be instantiated in one program unit and passed to other program units where it is used. When the object is no longer accessible, its resources (e.g., memory) can be reclaimed by the garbage collector.

This chapter uses a cluster of related Java classes to implement modules. Java classes should thus be designed and implemented to satisfy the principle of information hiding. That is, they should not reveal details of their internal implementation (i.e., their secrets). They should only reveal aspects needed for the abstraction.

According to the concepts and terminology used in the Types [1] chapter of the ELIFP textbook [3, Ch. 5], Java exhibits

- static typing

The Java compiler determines the type of an expression based on the explicitly declared (or inferred) types of its component subexpressions.

- nominal typing

In addition to the builtin types, the `class` and `interface` define types.

If class `Child` `extends` `Parent`, then type `Child` is considered a subtype of type `Parent`. Similarly, if a class or interface `Child` `implements` `IParent`, then type `Child` is considered a subtype of type `IParent`.

- polymorphic variables

Java variables declared with some type `Base` can hold objects of type `Base` or of any of its subtypes.

- polymorphic operations

Java dynamically binds method invocations to method implementations along the subtype (inheritance) hierarchy.

Java also overloads function calls and some builtin operators.

Java classes (especially those implementing ADTs) should be designed and implemented so that they satisfy the Liskov Substitution Principle [1,7,11]. That is, it should be possible to substitute a subclass instance for a superclass instance in any circumstances.

In the following sections, we examine how to use Java to implement abstract data types. We will elaborate on some of the concepts mentioned in this section.

## 2.3 Java Classes

As a language construct, a Java `class`{.java} is similar to a user-defined `struct` type in C or user-defined record *type* in Pascal. A `class` is a template for constructing data items that have the same structure but differing values (states). We say that an item constructed by a class is a *class instance* (or an *object*).

Like the C structure type or Pascal record type, a Java class can consist of several components. In C and Pascal, all the components are data fields. However, in Java, functions and procedures may be included as components of a class. These procedures and functions are called *methods*.

### 2.3.1 Class and Instance Methods

A method declared in a class may be either a class method or instance method.

- A *class method* is associated with the class as a whole, not with any specific instance.
- An *instance method* is associated with an instance of the class.

We declare a method as a *class method* by giving the keyword `static` in the header of its definition. For example, a `main` method of a program is a class method of the class in which it is defined.

```
public static void main(String[] args)
{ // beginning code for the program
}
```

If we *do not include* the keyword `static` in the header of a method definition, the method is an *instance method*. For example, consider methods that implement the bounded stack's `push` and `top` operations as specified in the chapter Data Abstraction Concepts [2]:

```
public void pop()
{ // code for pop operation
}

public Object top()
{ // code for top operation
}
```

Note that `pop()` is a *procedure* (i.e., it has return type `void`) method and `top` is a *function* method.

Scala note: The Scala language [8,10] does not have static members of classes. However, the methods of a Scala singleton `object` have basically the same characteristics described above for “class methods”. Often the “class methods” appear in the companion `object` for a class. i.e., the `object` with the same name as the class.

### 2.3.2 Class and Instance Variables

In a similar fashion, the variables (data fields) declared in a class may be either class variables or instance variables.

- A *class variable* is associated with the class as a whole; there is only one copy of the variable for the entire class. As with methods, the keyword `static` is used to declare a class variable.
- An *instance variable* is associated with an instance of the class; each instance has its own instance of the variable. As with methods, the absence of the keyword `static` denotes an instance variable.

An instance method has direct access to the instance variables of the class instance (object) to which it is applied. The instance's variables are implicit arguments of the method calls. (If needed to distinguish among names, the builtin variable `this` can be used to refer to the instance to which the method is applied.) The instance methods also have access to the class variables (if any).

Class methods only have access to the class variables. The methods do not

have any implicit arguments. In fact, class methods can be called without any instances of the class being in existence.

Scala note: The Scala language does not have static members of classes. However, the variables of a Scala singleton `object` have basically the same characteristics described above for “class variables”. Often the “class variables” appear in the companion `object` for a class. i.e., the `object` with the same name as the class.

