

Problem Solving through Programming with Greenfoot

Consider structuring it somewhat like the Kolling book using examples. However, I would give complete information on theory when something is introduced and I would also use very different examples. I want the examples to be better motivated and problem solving exercises. The goal is to have one example each week. So I walk in each week with a new example problem and we spend the week trying to solve it. In the book the chapters would also present parallel problems that could use similar methods and perhaps give quick solutions to those.

The problem with this approach is getting them up to speed in the beginning. Greenfoot might help with this. I need to look into the API some. There are 14-15 weeks in a class. First week is introduction material including a problem that gets solved without code. Every week after that is a new example and the needed elements are introduced. We have to start simple. I think that every so often I could throw in a game type example to keep things light.

Topics and examples

1. Introduction to programming and problem solving
 1. Motivate programming and why they should learn it.
 2. Have them solve something in Greenfoot without coding.
 3. Instantiate objects and call methods on them.
 4. Arguments/parameteres and return types from the GUI.
 5. Playing bejeweled with different levels of helper methods.
2. Basics of coding
 1. Learn sequential execution
 2. Variable declarations
 3. Edit a little code to change the behavior of something in Greenfoot
 4. Methods with arguments and return types in code. (Primitives and reference)
 5. Syntax and rules. Expressions, types, statements.
3. Conditionals
 1. Boolean logic
 2. Make things turn when they run into walls.
 3. Put in keyboard input. (This means looking at the API.)
 4. Mandelbrot set
4. Loops and Greenfoot Lists
 1. Cover for loop for repetition of a task.
 2. Greenfoot lists and using loops to go through them.
 3. For-each loop.
 4. Real meaning of inheritance. Cover super and instanceof operator.
5. Arrays and Lists
 1. Processing data
 2. Using arrays to store direction for movement in Greenfoot.
 3. Explain inheritance and polymorphism with lists.
6. File Processing
 1. Exceptions (required for I/O)
7. Recursion
 1. Wombat making optimal path to eat leaves.
 2. Optimization problems

- 3. Graphs
- 8. Maps
 - 1. Correlating data from multiple files.
- 9. GUIs
- 10. Graphics

Think of problems and sample programs:

I need to think of ways to get some business problems in here. These can be complex algorithms though. Optimization problems are really significant in business settings.

Greenfoot Examples

- 1. Games
 - 1. Tag
 - 2. Shooters
 - 3. Puzzle
 - 4. Arcade
- 2. Mazes
- 3. Simulations
 - 1. Ants
 - 2. Termites
 - 3. Swarms
 - 4. Physics
 - 5. Populations
- 4. Turtles
 - 1. Basic
 - 2. L-Systems
- 5. Game Playing AI and Competition
 - 1. Snipes
 - 2. Mine sweeper
- 6. Fractals
- 7. Business/Optimization
 - 1. TSP
 - 2. Max flow
 - 3. Minimum spanning tree
- 8. Scheduling
 - 1. Disney trip
 - 2. Classes
- 9. Treatment protocols

Standard Java

- 1. Text File processing/alteration
 - 1. Finding a pattern in a text file. Counting how many time it occurs or replacing it with something else.
- 2. Basic GUIs
 - 1. Tables for displaying data.
- 3. Business/Optimization

Preface

There are two significant goals of this text: to teach basic programming skills and to improve students' problem solving capabilities. This book is intended to provide a fairly gentle introduction to the topic of programming using examples that will relate to and interest to most college students. The book focuses on the Java programming language and begins in the Greenfoot environment, then moves on to the Eclipse environment later on. The teaching of programming is in many regards secondary to the broader goal of improving problem solving skills and serves as a vehicle to accomplish that goal.

This book is intended to be used in classes for introductory computer science students who don't already have a background in programming or for courses that teach programming to students who don't intend to major in computer science or the natural sciences. Using Greenfoot for the early chapters affords an easy and approachable way to understand the main concepts of object-oriented programming and allows us to build interesting, interactive programs with graphical output right from the beginning. The fact that both Greenfoot and Eclipse use standard Java means that this book does not have to cover two different languages with different syntax or semantics. In addition, Greenfoot has a wonderfully small and simple API which allows students to learn how to read and make use of a real API early on without throwing them into the vast expanse of the full Java API.

Early on, examples are selected that fit well with Greenfoot's visual orientation and have students writing games, doing simulations, or solving simple programs in the 2-D graphical environment. From there the examples grow into things that students can see as relevant to their lives or their course work outside of programming. The primary objective of this book is not so much to teach programming, but to instead to teach problem solving. Programming just happens to be a superb instrument for students to learn and hone these skills, and the computing environment helps them not only to formalize their thought processes, but break things down into their most fundamental pieces to ensure that all the problems they have to solve are small ones.

Late in the book students are moved away from Greenfoot as we begin to use Eclipse, a professional scale development environment. The first steps in Eclipse are taken writing applets, and students are introduced to the full Java API, having seen only parts earlier when certain standard classes were used in Greenfoot. After applets, full applications are brought into the picture.

Many of the topics in this text stand on their own and instructors can choose to skip certain sections to fit their teaching objectives for the class as a whole.

Chapter 1 - Introduction to programming and problem solving

1.1. Programming and Problem Solving

Welcome to the world of programming! The fact that you are reading this likely implies that you are taking an introductory programming class. At some point, you've probably heard someone (maybe yourself!) ask you the question, "Why I should bother to learn how to program?" We want to address that, but before we do we have to take a step back and answer some other questions first.

Let's begin with the question, "What is a program?" A program is a step-by-step set of instructions that tells a computer how to do something. A more general term is "algorithm" which is a step-by-step description of how to accomplish some task, on a computer or otherwise. This is actually a concept that everyone has been familiarized with, even if you have never programmed a computer. There are many different contexts in which you interact with step-by-step instructions for completing tasks, one of the most common being recipes. A recipe is simply a set of instructions that a person needs to follow in order to complete the preparation for a certain dish. Recipes aren't the only example either. Any set of instructions given to people for them to complete something, whether it be a worksheet for an elementary school student, a Lego building guide, or a tax form from the IRS, is basically an algorithm intended for a human.

One of the first differences you have to understand between writing a set of instructions for a person and one for the computer, is that the computer, while very powerful and fast, is rather stupid. It has little capacity to interpret things as something other than what is presented directly. As a result of this, the instructions given to them have to be very precise. Computers do exactly what you tell them to do, nothing more, nothing less. When you tell them to do something you also have to break it down into steps that they understand. The only steps the computer really understands are steps that are very small and exact. This has a significant impact on how we write programs, for example we don't use normal English to do it. It also makes programming a very valuable skill to learn even if programming computers is not part of your life plan. Understanding how work, will help in any field that uses them, and let's face it, they aren't going away soon.

Everything inside of a computer is actually stored as part of a vast set of numbers and the computer does nothing more than basic arithmetic operations on those numbers. The instructions that computers follow are numeric values. What numbers correspond to what instructions depends on the details of the computer you are using. Early programmers had to write everything in this numeric format that is commonly called *machine language* because it is the only language the computer really understands.

Writing programs in machine language is difficult, time consuming and rather unproductive. For these reasons, it wasn't long before people came up with other languages to write their programs in. To make these languages work they first had to write programs that would translate from their new languages to the machine language. These programs are called compilers, and allow us as humans to write programs in languages that are vastly more readable and that are easier for us to think in than trying to do everything at the level of the machine. They also have the advantage that while different machines might have different machine languages, a single program from a higher level language can run across many machines as long as each one has a compiler that produces the appropriate type of machine language.

There are literally thousands of different programming languages that have been developed over the decades. Many of them are still in use today. Some languages are good for certain tasks. Some people seem naturally drawn to one language over another. In this book we will be using the *Java programming language* originally created by Sun Microsystems. Java is a high level programming

language that is used in many different areas. It is freely available and has been implemented on many different systems. You can download Java for your computer by going to <http://java.sun.com>. To write Java programs all you need is a general text editor, like Notepad on a Windows system.

But just because you can do it that way doesn't make it optimal or even appealing. This book starts off using a programming environment called Greenfoot (<http://www.greenfoot.org>) that was developed for educational purposes. Greenfoot will help you to understand some of the aspects of the Java language and also provide you with a fairly simple way to make interesting programs with relatively little effort. Later, when we have developed our understanding of the main concepts of programming in general and Java specifically, we will spend some time in the Eclipse (<http://www.eclipse.org>) programming environment. Eclipse is a fully featured professional software development environment. Such environments have a great number of different tools built into them that can be very helpful to professional software developers and that can make programming a lot easier. Of course, you have to reach a certain level of understanding before they are really of significant assistance.

Now that we have addressed what a program is, we can return to our original question of why you should bother to learn to program. There are two very different, yet equally valid answers to this. We will start with the one that is most obvious to people. You should learn how to program because it is a valuable skill. If you are considering majoring in computer science, programming is going to be at the heart of pretty much everything you do. It is how you communicate with the computer. While computer science is about a lot more than programming, the ability to write programs and understand how they work is an essential undercurrent for the rest of the field.

Even if you have no interest in majoring in computer science or ever working in computer science, learning how to program is still a valuable skill because computers are so omnipresent. Nearly every activity you do in your daily life involves a computer, whether you realize it or not. When you think about computers in your life, you might only think of your desktop or laptop. However, your cell phone, iPod, and every other electronic gadget you might own these days is really a digital computer. There are computers in nearly every car you see on the road that control various aspects of the operation of the engine as well as the internal electronics.

So far we have just listed what you might call “client-side computers,” computer devices that people own and use, the ones you might interact with directly. Computers also control a lot of things in the background. They are responsible for street lights. They route packages that are sent to you and they oversee inventory control at your local stores. They even govern the the distribution of electricity and water. If all the computers in the world were to suddenly cease to function, us humans would find we have a very hard time getting anything done.

All of those computer systems run on programs. You have to have programs telling the computers what they should be doing at every instant they are in operation. Now, this might sound like a motivation for you to switch your major to computer science. In some ways it is, but the reality is that you need people with domain specific knowledge of different fields as well who can help determine how to best put computers to use in those areas. You see, even if you don't plan on programming for a living, you will be interacting with programs for a living. If you advance far enough in whatever career you pick, at some point you will be responsible for making decisions that relate to the acquisition and use of technology. The use of computers is only increasing over time, and even if you don't want to be the person writing the programs, you will at some point be working with those people to help them create software that improves how you do your job. The more knowledge you have about what those people are doing, the better you will be able to communicate with them and the more valuable you will be at your job.

There is also a less obvious reason to learn how to program that applies equally well to anyone.

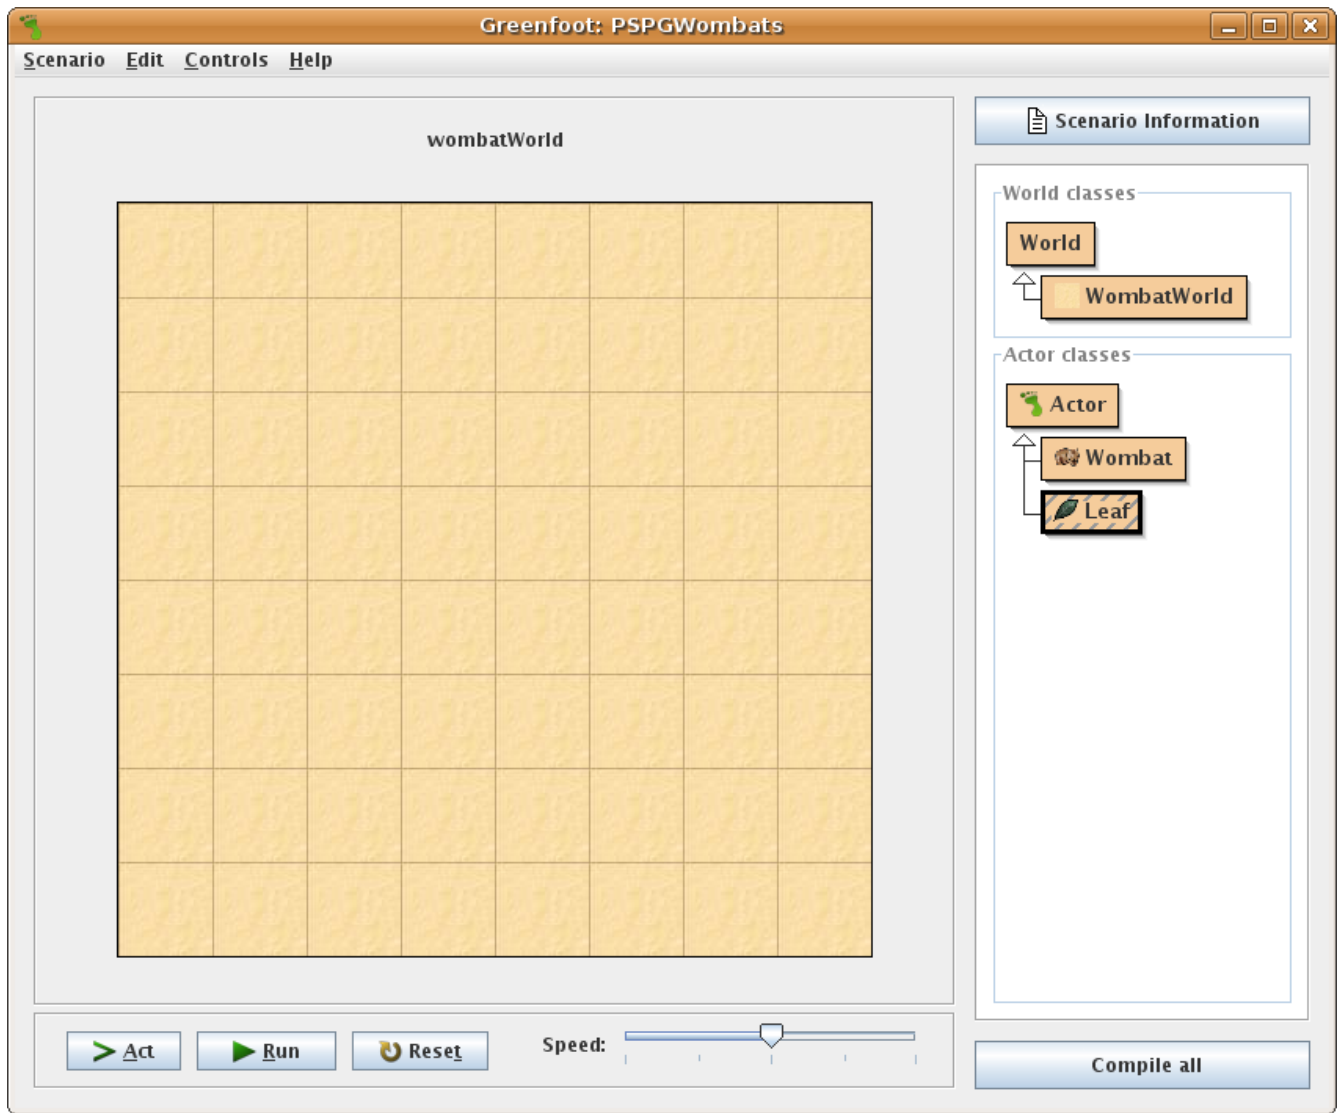
Remember what a program is? A set of instructions given to a computer to tell the computer how to complete some task. Most programs solve a type of problem for us. The first step in any programming task is figuring out *how* to solve the problem yourself in sufficient detail that you can give the instructions to a computer. Learning how to program is a great way for anyone to improve their problem solving capabilities. As we will see through the course of this book, a lot of the challenge in programming comes from the part that you do before involving a computer at all. You start with a problem and try to determine how you can go about solving that problem. When problems are large, the first thing you typically try to do is break them down into smaller pieces. In fact, a goal that we often strive for in computer science is to never do anything hard. If we have a problem that is big and hard we break it into smaller pieces over and over again until we get down to pieces that are easy to solve.

1.2. First Steps with Greenfoot

Now that we have established what programming is and why you should learn it, we can introduce the environment that we will be using to do our programming through much of this book. For this section you should run Greenfoot on your computer and play along. As with many skills, the only way to learn how to program is to sit down and do it. Abstract discussions can be beneficial, but you won't really understand what is going on until you have gone through the steps yourself. So with that idea in mind, you should start up Greenfoot. If you don't have it on your computer, simply go to <http://www.greenfoot.org/>, download it, and install it. Greenfoot will require that you have Java installed as well. If you get a message saying that you don't have Java, you can download it from <http://java.sun.com/>.

You should also download the zipped scenario file for this book. It contains all of the scenarios that we will be using. After downloading, you want to unzip the file in the scenarios directory where Greenfoot was installed. (If not, simply remember where you unzipped it so that you can get to it to open the different scenarios that we will use.)

When you bring up Greenfoot, go ahead and open up a scenario. To start with we will use the PSPGWombats scenario that was in the zip file. When you have loaded this, you should see a screen like figure 1.



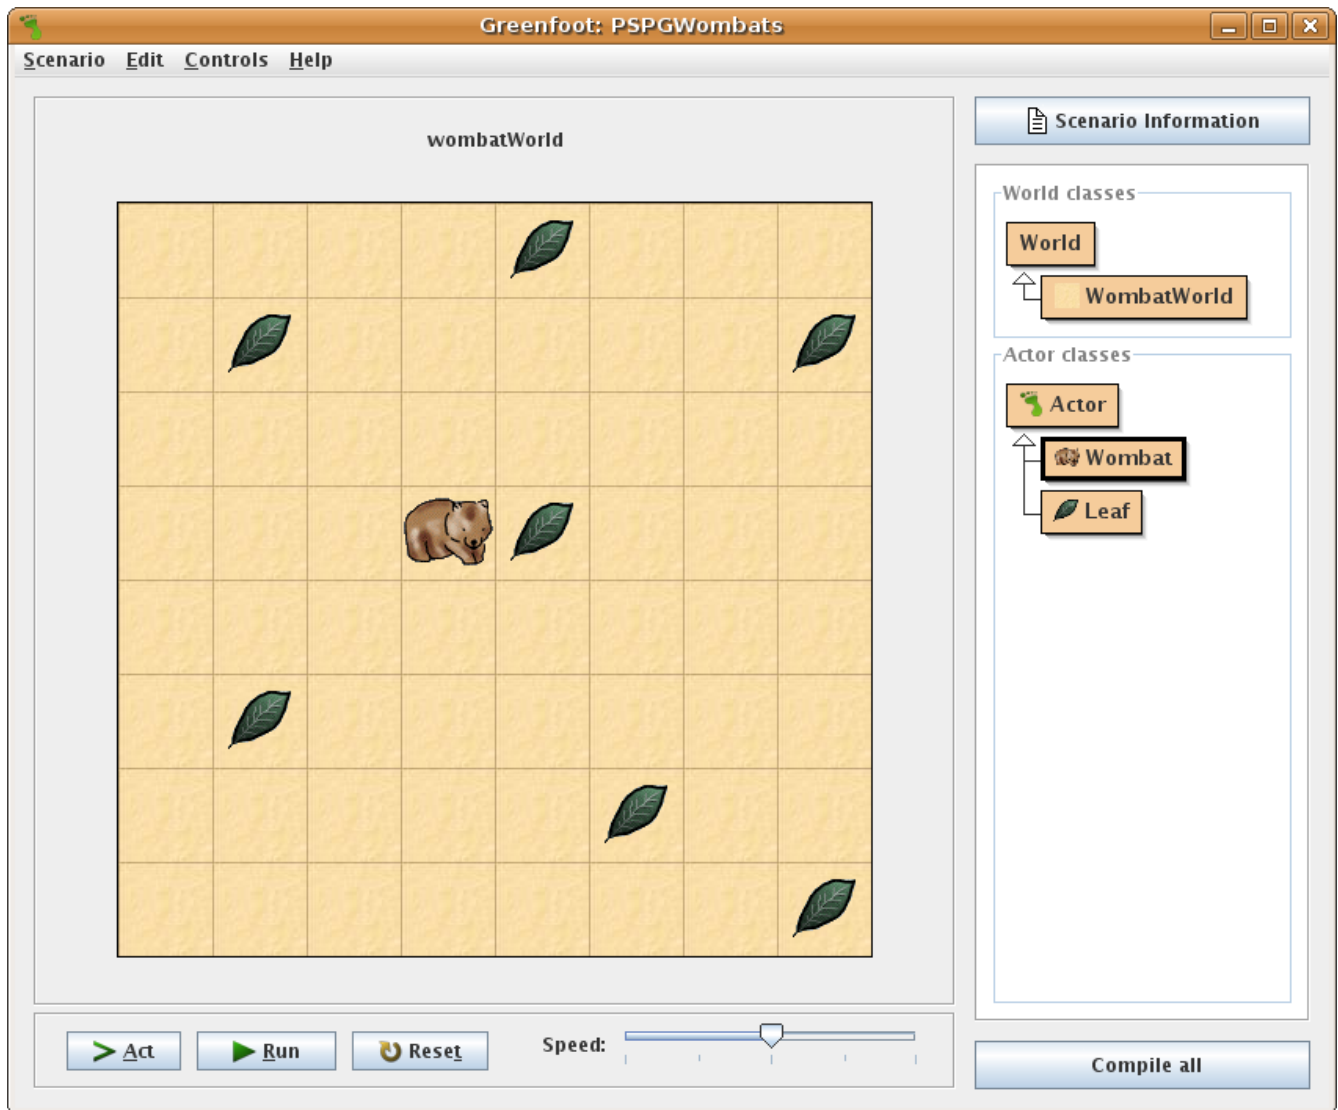
A *scenario* is simply the term used in Greenfoot for how it organizes programs. We will work with many different scenarios over the course of the text and you will even create some of your own. Each scenario contains the text files for pieces of code, images that are used in the scenario, and various other information that Greenfoot uses.