### 2.3.3 Public and Private Accessibility

The components of a class can be designated as `public` or `private`.

- The `public` components of the class are accessible from anywhere in the program (i.e., from any package).
- The `private` components are only accessible from inside the class.

As a general rule, the data fields of a class should be *private instance variables*, meaning that they are associated with a specific instance and are only accessible by the instance methods. This hides, or encapsulates, the data fields within the class instance.

Note: Actually, the instance methods of a Java class can access the instance variables of any instance of that class, not just the current instance.

In general, avoid public instance variables. They break the principle of information hiding, leading to potential entanglements among modules.

A public method of a class is a service provided by that instance to other parts of a program. The private methods of a class can be used in implementing the public methods.

Class methods and variables should be used sparingly. These are more or less the types of subprograms and global variables found in languages like C and Pascal. Their excessive use can greatly reduce the potential benefits that can be realized from object-oriented techniques.

Java note: There are two other types of accessibility, “friendly” and `protected`, but `public` and `private` are sufficient for our discussion of ADT implementations.

Scala note: Although similar in concept to that of Java, the accessibility features of Scala differ somewhat [8,10]. By default, all features are public in Scala, but accessibility can be restricted in a more fine-grained manner than in Java. The unmodified keyword `private` has the same meaning as in Java.

### 2.3.4 Primitive and Reference Variables

A Java variable is a strongly typed “container” in memory that is declared to hold either:

- a *value* of the associated primitive data type such as integers (`int`), floating point numbers (`double`), booleans (`boolean`), and single characters (`char`).
- a *reference* to (i.e., memory address of) an instance of the associated class (or other reference) type.

Java note: Although arrays are not class instances, array variables hold a reference to an instance of the array.

The class instances themselves are stored in the dynamically managed heap memory area. Java allocates memory from the heap to hold newly constructed instances of a class. Java’s garbage collector reclaims the memory for instances that are no longer needed by the program.

Note: Recent versions of Java can sometimes hide the differences between primitive values and references by automatically “boxing” primitive values as instances of the corresponding wrapper classes (e.g., `int` values as `Integer` instances). Scala goes further in that primitives and references are in the same type hierarchy. However, both languages run on the Java Virtual Machine, which makes a distinction between primitive values and references (i.e., pointers), so it is not possible to avoid the distinction entirely.

## 2.4 Implementing ADTs as Java Classes

If only one implementation of an ADT is needed, the following techniques can be used to implement an ADT using Java.

The implementation techniques discussed in this section implement the ADT in an imperative way. That is, instead of returning a new instance of the ADT with a modified state, a mutator operation usually modifies the state of the existing instance.

Caveat: The discussion of Java in this chapter does not use generic type parameters. For the `StackB` ADT (defined in the Data Abstraction Concepts), the type of the `Item` values stored in the stack can be a parameter of the `StackB` class.

TODO: Consider modifying this discussion to use an `Item` generic parameter.

1. **Use the Java `class` construct to represent the entire ADT.** If we want to allow access to the class from anywhere in the program, we will make the class `public`.

For the `StackB` ADT, we can use the following structure for the class:

```
public class StackB
{ // implementation of instance methods and data here
}
```

2. **Use an instance of the Java class to represent an instance of the ADT and, hence, variables of the class type to hold references to instances.**

For example, to declare a variable that can hold a reference to a `StackB` instance, we can use the following declaration:

```
StackB stk;
```

3. **As each component of the class is defined, ensure that the semantics of the ADT operations are implemented appropriately.** That is, make sure:

- an appropriate implementation (representation) invariant is defined to capture what it means for the internal state of an instance to be valid,
- the interface and implementation invariants are established (i.e., made true) by the constructors and preserved (i.e., kept true) by the mutator and accessor methods,
- each method's postcondition is established by the method in any circumstance when it is called with the precondition true.

The class and its methods should be documented with the invariants, preconditions, and postconditions.