The Greenfoot environment has several different areas in it. The largest part of the screen shows the graphical representation of the world. In the case of the Wombats scenario, this starts off with what looks like a sandy grid. To the right of the world display, Greenfoot lists the different **classes** that are part of this scenario. In the Wombats scenario we see the classes broken into two areas. On top are the “World classes” and on bottom are the “Actor classes”. The classes World and Actor are general and appear in all Greenfoot scenarios. The other classes below them, in this case WombatWorld, Wombat, Rock, and Leaf, are specific to this particular scenario.

In the way that Java works, classes are like blueprints and we can use these blueprints to build objects. One object of the WombatWorld class was automatically created and it is being drawn in the main display. We can also create objects from the actor classes and put those in the world. To see one way this is done in Greenfoot, right click on the Wombat class. You will see a pop-up menu with a

first option of “new Wombat()”. Select this option and then move your mouse over the image of the world. You will see that there is now a wombat on your mouse pointer. Simply click someplace in the world to put your new Wombat object down. This process of making new objects is called **instantiation** and we often refer to the new objects that are created as **instances**.

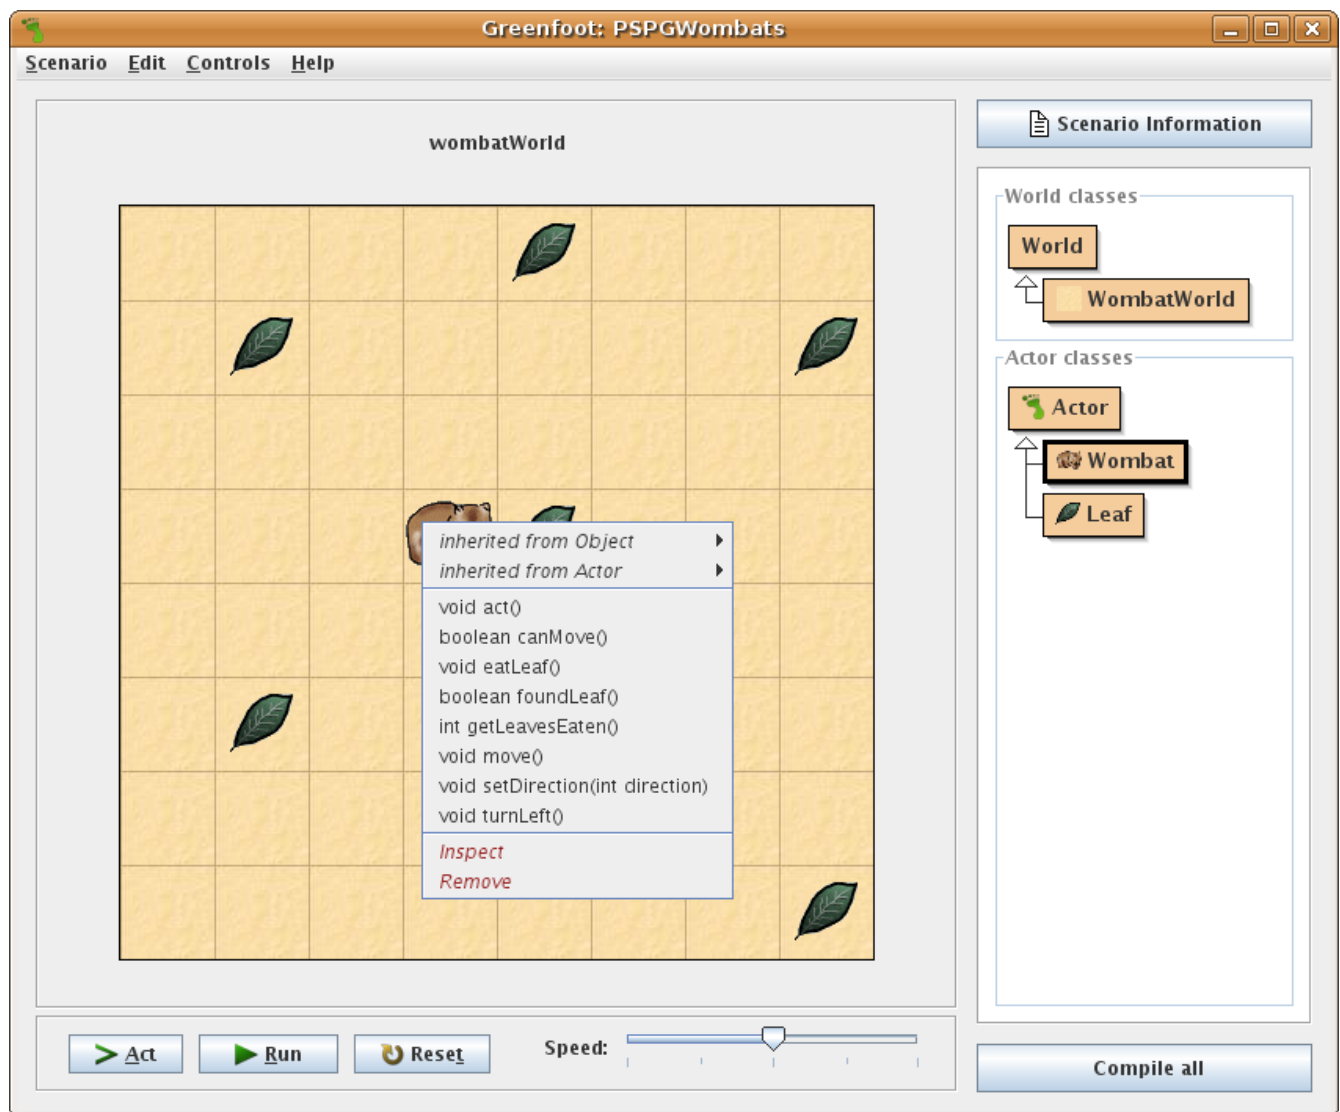
Creating a lot of instances this way is inefficient, so Greenfoot includes a shortcut that can be handy when you want to make a lot of the same type of object. Simply move your mouse over the world display and press the Shift key. An instance of whatever class is selected will appear by your mouse pointer. Simply click to add it to the world. Try doing this with the Leaf class. If we do it a few times, we can create a world that looks like the image below.



We now have a world with a set of actor objects in it. The question is, what can we do with that? That actually depends on the world class as well as the classes of the various objects we have put in the world. Just like objects in the real world, objects in our programmed world have various characteristics and behaviors, which are programmed into them. The characteristics, often called **properties** hold information about the objects or can represent memories the objects store. The

behaviors are called **methods** and they define what an object can do or allow us to ask an object for information about it. To learn about the properties or methods of a particular object, simply right click on it. Right clicking on the name of the world at the top gives you the options for the whole world (the world object, not all the options of the objects *in* the world).

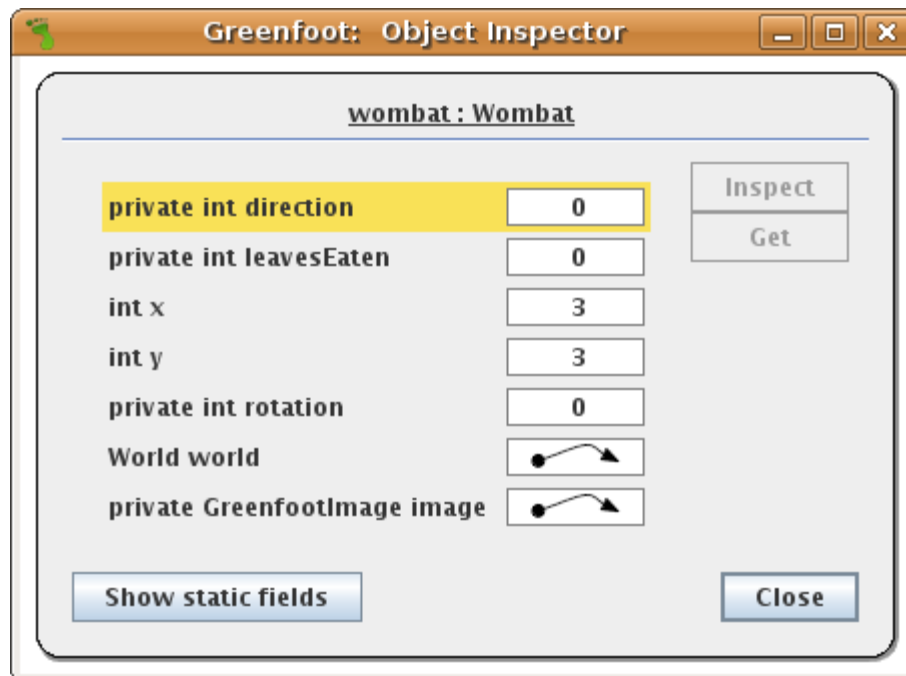
The figure below shows the menu of options that you get if you right click on a Wombat object. Most of the menu is a list of the different methods you can call on a Wombat. These methods are written in the notation and style of the Java language and as a result they might appear a bit odd at first. Nothing about them is very complex though. Notice that each method has at least two words with a set of parentheses. The first word tells us what type of information, if any, the method gives us back. The word *void* implies that no information is given back. The second word is the name of the method. So the method `void act()` tells the Wombat to act and doesn't give us back any information. The method `boolean canMove()` lets us ask the Wombat if it can move and it returns the answer to that question. Try calling some methods on the Wombat to see what happens.



One of the methods of the wombat, `setDirection`, is unlike any of the others in that it has something inside of the parentheses. As the name implies, this method lets us tell the

wombat to face in a different direction. The thing inside of the parentheses, “int direction”, is called a *parameter* and it allows us to give additional information to a method. In this case, we have to tell the wombat what direction we want it to face and we are passing in that information as an integer value. Try calling `setDirection` a few times and pass it some different values to see what happens.

At the end of the menu are two options in red that are not methods you can call. These are `Inspect` and `Remove`. The “remove” option does what one would expect, it takes the object out of the world. The “inspect” option lets you see all of the properties of a particular object. The figure below shows what we get if we select this option on our wombat. Each property of the wombat is shown on a row. It includes the types of property, the name of the property, and the value of the property. Some properties also include the keyword `private` that we won't worry about yet. We can see from the figure that the wombat stores a direction as an `int` and it has a value of 0. The type `int` is short for integer and can represent information that is whole numbers. The wombat stores four other ints as well, `leavesEaten`, `x`, `y`, and `rotation`. Below all of those we see two other properties called `world` and `image`. These are fairly self-explanatory and we will hold off on doing a full discussion of them until later. There are some other options in the Inspector that we will also ignore for the time being.



There are two other areas in the full Greenfoot view that we should talk about. Below the world display you find a set of controls with three buttons and a slide bar. The first button is labeled “Act”. When this button is pressed, Greenfoot goes through every object in the world and calls each one's `void act()` method in turn. We saw above that our Wombat had such a method. That method is actually common to all objects of type `Actor`. Click this button and see what happens.

Next to the `Act` button is a button with the label “Run”. When you click this button Greenfoot will basically do what happens with the `Act` button, but it will do it over and over again. How fast it does it depends on the setting on the “Speed” slide bar. Try clicking the `Run` button to see what happens. You can click it again to stop the acting.

Between the `Run` button and the slide bar is a button with the text “Reset”. If you click this button, the world will go back to its original state. For the Wombat's world that means that all your

actors will disappear and you will be left with the blank grid once again.

The last elements are buttons at the top and bottom on the right side of the window. At the top is a button that will show you information about the scenario. At the bottom is a button with the text “Compile all”. In the previous section we introduced the notion of compiling. We said a compiler was a program that could translate between a language that humans can understand fairly well to one that the computer can understand very well. Whenever we make changes to a program, we have to compile the program so that the new version can be run on the computer. In Greenfoot this is done by clicking the “Compile all” button.

Before we look at a specific problem, I want you to click the Run button and let the program run for a while. You should see the Wombat walk forward and eat a Leaf, then go around the edge eating any leaves that are on the edge. The Wombat will then continue to walk around the edge until you click on Pause. After you have watched the Wombat go around a few times go ahead and pause it. Now we want to call one last method on the Wombat. This time you should call the `getLeavesEaten()` method. Note that this method has a return type that isn't void. It returns an int. This makes sense because it should tell you the number of leaves this Wombat ate. What number does it tell you? Does this match with what you saw? Methods like this that have return types other than void are the way in which we can request information from an object.

1.3. First Problem

Now that you have some idea of how to use Greenfoot, we want to see how to solve a problem with it. As is fitting, we are going to solve a fairly simple problem. You probably noticed that the wombat has a method called `eatLeaf`. If this method is called when the wombat is on the same square as a leaf, the leaf disappears and the wombat adds one to its total of leaves eaten. As it happens, the `act` method will call `eatLeaf()` if the wombat happens to be on a square with a leaf. The problem I want you to solve is as follows: Given a setup like the one above, make the wombat eat all of the leaves.

If you clicked the Run button earlier you probably saw that our wombat is not all that intelligent. When its `act` method is called one of three things happens. It either eats a leaf, moves, or turns to the left. If it is on a leaf it will eat it. If it can move forward it will. If it can't move forward it will turn left. This simple behavior causes it to walk until it hits an edge, then go around the edge in a counter-clockwise direction. This will get some leaves, but not all of them. What do you have to do in order to get the wombat to eat all the leaves? You should only use the Act and Run buttons or directly call methods on the wombat by right clicking on it. Do not drag the wombat or the leaves around to accomplish this.

Direct the wombat so that it eats all the leaves. You might consider writing down each of the steps you followed to make it happen. What instructions did you give to the wombat? What order were they given in? Congratulations! What you have written basically constitutes your first program, though it would likely need to be cleaned up some before it could be understood by a computer. Let's go over your solution to the problem and see how good a job you did.

How complete and accurate are your instructions? If you were to give this to another person, would they be able to follow your steps and get the wombat to eat all the leaves? What if the person in question were a three year old with point and click skills?

How flexible is your solution? Imagine if we added a few leaves or moved some of the leaves around. Would your set of instructions still work? Technically you weren't asked to write something that would work on any set of leaves, only the set that you started with. However, one of the goals we strive for when we are programming is to make our programs as general as possible. Even if the program as a whole is very specific, we often like the pieces we break it into to be general. That way

we can reuse them later on. Writing general code can save you a whole lot of time in the long run.

Think for a minute about how you could go about making your instructions general so that they would work for any arrangement of leaves. Now remember that we want to give these instructions to a three year old. What things do you need to be able to say in your instructions? What ideas do you need to be able to express?

One last question about this problem: was the set of instructions you wrote *optimal*? To answer this we must first say what we mean by the word optimal. For this situation a fairly good meaning would be that we want to make as few method calls on the wombat as possible between the start of the program and when the last leaf is eaten. If you are clicking on buttons or using right clicks to call methods, this definition of optimal is actually quite practical.

Here again, you had not been told to find an optimal set of instructions, just some set. However, as we will see over time, one of the things computers can be very useful for is finding optimal solutions to problems. In this case you might have found a fairly good solution, but how could you be certain it was optimal? As a last point to consider, imagine being asked to make your solution both general and optimal. That is exactly what we will want to do with our programs most of the time.

1.4. A Second Problem

For our second sample problem we are going to need to use some scenarios that are not part of the standard Greenfoot install. You can download these from this book's website. There are three different scenarios that we are going to use: DotGame1, DotGame2, and DotGame3. Begin with DotGame1. Download the file and unzip it. If you can put it in the scenarios directory where Greenfoot is installed, that keeps things simple as all your scenarios will be in the same place. If you can't, just remember where you unzip it so that you can load it in Greenfoot.

When you load the scenario you will be presented with a grid of colored dots. One of them will have a box drawn around it. Click the run button and then “play” the game. It is a fairly standard theme where you have to line up colors. Make sure you understand what it is doing and how it works. You might consider dropping the speed some to make it easier to see what happens when you make moves.

Now that you have some idea of how this program works, load up DotGame2. This will start off looking pretty much identical. However, clicking the Run button on this version doesn't do anything. Unlike the last version, this isn't a functioning game. The act method, which we saw earlier and which is so important to making things happen automatically in Greenfoot, doesn't do anything in this version. Instead, with this version you have to right click on the world or actors in the world to make things happen. Spend a little time looking at the right click menus to see what different methods the world and the actors have. Once you feel that you have a grasp of them, try playing the game this way.

What methods do you have to call on what objects? What order do you have to call them in order to make things happen the way that they are supposed to? Take a few minutes and write down what you have to do in order to complete one “move” in the game. Pay attention to the order of what you have to call and how many times it has to be called. Some methods need extra information passed in as parameters. Which are those and what values do you need to pass in?

Now open up the third version. Do exactly what you did for the second version. Look at the methods that are available and figure out how to play the game. What has happened between the second and third versions? What does it look like now if you write down the set of instructions that you have to call in order to make the game play?

For this last version we !!!

The same basic thing is happening when you go through and play the game in all three versions of this program. The difference is that in the first version all of the instructions have been put into the program. In the second version, some of those instructions were taken out. Some of the methods were basically removed. As a result of this, you have to go in and duplicate what the earlier code was doing. In the third version, even more functionality was removed and you have to manually make calls on different methods had been called by the stuff that was taken out.

What you are seeing here is a specific example of how problems are broken down into pieces. Even the steps in the third version are at a fairly high level. They are composed of smaller steps that deal with individual pieces in the game. As a last step, take the instructions you wrote down for version 2 and compare them to the instructions for version 3. What do you notice when you do this?

So what was the point of this exercise? The first objective was to get you to think through solving a problem in a systematic, step-by-step way. If you went through and really played the game in each scenario and wrote down what methods you had to call, with what parameters, and in what order, then you will have accomplished that objective. There is another objective as well. Unlike the Wombat moving around and eating leaves, this example clearly has different levels of complexity. You can see how the problem breaks into a few high-level steps and then explore how each of those steps can be further broken down. You should have seen how a single method call that you made in version 2 gets split out into many different calls in version 3. Further, you can hopefully picture in your head how each of the calls in version 3 would be further broken down.

This approach to problem solving, where we begin with the full problem and break it into a few large pieces, then break down each of those pieces, and so on is often called top-down design. This is in contrast to bottom-up design where you start with certain basic operations that you know are going to be significant in solving a problem and put them together into successively larger pieces until you get a solution for the whole thing. In real-world problem solving we often use a combination of these two approaches, using whichever fits better at any given time.

!!! Include talking about the order of execution with the outlines. You have to crack the egg before you blend it.

Exercises

1. Don't know.
2. There were lots of questions in the chapter that I could refer back to and make into formal exercises.
3. There were several things we discussed about at lunch at Freddie's one fine day. Did you make note of them somewhere?

Chapter 2 - Basics of coding

What we did in the last chapter was helpful for getting familiar with Greenfoot and it was illustrative of the type of thinking that we do when we are programming, but it was not programming. Technically, we haven't yet seen any of the Java programming language. In this chapter we will start to learn Java and really program.

2.1. Looking at Code

Let's start Greenfoot and open up the PSPGWombats scenario again. To look at the Java code that defines a particular class in Greenfoot you can either right click on the class box on the right side of the Greenfoot display and select "Open editor", or you can double click on the class. We want to look at three different classes in this scenario and we will move in order of increasing complexity beginning with the Leaf class.

When you bring up the editor for Leaf you should get a window that looks like the figure below. If you recall, the Leaf doesn't really do much of anything, and perhaps now you can see why, there isn't much to this class. Let's go through this step-by-step to see what we have. Lines 8-16 actually contain the class itself. We know this because on line 8 it tells us that it is declaring class Leaf. Notice how some of the words are colored. Most programming editors do this for you because it helps you follow code and also to see when you have made an error. The words that are colored are keywords that have to be spelled and capitalized properly. If you mistype one it won't be colored in properly which is easy for you to see.

```
1 import greenfoot.*; // (World, Actor, GreenfootImage, Greenfoot and MouseInfo)
2
3 /**
4  * This class represents a leaf in the Wombat world.
5  *
6  * @author mlewis
7  */
8 public class Leaf extends Actor
9 {
10     /**
11     * The leaf does not do anything.
12     */
13     public void act()
14     {
15     }
16 }
17
```

We will go through what each piece of this code means, but first we get to see our first “rule” for Java code. These will appear frequently in the text. They give you a fairly formal specification for how something should be written in Java. The text in italics refer to other rules that appear elsewhere in the chapter. Things that are in square brackets are optional. This rule is actually a simplification of the full rule for the Java language. It has been simplified to a form that better fits our needs early on. For that reason, this rule, along with some others in this chapter, will be changed later to provide more complete specifications.

Class:

```
public class Name [extends ClassType] {  
    [Constructors]  
    [Methods]  
    [MemberData]  
}
```

The keywords on line 8 are public, class, and extends. They also appear in the first line of our rule. The keyword public is discussed below with its usage on line 13. The class keyword just tells Java that this is a class. The name of the class, Leaf, follows the keyword. The extends keyword implies something called inheritance and the name of the class we are inheriting from, in this case Actor, follows it. Inheritance has a number of different implications, but for now all you need to know is that this says that the leaf is an Actor. This is what we want because anything we want to put into the Greenfoot world has to be an Actor. Inheritance will be discussed in detail in chapter 4.