4. **Represent the ADT's constructors by Java constructor methods. In most circumstances, also include a parameterless default constructor.**

A Java constructor is a method with the same name as the class. It does not have a return type specified. Upon creation of an instance of the class, the constructor initializes the instance's state so that the class invariants are established.

A constructor is normally invoked by the Java operator `new`. The operator `new` allocates memory on the heap for the instance, calls the constructor to initialize the new instance, and then returns a reference to the new instance.

For example, we can represent the ADT operation `create` by the constructor method `StackB`.

```
public class StackB
{   public StackB(int size)
    {   // initialization code
    }

    // rest of StackB methods and data ...
}
```

A user of the `StackB` class can then declare a variable and initialize it to hold a reference to a new stack with a capacity of 100 items as follows:

```
StackB stk = new StackB(100);
```



The expression `new StackB(100)` allocates a `StackB` instance in the heap storage and calls the constructor above to initialize the data fields encapsulated within the instance.

5. **Represent the ADT operations by instance methods of the class.** Thus the state of the ADT instance, which is given explicitly in the ADT signatures, becomes an *implicit argument* of all method calls. Mutators also have the state as an implicit return.

We can apply a method to a class instance by using the selector (i.e., “dot”) notation. This notation is similar to the notation for accessing `record` components in Pascal.

For example, in the case of the `StackB` ADT we can represent the operations as instance methods of class `StackB`. The explicit `StackB` parameters and return values of the operations thus become implicit.

Suppose we want to push an item `x` onto the `stk` created above. We can do that with the following code:

```
if (!stk.full())
    stk.push(x);
```

We can then examine the top item and remove it:

```
if (!stk.empty())
{   it = stk.top();
    stk.pop();
}
```

6. **Make the constructors, mutators, accessors, and destructors public methods of the class.** That is, precede the method’s definition by the keyword `public`.
7. **Represent the ADT mutator operations by Java procedure (i.e., `void`) methods, except those mutator operations that explicitly require new instances to be generated (e.g., a copy or clone operation).**

For example, the `pop` method of `StackB` would have the following structure:

```
public void pop()
{   // code to implement operation
}
```

A mutator method modifies the encapsulated state of the class instance (which is the implicit argument of the method). In any circumstance in which its precondition and the class invariants hold on entry, the method must establish its postcondition and reestablish the invariants upon exit. (The invariant might not hold in the middle of the method’s execution.)

Comment: Implementing mutator operations as procedure calls that modify the stored state is really an optimization. All mutators can be implemented in the applicative style, returning a modified copy of the instance. This implementation might, however, be inefficient in use of processor time and memory.

8. **For certain mutator operations (e.g., copy or clone), implement the corresponding Java methods to return new instances of the class rather than to modify the current instance (i.e., their implicit arguments).**

Any mutator method must, of course, establish its postcondition and reestablish the invariants for the current instance. In addition, these applicative mutators must also establish the invariants for the new instance returned.

9. **Represent the ADT accessor operations by Java function methods of < the proper return type.**

For example, the `empty` method of `StackB` would have the following structure:

```
public boolean empty()
{ // code to implement operation
}
```

An accessor method accesses the encapsulated state of the class instance (which is the implicit argument) and computes a value to be returned. In any circumstance in which its precondition and the class invariants hold on entry, the method must establish its postcondition and reestablish the invariants upon exit. (The invariant might not hold in the middle of the method's execution.)

10. **If necessary for deallocation of internal resources, represent the ADT destructor methods by explicit Java procedures; in most cases, however, just allow the automatic garbage collection to reclaim instances that are no longer being used.**

For example, in the `StackB` class, we might include an explicit destroy operation that releases the storage resources and disables further use of the instance.

```
public void destroy()
{ // code to free resources
}
```

Java note: The Java framework allows a `finalize()` method to be included in each class. This method is called implicitly whenever the garbage collector detects that the instance is no longer in use. However, since it is difficult to predict when (if ever) this method will be executed, it is safer

to include explicit destructors when resources are in short supply and must be explicitly managed.