Lines 9 and 16 have curly braces on them. Because the open curly brace on line 9 follows the class declaration, and the close curly brace on line 16 matches up with it, these two signify the range of everything that is in the class. Java uses curly braces for this purpose in many contexts. Recall the outline format that was discussed for instructions in the previous chapter? The curly braces are the formal way that we tell Java that certain things are grouped together in a certain **scope**. Note that lines 10-15 are indented, just like they would be in an outline. This isn't required. It turns out that Java doesn't pay attention to spacing, indentation, or even new lines, but indentation makes code much easier for humans to read. For this reason, all the code in this book will be indented such that nested lines are tabbed in one level beyond the code that it sits inside of. It is highly recommended that you do this for your own code as well. While Java will understand the code with or without indentation, you will find in time that your ability to understand it is greatly enhanced if everything is properly indented.

Lines 13-15 represent the only things in this particular class. They define a method called act. You saw the act method in the last chapter when we called it on the wombat. When you pressed the Act button you noticed that the leaves didn't do anything. Now you can see why. The act method is blank. Once again we see the curly braces as well. In this case, nothing is nested inside of these curly braces. The word public appears here again as well as a new keyword, void. At the end of the line, the word act is followed by parentheses. Here is a basic rule for a method. This act method isn't using much of it.

Method:

```
public Type name([Type1 name1[,Type2 name2, ...]]) {  
    [Statements]  
}
```

We want to go through this one piece at a time and we will start with public. The public keyword tells Java that something is open and available to any code. To understand what this means, think of a television set. The main use of a television is watching, but you also have to occasionally tell it to do something. What can you tell a television? Well, turning on and off is a big one. Changing channels and/or input sources and changing the volume are also common options. Most modern TVs also have a number of other settings. Each of these things can be done by pressing a button or a set of buttons either on the TV itself or on the remote of the TV. There is a lot of stuff hidden inside the TV as well. In theory you could change a lot of different things if you took off the case and started poking around. Of course, odds are good that if you did so you would break something and the TV would never work quite right again.

The buttons on the TV and the remote are the public elements of the TV. They are the things that anyone else is allowed to interact with. They are public in part because they are safe. The makers of the TV only exposed things to you that won't give you the power to break the TV. All the stuff hidden away inside of the TV is private. We will see shortly that private is also a keyword in Java and, just like the TV, signifies things that other people should not be messing with, often because messing with it could put an object into a state in which it doesn't work well.

After public on line 13 is the word void. If we look at our rule for methods, this is in the place called *Type*. This is the return type of the method. It tells us what information the method gives back to us when it is called. The type void indicates that no information is returned. Methods with a void return type simply do something, they don't return any information to us. This makes sense for the act method. When we call act, something should happen, but in general act doesn't need to tell us anything when it is done. The rule for *Type* is below. It uses a new symbol for the rules. The | symbol, called a pipe, is like the word "or". It tells us that one thing or another can go in a certain place. Looking at these rules we see that for act *Type* is a *PrimitiveType*, and it is using the last option under *PrimitiveType* of void.

Type:

PrimitiveType | *ClassType*

PrimitiveType:

boolean | char | int | double | void

The type void can only be used as the return type of a method, not for a variable declaration.

ClassType:

This is the name of a class. It can be one that you have written or one that is in an library. Class names typically start with a capital letter and have each new word capitalized. They follow the normal rules for Java names.

Both the class and the method have referred to another rule called *Name*. When this appears it simply expects a valid Java name. The rule below presents the options you have for making Java names. While the rule is incredibly flexible, you typically won't use all of that flexibility, nor would you want to. It is best to give things names that are easy to remember and that have meaning to them. That way when you come back to the same code you will be able to more easily understand it. Java also distinguishes between upper and lower case letter so the names count, COUNT, and CoUnT are all different. It is generally not advisable to create names that only differ in their capitalization though as it can get very confusing. There is something of a standard that people try to follow when it comes to

naming in Java. Class names begin with a capital letter. The names for non-types, like methods and variables, which we will see later, start with a lower-case letter. In both cases, the first letter of the second, third, etc. word in the name is capitalized. This style is given the name camel because the capital letters on each word look like humps on a camel. All the examples in this book will adhere to the style described here.

Name:

Java is fairly flexible about naming. Valid Java names have to start with a letter or underscore followed by any number of letters, underscores, or numbers.

The parentheses after the word `act` on line 13 indicate that this is a method. From the Method rule we can see that it is possible to put parameters inside of the parentheses. We will see these a lot more later. One of the things you will come to learn about programming languages is that they often have a lot of punctuation¹ and they are very picky about it. So as we go through looking at various pieces of code, pay particular attention to the placement of punctuation.

What about the lines before line 8? The first line has the word `import` in red followed by “`greenfoot.*;`” and then some stuff in gray following two slashes. The `import` statement tells Java where to go looking for things. In this case, it is telling Java that if it ever sees something it doesn't know about, it should go look in the `greenfoot` library. This makes sense as this class is part of a `Greenfoot` program. Most of our classes that we write in `Greenfoot` will have a line like this at the top.

And what about the stuff in gray after the double slashes? That is a comment. The double slashes signify a single line comment in Java. Anything following them on that line is a comment. You can type whatever you want in comments because the Java compiler ignores them. Lines 3-7 are also comments. These are multi-line comments that begin with `/*` and end with `*/`. The fact that this one begins with `/**` just means that this is a special type of comment and is a detail that we won't worry about for now. The compiler doesn't care. As soon as it sees the `/*`, it ignores everything until it gets to a `*/`. For this reason, you will never see a comment appear in any rules. They could go pretty much anywhere except in the middle of a word.

Comments can be a vital part of programs. They are typically used as documentation to help make it clear what a piece of code is doing. There are different philosophies on how much and when comments should be used though. They range from very minimal use of comments to extensive commenting that documents all aspects of a piece of code. The logic behind the choice of how to comment will be discussed in various parts of this book after you have a stronger foundation in programming to understand them. As a beginner though, you should probably adopt the commenting strategy that is accepted by the person giving you your grade.

Before we move on to another class, you should look at the editor window shown in the last figure. In addition to the menus, there are a set of buttons that include commonly used editor tasks such as cut and paste and a button to compile this file. You might want to familiarize yourself with the options in the menus. Some of them can help to make your life much easier when we get to the stage of writing code.

So that was the `Leaf` class. It was simple. It didn't really do anything. That is what made it a good place to start. Now we move on to the `WombatWorld` class. This class is a fair bit longer. We will just look at the code and not view it as an image in the editor. It starts off much the same as `Leaf`. There is an `import` statement of `greenfoot.*` at the top, a comment, and a public class declaration. While the `Leaf` class inherited from `Actor`, the `WombatWorld` class inherits from `World`. This means

¹ Some languages have less punctuation, but they typically have keywords in place of the punctuation. For example, the Pascal language used “begin” and “end” in basically the same way that Java uses “{“ and “}”.

that WombatWorld is a World.

```
import greenfoot.*; // (World, Actor, GreenfootImage, and Greenfoot)

/**
 * A world where wombats live.
 *
 * @author mlewis
 */
public class WombatWorld extends World {
    /**
     * Create a new world with 8x8 cells and
     * with a cell size of 60x60 pixels
     */
    public WombatWorld() {
        super(8,8,60);
    }

    /**
     * Populate the world with a fixed scenario of wombats and leaves.
     */
    public void populate() {
        addObject(new Wombat(),3,3);
        addObject(new Leaf(),4,0);
        addObject(new Leaf(),1,1);
        addObject(new Leaf(),7,1);
        addObject(new Leaf(),4,3);
        addObject(new Leaf(),1,5);
        addObject(new Leaf(),5,6);
        addObject(new Leaf(),7,7);
    }

    /**
     * Place a number of leaves into the world at random places.
     * The number of leaves can be specified.
     * @param howMany The number of leaves to add.
     */
    public void randomLeaves(int howMany) {
        for(int i=0; i<howMany; i++) {
            Leaf leaf = new Leaf();
            int x = Greenfoot.getRandomNumber(getWidth());
            int y = Greenfoot.getRandomNumber(getHeight());
            addObject(leaf, x, y);
        }
    }
}
```

The WombatWorld class really has stuff inside of it. To be specific, it has a constructor and two methods inside of it. The constructor is something new that we didn't see in the Leaf and a rule for it is presented below. In general, constructors are very much like methods. The main differences are that they don't have a return type, not even void, and they have to have the same name as the class they are in. So in this case the constructor has the name WombatWorld.

Constructor:

```
public ClassName([Type1 name1[,Type2 name2, ...]]) {  
    [Statements]  
}
```

The purpose of the constructor is to set up values the way we want them when an object of this class is created. The constructor is also the first place where we have seen a statement inside of the curly braces. In this case this is a special statement using the keyword `super`. We won't get into the details of `super` at this point, but you should know that this line of code makes it so that our `WombatWorld` is an 8x8 grid with each grid being 60x60 pixels.

It should be explicitly mentioned at this point that you should not worry if everything you see doesn't make sense right now. We are going through existing code to give you a feel for what you can put into Java code and what you can do with it. Some features you aren't meant to fully understand yet. Shortly we will start writing our own new code. When we do that it will be done at a pace where each step should be explicit and completely understandable. Making everything work the way we want in `PSPGWombat` simply invokes a few elements that we won't get into the full details of for a few chapters.

The other two methods in `WombatWorld` are more interesting. The first one, `populate`, sets up the world in the configuration that was shown in chapter 1. It has eight statements in it. The rule for a statement is given below and shows that there are several different options for statements. All the statements in `populate` happen to be method calls so that is the only specific statement rule that is presented. The others will wait for later. One thing to note about statements is that they tend to end in a semicolon. The one exception to that rule is the *CodeBlock*. Remember that Java doesn't pay much attention to white space, including new lines. Just like a sentence in English can be spread across multiple lines and only ends when you get to the proper punctuation mark, statements in Java can also span multiple lines and only end when you get to the semicolon. Unlike English though, we will typically put one statement on a line in Java and will only split statements across lines if they happen to get really long.

Statement:

VariableDeclaration | *Assignment* | *MethodCall* | *ReturnStatement* | *CodeBlock*

MethodCall:

```
[objectExpression .]methodName([expression1 [, expression2, ...]]); |  
ClassName.methodName([expression1 [, expression2, ...]]);
```

The `objectExpression` is any expression whose type is an object type. If no `objectExpression` is given, the method is called on the current object. The `methodName` must be a valid name of a method on that type. The expression is often just the name of a variable, but it can itself be a method call that returns an object.

The second format is used for static methods. For these we don't need an object, we just put the name of the class in front of the dot. The `Greenfoot` class has only static methods in it so you call them with a format like `Greenfoot.getRandomNumber(3);`.

The eight statements in `populate` all call the same method, `addObject`. There are three different arguments passed into `addObject` in all of these calls. The first is a new object that is an `Actor`. The second and third are numbers, integers to be more specific. The meaning of `addObject` is fairly clear.

It is supposed to add an object into the world. Note how it uses the camel naming scheme and because it is a method, it starts with a lower case letter. The meaning of the arguments that are passed in isn't quite as clear. They happen to be the location that we want the object to be placed at. The x-coordinate is first and the y-coordinate is second. How do we know that though and how can we find out what other methods we might be able to call? The solution to these problems is the API. That stands for Application Programmers Interface. What it really means is that this is the help documentation that tells you all the methods and other things that are in particular classes.

Because we are programming in Greenfoot, methods like `addObject` can be found in the Greenfoot API, which you can find under “For Programmers” at the Greenfoot web site (<http://www.greenfoot.org/>). There only five classes in the Greenfoot API. This is one of the best things about Greenfoot, the API is rather simple. So which one of those five classes is the `addObject` method in? If you read the rule for *MethodCall* and compare it to the calls to `addObject` you see that because we don't have `something.addObject` that these calls are being made on the current object, which, because we are in the class `WombatWorld`, has type `WombatWorld`. Looking at the `WombatWorld` class we see only three things, and `addObject` is not one of them. The solution is at the top of the class where it says `WombatWorld extends World`. Earlier we said that this implies that `WombatWorld` is a `World`. One implication of this is that `WombatWorld` implicitly has all of the methods that are part of `World`. The figure below shows a screen shot of the `World` API page and sure enough, there is an `addObject` method in it.

The screenshot shows a Mozilla Firefox browser window displaying the JavaDoc page for the `greenfoot` package, specifically the `World` class. The browser's address bar shows the URL `http://www.greenfoot.org/doc/javadoc/`. The page layout includes a navigation menu at the top with links for `Package`, `Class`, `Tree`, `Deprecated`, `Index`, and `Help`. Below this, there are links for `PREV CLASS`, `NEXT CLASS`, `FRAMES`, `NO FRAMES`, `SUMMARY: NESTED | FIELD | CONSTR | METHOD`, and `DETAIL: FIELD | CONSTR | METHOD`.

The main content area is titled `greenfoot` and `Class World`. It shows the class hierarchy: `java.lang.Object` is the superclass, and `greenfoot.World` is the subclass. The class is defined as a `public abstract class World` that extends `java.lang.Object`.

A descriptive paragraph follows: "World is the world that Actors live in. It is a two-dimensional grid of cells. All Actor are associated with a World and can get access to the world object. The size of cells can be specified at world creation time, and is constant after creation. Simple scenarios may use large cells that entirely contain the representations of objects in a single cell. More elaborate scenarios may use smaller cells (down to single pixel size) to achieve fine-grained placement and smoother animation. The world background can be decorated with drawings or images."

Metadata includes `Version: 2.1` and `Author: Poul Henriksen, Michael Kolling`. A `See Also:` section points to the `Actor` class.

Two summary tables are provided:

- Constructor Summary:**

<code>World</code> (int worldWidth, int worldHeight, int cellSize)	Construct a new world.
--	------------------------
- Method Summary:**

void <code>act</code> ()	Act method for world.
void <code>addObject</code> (<code>Actor</code> object, int x, int y)	Add an Actor to the world.
<code>GreenfootImage</code> <code>getBackground</code> ()	Return the world's background image.
int <code>getCellSize</code> ()	Return the size of a cell (in pixels).
java.awt.Color <code>getColorAt</code> (int x, int y)	Return the color at the centre of the cell.
int <code>getHeight</code> ()	Return the height of the world (in number of cells).

The `addObject` method is the second method in the method summary table. To the left of the method name is listed the return type of the method. This says that `addObject` returns `void`, which means that it doesn't give us back any information. In the parentheses to the right of the method name you find the argument list. This tells us that we are supposed to pass three things into `addObject`. The first is something of type `Actor` which in this scenario means either a `Wombat` or a `Leaf`. After that we see that we are supposed to pass in two integers, `x` and `y`. The API makes it clear what we are passing into the calls to `addObject` and what each one means. The API is a very helpful tool and we will be referring back to it frequently.

There is only one more thing to populate that we haven't explained. The first argument to each call of `addObject` starts with the keyword `new`, followed by a class name and parentheses. The `new` keyword tells Java that we want to make a new object. The class name is the name of the class that we want to make a new one of. The parentheses are required and can have arguments in them. Using `new` is what actually calls constructors like the one we talked about earlier. If our constructor had

parameters that it takes in, then we would need to provide corresponding arguments when we used `new`.

The last method in `WombatWorld`, `randomLeaves`, is a bit more complex and contains some elements we won't talk about for a few chapters. In particular, the `for` loop won't be discussed until chapter 4. At this point it is sufficient to say that it allows us to tell Java to do something multiple times. Unlike earlier methods, the `randomLeaves` method takes in a parameter. When we call `randomLeaves`, we have to provide it with an integer telling it how many leaves we want to add. This value is stored in the variable with the name `howMany`.

Inside of the `for` loop there are three other variables that are declared so at this point it is probably worth taking a minute to talk about variables and introducing two of the other statement types. Variables in programming are conceptually very similar to their counterparts in math that you learned about in algebra. They are names that we can use in an abstract way to refer to values. They are not perfectly analogous to math variables however because programming variables can change their values as a program runs. A good mental image for a variable is to think of it as a name that we put on a box that can hold a value in it. We are able to look in the box to see the value, or change the value that is in it.

VariableDeclaration:

```
Type name[=expression];
```

This declares a variable of the given type with the specified name. You can think of variables as boxes that can hold a value. The box is labeled with the name we give it and it can only hold things of the proper type. What you can have for type is described after all the statement types are introduced.

A variable can only be used in a certain “scope”. The scope of a variable runs from the point at which it was declared down to the close curly brace that ends the code block it is declared in. Consider the following example method.

```
public void foo() {  
    int a=5;  
    {  
        int b=8;  
        // Point 1  
    }  
    // Point 2  
}
```

At point 1 you could use either `a` or `b`. However, the scope of `b` is only inside the inner curly braces so you can't use `b` at point 2, only `a`.

Assignment:

```
name =expression;
```

This stores the value of an expression in the variable with the given name. The variable must be in scope. So it either has to be declared in the current method above this line or it should be a class level variable. In order for this to work, the types need to match.

Here are some example variable declarations and assignments that could be made from them.

```
int a;
int b;
double d;    // This type will be discussed later in the chapter.
String s;    // This type will be discussed later in the chapter.

a=6+7;
b=a-4;
d=2*3.14159;
s="This is a test.";
```

Another concept that is very important to variables in programming is type. We have already seen the rules for types and have had it pop up in other contexts. The first word in a variable declaration with our current rule is the type of the variable we are declaring. This is significant, because variables can only store values of the correct type. The first variable declared inside of the for loop has the type Leaf and the name leaf while the other two both have type int and names x and y. If you were to try to put an integer value in the variable leaf you would get an error. Similarly, if you tried to put a value of type Leaf in x or y, you would also get an error. We will see a lot more about how this matters over time.

At first it might seem confusing to have a type Leaf and a variable called leaf. This is one case of having two names that only differ by capitalization that is fairly common. This particular usage doesn't cause confusion as long as you follow standard naming schemes because any name that starts with a capital letter will be a type while those that start with lower case letter will be either variables or methods.

All three of the variable declarations in randomLeaves include the optional initialization. That is to say that the variable name is followed by an equal sign and an expression of the proper type. This gives puts an initial value into the variable. In the case of leaf, the initial value is simply a new leaf. For x and y the expressions are a bit more complex. The expressions for x and y use method calls to get values. Both have a call to a method named getRandomNumber that is in the class Greenfoot. Take a minute to go look in the API at the Greenfoot class. This class is called a utility class. All the methods in it are static which means that we call them on the whole class, not on some instance of the class. Basically, you will always write Greenfoot.methodName(...) when you call methods from the Greenfoot class.

The getRandomNumber method takes a single argument that is an int. In this case, we want to get numbers between 0 and the size of the world. In the constructor we saw that this world is made to be 8x8. As such, this code could have called Greenfoot.getRandomNumber(8). However, it is generally considered bad form to put in values like this that appear as “magic numbers”. If we put in the value 8 and later wanted to change the size of the world, we would have to find all of the 8s in the code and change them appropriately. Instead, the World class has methods in it called getWidth() and getHeight(). These return to us the width and height of the world. So if we alter the size in the constructor, this code will continue to work. This should remind you of the concept of generality that was discussed in the first chapter some. Using the number 8 would have made the code overly specific. By using these methods the code is made more general. That benefits us later because we can make changes to other things without having to change this code as well.

Assignment statements are also introduced in the rule above. They change the value of a variable that has already been declared. When a method is executed, the lines in that method are

executed in roughly top-down order². If one line declares a variable and give it a value it will have that value until we get to another line that does an assignment and changes the value. The variable will then contain that new value until another assignment statement is reached.

Expression:

An expression in Java is anything that has a type and value. Below is a list of options.

- Variable name – The name of a variable evaluates to the value currently stored in that variable.
- Numeric literal – Numbers are expressions. You can represent integer values or decimal fractions as you normally would. Very large or very small numbers are written using scientific notation of the form 1.234e56 which is the number 1.234×10^{56} , a rather large number indeed.
- String literal – String literals can be any sequence of characters enclosed in double quotes. For example, “Hi mom.” is a string literal.
- null – This is the value for any object type when we don't have a valid object.
- Method call – A method call for any method that returns something other than void is an expression.
- Operator expression – Java supports different mathematical and logical operators. Most will join two different expressions. Some operate on one expression and one uses three expressions. This is a list that includes most of the ones we are likely to use.
 - A+B : Standard addition for numeric values. If either A or B is a String the plus sign will do string concatenation.
 - A-B : Subtraction for numeric values.
 - A*B : Multiplication for numeric values.
 - A/B : Division for numeric values.
 - A%B : Modulo operator that returns the remainder after division for integer values.
 - A=B : This is the assignment we saw above. It is also an expression which takes on the value that is being stored. In this case, A needs to be the name of a variable while B can be any expression of the right type.
 - new *ClassName*(args) : The new operator is used to make objects. This will make a new object of type *ClassName* using the constructor that is matched by the arguments passed in.

The last line in the for loop in randomLeaves calls the same addObject method we saw earlier to add the new leaf into the world at the randomly generated location. You should try calling these two methods in Greenfoot to see what they do. You could also try altering the constructor to have a different sized world and see what happens. Do you notice any interesting behaviors? We talked about the fact that randomLeaves is general. Is populate? Why or why not?