11. Use **private** data fields of the Java class to represent the encapsulated state of the instance needed for a particular implementation. By making the data fields **private** they are still available to the instance's methods, but are not visible outside the class.

For example, the `StackB` class might have the following data fields:

```
public class StackB
{ // public operations of class instance

    // encapsulated data fields of class instance
    private int topItem; // Pointer to next index for insertion
    private int capacity; // Maximum number of items in stack
    private Object[] stk; // the stack
}
```

12. Do not use **public** data fields in the class. These violate the principle of information hiding. Instead introduce appropriate accessor and mutator methods to allow manipulation of the hidden state.
13. Include, as appropriate, **private** methods to aid in implementation.

Functionality common to several methods can be placed in separate functions and procedures as needed. However, since these are **private**, they can only be accessed from within the class and thus can be changed without affecting the public interface of the class.

14. Add any other methods needed to make the ADT fit into the Java environment.

For example, it is frequently useful to add public `toString` and `clone` methods. The `toString` method returns a Java `String` reflecting the "value" of the instance in a format suitable for printing. The `clone` method creates a new instance that has the same value as the current instance.

15. In general, avoid use of class (i.e., **static**) variables. Since a class variable is shared among all instances of the class, it may be difficult to preserve the invariants for individual instances as the value of the class variable changes.

However, it is a good programming practice to use class *constants* where appropriate. These are data fields declared with both the **static** and **final** modifiers. Their values may be initialized but cannot be changed thereafter.

These constants may be declared `private` if usage is to be restricted to the class or `public` if the users of the class also need access.

By convention, the names of constants are normally written with all uppercase letters. For example, the following defines a symbolic name for the integer code used for Sunday as a day of the week in the `Day` class defined later.

```
public static final int SUNDAY = 1;
```

Caveat: When this set of notes was originally written, Java did not yet have generics. So the examples below handle the type parameters of the ADT in other ways. A Java generic provides a class facility that can be parameterized with types (like the C++ `template` or Ada `generic` mechanisms).

For example, in the implementation below we represent the set `Item` of the `StackB{.java}` ADT by the class `Object`. As we will see when we discuss inheritance, the `Object` type will allow us to store an instance of any class on the `StackB`. With this definition, any data of a reference type can appear in the stack, but values of the primitive types cannot. A better implementation would have `Item` as type parameter of the class.

The next section gives a Java implementation of the `StackB` ADT. A similar constructive definition and two implementations of a `Queue` ADT are available in a separate document.

TODO: Beter integrate `SoftwareInterfaces` into this document collection.

#### 2.4.1 Java Implementation of Bounded Stack

TODO: Reconstruct (or find) the following source code and ensure that it executes on the current Java platform. Insert link.

In this section, we give an implementation of the `StackB` ADT that uses an array of objects and an integer “pointer” to represent the stack. (This implementation does not use Java generic classes.)

This implementation is not robust; each operation assumes that its precondition holds. A more robust implementation might check whether the precondition holds and throw an exception if it does not.

Remember that the invariants are implicitly pre- and postconditions of all mutator and accessor methods, postconditions of the constructor, and preconditions of the destructor.

```
// A Bounded Stack ADT
public class StackB
{ // Interface Invariant: Once created and until destroyed, this
  //   stack instance has a valid and consistent internal state

  public StackB(int size)
```

```

// Pre:  size >= 0
// Post:  initialized new instance with capacity size && empty()
{   stk = new Object[size];
    capacity = size;
topItem = 0;
}

public void push(Object item)
// Pre:  NOT full()
// Post:  item added as the new top of this instance's stack
{   stk[topItem] = item;
    topItem++;
}

public void pop()
// Pre:  NOT empty()
// Post:  item at top of stack removed from this instance
{   topItem--;
    stk[topItem] = null;
}

public Object top()
// Pre:  NOT empty()
// Post:  return item at top of this instance's stack
{   return stk[topItem-1];
}

public boolean empty()
// Pre:  true
// Post:  return true iff this instance's stack has no elements
{   return (topItem <= 0);
}

public boolean full()
// Pre:  true
// Post:  return true iff this instance's stack is at full capacity
{   return (topItem >= capacity);
}

public void destroy()
// Pre:  true
// Post:  internal resources released;  stack effectively deleted
{   stk = null;
capacity = 0;
    topItem = 0;
}