The last class we want to look at in PSPGWombats is the Wombat class itself. The code for it is shown below. This class contains a lot more methods than the WombatWorld class did. We won't bother with going through each and every one here in the text though you should try to look through them and see if you understand what they are doing. There are a few new features that are worthy of discussion and we will go through those.

```
import greenfoot.*; // (World, Actor, GreenfootImage, and Greenfoot)
```

² Some statement types, like the for loop, can alters this a bit and cause code to repeat statements or skip over some statements, but for normal statements, the way the program runs is to do the steps in order one at a time.


```

/**
 * Wombat. A Wombat moves forward until it can't do so anymore, at
 * which point it turns left. If a wombat finds a leaf, it eats it.
 *
 * @author mlewis
 */
public class Wombat extends Actor {
    private int direction;
    private int leavesEaten;

    public Wombat() {
        setDirection(0);
        leavesEaten = 0;
    }

    /**
     * Do whatever the wombat likes to do just now.
     */
    public void act() {
        if(foundLeaf()) {
            eatLeaf();
        }
        else if(canMove()) {
            move();
        }
        else {
            turnLeft();
        }
    }

    /**
     * Check whether there is a leaf in the same cell as we are.
     */
    public boolean foundLeaf() {
        Actor leaf = getOneObjectAtOffset(0, 0, Leaf.class);
        return leaf != null;
    }

    /**
     * Eat a leaf.
     */
    public void eatLeaf() {
        Actor leaf = getOneObjectAtOffset(0, 0, Leaf.class);
        if(leaf != null) {
            // eat the leaf...
            getWorld().removeObject(leaf);
            leavesEaten = leavesEaten + 1;
        }
    }

    /**
     * Move one cell forward in the current direction.
     */
    public void move() {
        if (canMove()) {

```

```

        setLocation(getX()+getXOffset(direction),getY()
                    +getYOffset(direction));
    }
}

/**
 * Test if we can move forward. Return true if we can, false otherwise.
 */
public boolean canMove() {
    World myWorld = getWorld();
    int x = getX()+getXOffset(direction);
    int y = getY()+getYOffset(direction);
    // test for outside border
    if (x >= myWorld.getWidth() || y >= myWorld.getHeight()) {
        return false;
    }
    else if (x < 0 || y < 0) {
        return false;
    }
    return true;
}

/**
 * Turns towards the left.
 */
public void turnLeft() {
    direction=direction-1;
    if(direction<0) {
        direction=3;
    }
    setDirection(direction);
}

/**
 * Sets the direction we're facing.
 */
public void setDirection(int dir) {
    direction = dir;
    setRotation(90*direction);
}

/**
 * Tell how many leaves we have eaten.
 */
public int getLeavesEaten() {
    return leavesEaten;
}

/**
 * Returns the x offset for the specified direction.
 * @param dir The direction we want an offset for.
 * @return The x offset.
 */
public int getXOffset(int dir) {
    if(dir==0) {

```

```

        return 1;
    } else if(dir==2) {
        return -1;
    } else {
        return 0;
    }
}

/**
 * Returns the y offset for the specified direction.
 * @param dir The direction we want an offset for.
 * @return The y offset.
 */
public int getYOffset(int dir) {
    if(dir==1) {
        return 1;
    } else if(dir==3) {
        return -1;
    } else {
        return 0;
    }
}
}

```

The first two lines inside of the Wombat class provide us with something new. These look a lot like the pattern for variable declarations and indeed, that is very close to what they are. These lines declare class member variables. They are also the first place we have seen the keyword `private` which was alluded to in the discussion of `public`. The rule for member data is shown below.

Member Data:

```
private Type name[=expression];
```

This defines a class level variable with the specified name and type. The initial value doesn't have to be specified.

Constructors and methods have a sequence of statements inside of curly braces. As a general rule, curly braces are used in Java to group lines of code. When a method is called in Java, the statements in that method are executed one at a time in order from the top of the method to the bottom. At this point we will only consider five types of statements in Java. So any place that you can have a statement, you can have one of these five. More statement types will be introduced later on.

Here is a simple little example of a class. This one doesn't do much, but it demonstrates the three things that you can put inside of a class.

```
public class SimpleClass {
    public SimpleClass(int v) {
        value=v;
    }
    public int getValue() {
        return value;
    }
}

```

```
}  
private int value;  
}
```

So what is the real difference between the variables we saw earlier, which are often called local variables, and the member variables we are seeing here? The difference is something called **scope**. The scope of a variable determines the region of code over which it can be used and how long it is kept around. Local variables exist inside of the method that they are declared in. Technically, a local variable can only be used from the point at which it is declared down to the end of the code block, denoted with curly braces, in which it was declared. Any attempt to use it outside of that region will be an error. What is more, the variable ceases to exist when you exit that scope and a new one is created when you get back to the declaration point. One of the big implications of this is that a local variable declared in one method can't be used in another. You can declare a variable using the same name, but it will get a separate block of memory. It is somewhat like having two people with the same name. Sharing the same name doesn't make them the same person.

Member data also has a scope inside of the curly braces it is declared in. Since member data is declared outside of any method, and it inside the curly braces of the class as a whole, that is the scope that it has. You should think of member data as the memories of objects of the given class. In this case, every Wombat remembers its direction as well as how many leaves it has eaten. Both of these are integer values. Member variables are initialized by default. Because these are numbers, they are set to be zero. There is also code in the constructor for Wombat to explicitly initialize those values.

If you look through the class code some you will see that both direction and leavesEaten appear in a number of different methods. This fits with the fact that member variables have a scope that goes through the full class. It begs a question though. When should a variable be a member variable and when should it be local. The general rule that we will follow is to make variables have the smallest scope that they can. So while all variables could be made into member variables, doing so generally leads to problems and is considered poor form. We will only make variables into member variables when they are used across multiple methods and the value needs to be remembered from one method to another, or across multiple calls to a single method.

So what other things are new in the Wombat class? Obviously it is longer, and it has a bunch of different methods in it. This is an example of problem decomposition. Look at the act method, it is right below the constructor. It uses the if construct that we won't go into until next chapter, but you should be able to read it and figure out what it is doing. The act method doesn't really do anything directly. Instead, it has calls to five other methods that it could make, depending on the situation, and those other methods do the work. The names for those methods are straightforward and easy to read which makes the act method itself easy to read. It also means that when the act method was written, we didn't have to figure out how to do everything. Instead, we broke it into a few key steps, then went and wrote methods to accomplish each of those key steps.

The first method that act calls is inside of parentheses on the first if. It is the foundLeaf method which can be found right after the act method. This method shows us two other new concepts that we want to look at briefly. Both new concepts stem from the fact that this is the first method we have seen that has something other than void as the return type. In this case, the return type is boolean. We will go into the details of boolean in the next chapter. The simple explanation is that it is a type that is either true or false. Those appear as keywords in the canMove method which also returns a boolean. Having a return type other than void forces this method to include a new type of statement, the return statement. The rule for return statements and the last statement type we haven't yet discussed appear below.

ReturnStatement:

```
return [expression];
```

This statement stops the execution of a method and returns the value of the specified expression. If the method has a void return type, then there isn't an expression. Otherwise there must be an expression and the expression type must match the return type of the method or be convertible to it.

CodeBlock:

```
{ [statements] }
```

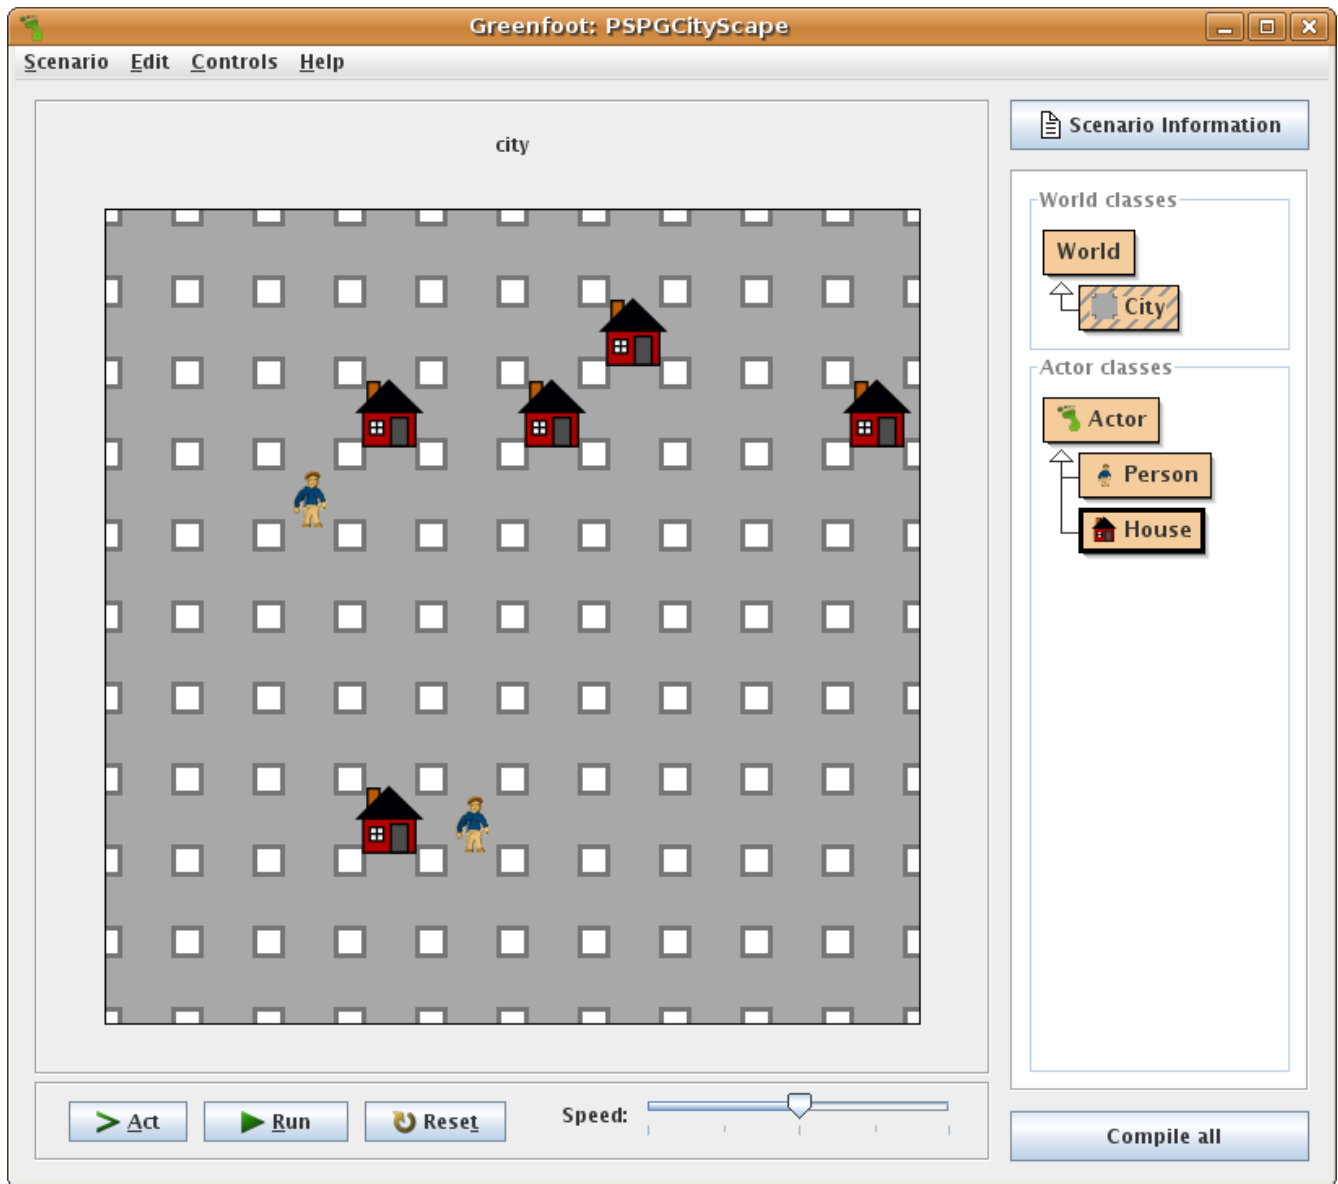
Any place a single statement is expected in Java, it can be replaced by multiple statements inside of curly braces. They don't all have to be on the same line as I have typed it here.