```

```

// Implementation Invariant for informal model:
//    0 <= topItem <= capacity &&
//    stack is in array section stk[0..topItem-1]
//    with the top at stk[topItem-1], etc.

// Implementation Invariant for more formal model representing stack
// as tuple (integer max, sequence stkseq)
//    m == capacity && 0 <= topItem <= capacity &&
//    stackInArray(stk, topItem, stkseq)
//    where stackInArray(arr, t, ss) = if t == 0 then ss == []
//                                     else arr[t-1] == head(ss)
//                                     && stackInArray(arr, t-1, tail(ss))

private int topItem; // Pointer to next index for insertion
private int capacity; // Maximum number of items in stack
private Object[] stk; // the stack
}

```

## 2.5 Better Approach to Implementing ADTs in Java

TODO: Reconstruct (or find) the following source code and ensure that it executes on the current Java platform. Inset link.

If several different implementations of an ADT are needed, then the Java specification of an ADT's interface should be separated from the class implementation. The interface specification can be reused among several classes and various implementations of the interface can be used interchangeably.

This can be done as follows.

1. Define a Java **interface** that specifies the type signatures for the ADT's mutator and accessor (and, if needed, destructor) operations. These method signatures should have the same characteristics as described above in the discussion of class-based specification.
2. Specify and document the **interface** by the interface invariants, preconditions, and postconditions that must be supported by any implementation of the ADT. There are no implementation invariants for an interface, but individual classes that implement the **interface** will have them.

For example, a bounded stack **interface** might be specified as follows:

```

public interface StackADT
{ // Interface Invariant: Once created and until destroyed, this
  //   stack instance has a valid and consistent internal state

  public void push(Object item);

```

```

// Pre:  NOT full()
// Post: item added as the new top of this instance's stack

...

public Object top();
// Pre:  NOT empty()
// Post: return item at top of this instance's stack

...
}

```

3. Provide one or more concrete classes that implement the interface.

For example, an array-based StackADT could be implemented similarly to the StackB definition given in the previous section.

```

public class StackInArray implements StackADT
{ // Interface Invariant: Once created and until destroyed, this
  //   stack instance has a valid and consistent internal state

public StackInArray(int size)
// Pre:  size >= 0
// Post: initialized new instance with capacity size and empty()
{  stk = new Object[size];
  capacity = size;
   topItem = 0;
}

public void push(Object item)
// Pre:  NOT full()
// Post: item added as the new top of this instance's stack
{  stk[topItem] = item;
   topItem++;
}

...

public Object top()
// Pre:  NOT empty()
// Post: return item at top of this instance's stack
{  return stk[topItem-1];
}

...
}

```

```

// Implementation Invariant for informal model:
//    0 <= topItem <= capacity &&
//    stack is in array section stk[0..topItem-1]
//    with the top at stk[topItem-1], etc.

// Implementation Invariant for more formal model representing stack
// as tuple (integer max, sequence stkseq)
//    m == capacity && 0 <= topItem <= capacity &&
//    stackInArray(stk, topItem, stkseq)
//    where stackInArray(arr, t, ss) =
//        if t == 0 then ss == []
//        else arr[t-1] == head(ss)
//        && stackInArray(arr, t-1, tail(ss))

private int topItem; // Pointer to next index for insertion
private int capacity; // Maximum number of items in stack
private Object[] stk; // the stack
}

```

4. Declare variables of the ADT's **interface** type to hold instances of any concrete class that implements the **interface**. Any of the operations defined in the **interface** can be applied to the instance to which this variable refers.

For example, a variable of type `StackADT` can hold instances of any concrete class that implements the **interface** `StackADT`.

```

StackADT theStack = new StackInArray(100);
theStack.push("Hello World");

```

For an ADT specification and implementations that follow this approach, see the description of the Ranked Sequence ADT case study given in a separate document. In addition to Java interfaces, the Ranked Sequence case study uses other Java features such as exceptions, enumerations, packages, and Javadoc annotations.

TODO: Integrate SoftwareInterfaces into this document collection.

### 2.5.1 Java Class Implementation for Day

TODO: Reconstruct (or find) the following source code and ensure that it executes on the current Java platform. Inset link.

The following implementation of the Day ADT is adapted from the like-named class in Horstmann and Cornell's book *Core Java* [6].

This implementation represents the calendar as three integers. It converts the dates to and from Julian dates to do some of the operations.



```

// This class implementation is adapted from the Day class in
// Horstmann and Cornell, Core Java 1.2: Volume I - Fundamentals
// (Fourth Edition), Prentice Hall, 1999.

import java.util.*;
import java.io.*;

public class Day
{
    // Interface Invariant: Once created and until destroyed, this
    // instance contains a valid date. getdate() != 0 &&&
    // 1 <= getMonth() <= 12 &&& 1 <= getDay() <= #days in getMonth().
    // Also calendar date getMonth()/getDay()/getYear() does not
    // fall in the gap formed by the change to the modern
    // (Gregorian) calendar.

    // Constants for days of the week

    public static final int SUNDAY    = 1;
    public static final int MONDAY    = 2;
    public static final int TUESDAY   = 3;
    public static final int WEDNESDAY = 4;
    public static final int THURSDAY  = 5;
    public static final int FRIDAY    = 6;
    public static final int SATURDAY  = 7;

    // Constructors

    public Day()
    // Pre: true
    // Post: the new instance's day, month, and year set to today's
    // date (i.e., the date of creation of the instance)
    //
    // Implementation uses GregorianCalendar class from the Java API
    // to get today's date.
    //
    { GregorianCalendar todaysDate = new GregorianCalendar();
      year  = todaysDate.get(Calendar.YEAR);
      month = todaysDate.get(Calendar.MONTH) + 1;
      day   = todaysDate.get(Calendar.DAY_OF_MONTH);
    }

    public Day(int y, int m, int d)
        throws IllegalArgumentException
    // Pre: y != 0 &&& 1 <= m <= 12 &&& 1 <= d <= #days in month m
    // (y,m,d) does not fall in the gap formed by the

```

```

//          change to the modern (Gregorian) calendar.
// Post:  the new instance's day, month, and year set to y, m,
//        and d, respectively
// Exception:  IllegalArgumentException if y m d not a valid date
{   year = y;
    month = m;
    day = d;
    if (!isValid())
        throw new IllegalArgumentException();
}

// Mutators

public void setDay(int y, int m, int d)
    throws IllegalArgumentException
// Pre:   y != 0 &&& 1 <= m <= 12 &&& 1 <= d <= #days in month m
//        (y,m,d) does not fall in the gap formed by the
//        change to the modern (Gregorian) calendar.
// Post:  this instance's day, month, and year set to y, m,
//        and d, respectively
// Exception:  IllegalArgumentException if y m d not a valid date
{   year = y;
    month = m;
    day = d;
    if (!isValid())
        throw new IllegalArgumentException();
}

public void advance(int n)
// Pre:   true
// Post:  this instance's date moved n days later. (Negative n
//        moves to an earlier date.)
{   fromJulian(toJulian() + n);
}

// Accessors

public int getDay()
// Pre:   true
// Post:  returns the day from this instance, where
//        1 <= getDay() <= #days in this instance's month
{   return day;
}

public int getMonth()
// Pre:   true

```

```

// Post: returns the month from this instance's date, where
//       1 <= getMonth() <= 12
{   return month;
}

public int getYear()
// Pre: true
// Post: returns the year from this instance's date, where
//       getYear() != 0
{   return year;
}

public int getWeekday()
// Pre: true
// Post: returns the day of the week upon which this instance
//       falls, where 1 <= getWeekday() <= 7;
//       1 == Sunday, 2 == Monday, ..., 7 == Saturday
{   // calculate day of week
    return (toJulian() + 1) % 7 + 1;
}

public boolean equals(Day dd)
// Pre: dd is a valid instance of Day
// Post: returns true if and only if this instance and instance
//       dd denote the same calendar date
{   return (year == dd.getYear() && month == dd.getMonth()
        && day == dd.getDay());
}

public int daysBetween(Day dd)
// Pre: dd is a valid instance of Day
// Post: returns the number of calendar days from the dd
//       instance's date to this instance's date, where
//       equals(dd.advance(n)) would hold
{   // implementation code
    return toJulian() - dd.toJulian();
}

public String toString()
// Pre: true
// Post: returns this instance's date expressed in the format
//       "Day[year,month,day]"
{
    return "Day[" + year + "," + month + "," + day + "];"
}

```

```

// Destructors -- None needed

// Private Methods

private boolean isValid()
// Pre: true
// Post: returns true iff this is a valid date
{
    Day t = new Day();
    t.fromJulian(this.toJulian());
    return t.day == day && t.month == month
        && t.year == year;
}

private int toJulian()
// Pre: true
// Post: returns Julian day number that begins at noon of this day
//
// A positive year signifies A.D., negative year B.C.
// Remember that the year after 1 B.C. was 1 A.D. (i.e., no year 0).
//
// A convenient reference point is that May 23, 1968, at noon
// is Julian day 2440000.
//
// Julian day 0 is a Monday.
//
// This algorithm is from Press et al., Numerical Recipes
// in C, 2nd ed., Cambridge University Press 1992.
//
{
    int jy = year;
    if (year < 0)
        jy++;
    int jm = month;
    if (month > 2)
        jm++;
    else
    {
        jy--;
        jm += 13;
    }
}
int jul = (int) (java.lang.Math.floor(365.25 * jy)
    + java.lang.Math.floor(30.6001*jm) + day + 1720995.0);

int IGREG = 15 + 31*(10+12*1582);
// Gregorian Calendar adopted Oct. 15, 1582

if (day + 31 * (month + 12 * year) >= IGREG)
    // change over to Gregorian calendar

```

```

    {   int ja = (int)(0.01 * jy);
        jul += 2 - ja + (int)(0.25 * ja);
    }
    return jul;
}

private void fromJulian(int j)
// Pre: true
// Post: this calendar Day is set to Julian date j
//
// This algorithm is from Press et al., Numerical Recipes
// in C, 2nd ed., Cambridge University Press 1992
//
{   int ja = j;

    int JGREG = 2299161;
        /* the Julian date of the adoption of the Gregorian
           calendar
           */

    if (j >= JGREG)
        /* correct for crossover to Gregorian Calendar */
    {   int jalpha = (int)((float)(j - 1867216) - 0.25)
            / 36524.25;
        ja += 1 + jalpha - (int)(0.25 * jalpha);
    }
    int jb = ja + 1524;
    int jc = (int)(6680.0 + ((float)(jb-2439870) - 122.1)
            /365.25);
    int jd = (int)(365 * jc + (0.25 * jc));
    int je = (int)((jb - jd)/30.6001);
    day = jb - jd - (int)(30.6001 * je);
    month = je - 1;
    if (month > 12)
        month -= 12;
    year = jc - 4715;
    if (month > 2)
        --year;
    if (year <= 0)
        --year;
}

// Implementation Invariants:
//   year != 0 && 1 <= month <= 12 && 1 <= day <= #days in month
//   (year,month,day) not in gap formed by the change to the
//   modern (Gregorian) calendar

```

```
        private int year;  
        private int month;  
        private int day;  
    }
```

## 2.6 What Next?

TODO

## 2.7 Exercises

TODO

## 2.8 Acknowledgments

This chapter was originally part of my Data Abstraction Concepts notes [2]. See the Acknowledgments section of that chapter for more information on its development. For the Lua-based offering of CSci 658 in Fall 2013. I separated most of the Java-specific material into this chapter to make the content of the Data Abstraction chapter more language independent.

In this chapter, I use a Java programming style influenced by Horstmann and Cornell’s book *Core Java* [6]. I adapted the `Java Day` design and implementation from the like-named class in Chapter 4 of Horstmann and Cornell’s of that book [6].

In Summer 2017, I modified this chapter slightly to link it into the revised (Pandoc Markdown) version of the Notes on Data Abstraction chapter and then reformatted it to use Pandoc Markdown in Spring 2018.

I retired from the full-time faculty in May 2019. As one of my post-retirement projects, I am continuing work on possible textbooks based on the course materials I had developed during my three decades as a faculty member. In January 2022, I began refining the existing content, integrating separately developed materials together, reformatting the documents, constructing a unified bibliography (e.g., using `citeproc`), and improving my build workflow and use of Pandoc.

In 2022, I also added the section “Java as an Object-Oriented Language” to better tie this chapter to the concepts and terminology used in the ELIFP textbook [3], especially chapters 3 and 5.

I maintain this chapter as text in Pandoc’s dialect of Markdown using embedded LaTeX markup for the mathematical formulas and then translate the document to HTML, PDF, and other forms as needed.

## 2.9 Concepts

TODO

### 2.10 References

- [1] H. Conrad Cunningham. 2019. *Type system concepts*. University of Mississippi, Department of Computer and Information Science, University, Mississippi, USA. Retrieved from <https://john.cs.olemiss.edu/~hcc/docs/TypeConcepts/TypeSystemConcepts.html>
- [2] H. Conrad Cunningham. 2022. *Data abstraction concepts*. University of Mississippi, Department of Computer and Information Science, University, Mississippi, USA. Retrieved from <https://john.cs.olemiss.edu/~hcc/docs/DataAbstraction/DA01/DataAbstractionConcepts.html>
- [3] H. Conrad Cunningham. 2022. *Exploring programming languages with interpreters and functional programming (ELIFP)*. University of Mississippi, Department of Computer and Information Science, University, Mississippi, USA. Retrieved from <https://john.cs.olemiss.edu/~hcc/docs/ELIFP/ELIFP.pdf>
- [4] H. Conrad Cunningham. 2022. Object-based paradigms. In *Exploring programming languages with interpreters and functional programming (ELIFP)*. University of Mississippi, Department of Computer and Information Science, University, Mississippi, USA. Retrieved from [https://john.cs.olemiss.edu/~hcc/docs/ELIFP/Ch03/03\\_Object\\_Paradigms.html](https://john.cs.olemiss.edu/~hcc/docs/ELIFP/Ch03/03_Object_Paradigms.html)
- [5] H. Conrad Cunningham. 2022. *Object-oriented software development*. University of Mississippi, Department of Computer and Information Science, University, Mississippi, USA. Retrieved from <https://john.cs.olemiss.edu/~hcc/docs/OOSoftDev/OOSoftDev.html>
- [6] Cay S. Horstmann and Gary Cornell. 1999. *Core Java 1.2: Volume I—Fundamentals*. Prentice Hall, Englewood Cliffs, New Jersey, USA.
- [7] Barbara Liskov. 1987. Keynote address—Data abstraction and hierarchy. In *Proceedings on object-oriented programming systems, languages, and applications (OOPSLA '87): addendum*, ACM, Orlando, Florida, USA, 17–34.
- [8] Martin Odersky, Lex Spoon, and Bill Venners. 2008. *Programming in Scala* (First ed.). Artima, Inc., Walnut Creek, California, USA.
- [9] Martin Odersky, Lex Spoon, and Bill Venners. 2021. *Programming in Scala* (Fifth ed.). Artima, Inc., Walnut Creek, California, USA.
- [10] Scala Language Organization. 2022. The Scala programming language. Retrieved from <https://www.scala-lang.org/>
- [11] Wikipedia: The Free Encyclopedia. 2022. Liskov substitution principle. Retrieved from [https://en.wikipedia.org/wiki/Liskov\\_substitution\\_principle](https://en.wikipedia.org/wiki/Liskov_substitution_principle)