Note that the rule for the return statement implies that the expression is optional. This is only the case if the method returns void. Otherwise, there has to be an expression and the type of that expression has to match the return type of the method. A method can contain more than one return statement. The `canMove` method demonstrates this. Whenever a return statement is reached, the flow of control leaves that method and goes back to whatever called the method. So a return statement with no expression is an effective way to break out of a void method and not do anything that follows it.

None of the other aspects of Wombat are significant to us at this point. You actually have rules to understand everything here with the one exception of the if statements and the conditions that are in parentheses after them. At this point we have spent enough time king at existing code. It is time to write some of our own from scratch.

2.2. Constructing New Code

We aren't going to write our code in the Wombat scenario. Instead, we will start working with a new scenario called `PSPGCityScape`. This is a very basic, simple scenario that we will build up over time to demonstrate a number of different aspects of programming and how we can do things in Greenfoot. When you load up the scenario and compile it, you will get a scene with two people and five houses that might look something like the image below, but might be a bit different. Look at the code in the `City` class to see how this is being set up. Why can yours be different from the figure? Do you get the same setup if you Reset the scenario? What is it in the code that causes that?



Try clicking Act or Run. What happens? Nothing. If you go into the code you will see why. The Person and the House classes have nothing in them except for an empty act method. In fact, the code that you are seeing is the default code generated by Greenfoot when you make a new Actor class. We are going to modify the Person class and make it so that the person moves when the act method is called. First you should modify the comments so that it says that you are the author and it describes the class that you are writing. To make sure that you didn't alter anything that breaks the code, press the Compile button on the editor or the Compile All button on the main Greenfoot window. If there is an error you will be told about it.

Now we want to modify the act method in Person to make it so that the person moves one square to the right each time that act is called. How are we going to do this? How can we make the person move? There clearly aren't any methods in the Person class that will do this for us. Remember in the last section where we saw that WombatWorld got to use the methods in World because it inherited from that class? Person inherits from Actor. Look up the Person class in the API and see if there are any methods that could help with moving an Actor. The figure below shows the summary of

methods in the Actor page of the API.

The screenshot shows a web browser window titled "Actor (Greenfoot API) - Mozilla Firefox" with the URL "http://www.greenfoot.org/doc/javadoc/". The page displays the "Constructor Summary" and "Method Summary" for the Actor class. The "Constructor Summary" section includes the `Actor()` constructor. The "Method Summary" section lists 17 methods with their signatures and descriptions:

Method Signature	Description
<code>void act()</code>	The act method is called by the greenfoot framework to give objects a chance to perform some action.
<code>protected void addedToWorld(World world)</code>	This method is called by the Greenfoot system when the object has been inserted into the world.
<code>GreenfootImage getImage()</code>	Returns the image used to represent this Actor.
<code>protected java.util.List getIntersectingObjects(java.lang.Class cls)</code>	Return all the objects that intersect this object.
<code>protected java.util.List getNeighbours(int distance, boolean diagonal, java.lang.Class cls)</code>	Return the neighbours to this object within a given distance.
<code>protected java.util.List getObjectsAtOffset(int dx, int dy, java.lang.Class cls)</code>	Return all objects that intersect the given location (relative to this object's location).
<code>protected java.util.List getObjectsInRange(int radius, java.lang.Class cls)</code>	Return all objects within range 'radius' around this object.
<code>protected Actor getOneIntersectingObject(java.lang.Class cls)</code>	Return an object that intersects this object.
<code>protected Actor getOneObjectAtOffset(int dx, int dy, java.lang.Class cls)</code>	Return one object that is located at the specified cell (relative to this objects location).
<code>int getRotation()</code>	Return the current rotation of the object.
<code>World getWorld()</code>	Return the world that this object lives in.
<code>int getX()</code>	Return the x-coordinate of the object's current location.
<code>int getY()</code>	Return the y-coordinate of the object's current location.
<code>protected boolean intersects(Actor other)</code>	Check whether this object intersects with another given object.
<code>void setImage(GreenfootImage image)</code>	Set the image for this object to the specified image.
<code>void setImage(java.lang.String filename)</code>	Set an image for this object from an image file.
<code>void setLocation(int x, int y)</code>	Assign a new location for this object.
<code>void setRotation(int rotation)</code>	Set the rotation of the object.

If you go through this list you notice that there are two main classes of methods. There are 11 methods that begin with the word get and four methods that begin with the word set. The get methods allow us to ask the Actor for information about it. The set methods allow us to set values in the Actor. There are three methods that start with neither get nor set, which do various things. One of them is the act method.

Telling an actor to move is giving it an instruction. So the get methods are ruled out fairly quickly as they shouldn't change the actor at all, they should just provide us with information about the actor. One of the set methods is named setLocation which seems promising for making the person move. We are supposed to pass it an x and a y coordinate and it should move the actor to that location. Let's go ahead and try this with the simplest implementation possible. We will just call setLocation and give it constant values for x and y, perhaps something like the following.

```
public void act()
{
    setLocation(0,0);
}
```

Put this code into your Person class, compile all the classes, and click the Act button. What happens? You should see that both of the people move. They both move to the top left corner on top of each other so you only see one of them. Is this what you were expecting to have happen? Try changing the values so instead of 0,0 you have some other numbers. What happens then?

It shouldn't take long for you to realize that the x coordinate move the player from left to right and the y coordinate moves the player from top to bottom. This second fact might surprise you, after all, you always learned in math that the y axis goes from bottom to top. Computer graphics tend to be oriented the opposite way for historical reasons. The earliest “displays” were printing on paper and the paper always scrolled by so that the printing was done from top to bottom. The other thing that might surprise you is that both the x and y values start at 0, not at 1 the way you normally count. We will see that it is very common in Java to start counting things at zero. This also has historical roots.

So we got our people to move and we learned some things about the coordinates that are used in Greenfoot. However, we didn't get the people doing what we wanted. We wanted to make them walk to the side. Before we do that though, let's take one small step up in complexity for our program. The first version was boring because the people always jumped to a fixed location and it was always the same location. Why don't we make it so that the people jump to random locations? We can do this using code that is very similar to things we saw in the PSPGWombat example. We are going to call the method `Greenfoot.getRandomNumber()`. In addition to using this method, we should go ahead and introduce some variables in the act method. In this particular example we don't have to, but if we don't, our one line of code will get really long. Introducing variables, even when they aren't needed, can help make the code more readable. Make your act method look like this.

```
public void act()
{
    int x=Greenfoot.getRandomNumber(10);
    int y=Greenfoot.getRandomNumber(10);
    setLocation(x,y);
}
```

Compile this code and click the Act button. What happens? What happens if you click in again? Does this behavior make sense? It is what we wanted? Is this the ideal code for doing it?

The answer to that last question is, “No.” Why not? What is wrong with it? The problem is that it isn't general. I you recall the discussion in the last section we looked at the `randomLeaves` method in `WombatWorld` and discussed how the random number generation was done in a way that was general. The problem with this code is that we have hard coded in the number 10 because we know from the `City` class that the world is a 10x10 grid. What happens if someone changes the size of the world though? If they change the size of the world and don't change this act method in the `Person` class, our people will behave in ways that they shouldn't. How they misbehave depends on how the world size is changed, but whatever is done, we will have a bug³.

3 The term bug is commonly used in reference to errors that occur in programs. This term also has historical roots. The earliest computers were vast devices that filled large rooms and were often programmed by actually changing the wiring. One of these early computers kept producing erroneous output. Eventually it was discovered that a moth had flown into

We can make this code more general in the same way that randomLeaves was made more general. Instead of hard coding the numbers for the size of the world, we should ask for those values with a method call. In the WombatWorld class this was done by calling the getWidth() and getHieght() methods. This code tries to do that.

```
public void act()
{
    int x=Greenfoot.getRandomNumber(getWidth());    // Error: Actor doesn't
        have a getWidth() method, World does.
    int y=Greenfoot.getRandomNumber(getHeight());    // Error: Actor doesn't
        have a getHeight() method, World does.
    setLocation(x,y);
}
```

Unfortunately, this code doesn't work. That is because our act method is in Person which is an Actor. The getWidth() and getHeight() methods are in World. In order to call them, we have to have an object of type World. We don't want just any world either, we want the world that our person is part of. Fortunately, the Actor class has a method that will do this for us. If you go back and look in the API for Actor you will see that there is a method called getWorld() that returns an object of type World. Once we have that object, we can call getWidth() and getHeight() on it. This working code is shown here.

```
public void act()
{
    int x=Greenfoot.getRandomNumber(getWorld().getWidth());
    int y=Greenfoot.getRandomNumber(getWorld().getHeight());
    setLocation(x,y);
}
```

This is using the longer form of the method call where we have to explicitly provide an object to call the method on and use the dot notation to say that we are doing something with that object. Put this code in and compile the program then click the Act button. The behavior should be the same as what you saw before, but this code is now more general because it will work for any sized world.

Now we are ready to write the code we set out to add to our person to make it so that the person “walks” to the right. The question is, how do we do this? We know that locations with bigger x values are to the right. The problem is that we want the player to take one step each time act is called. So we want the x part of the location to be one bigger than what it already was. So to do this we need to get the current position in both x and y. The x we want to make bigger and the y we want to leave the same. To get the current location we need to consult the API again. It turns out there are methods in the Actor class called getX() and getY() that tell us the current location of the actor. So instead of setting our variables x and y equal to constants or random numbers, we can set them to the current position with a little math for offsets. This gives us the following code.

```
public void act()
{
    int x=getX()+1;
```

the machine and was causing a short that was altering the calculated results. Unlike the moth, modern bugs aren't the results of chance. They are errors put there by programmers who make mistakes. Despite this, the use of the term bug is still very common.

```
    int y=getY();
    setLocation(x,y);
}
```

You should put this into your Person class, compile it, and see what happens when you run it. Do the two people in the city walk to the right? They should. Make sure you understand why they do.

2.3. Types in Java

We have talked about types a few times to this point and we have seen a few used. The first type that we saw was the int type. This should be used when you need an integer value, that is a number, positive or negative, with no fractional part. We have also seen the boolean type which can either be true or false. We will see a lot more of the boolean type in the next chapter when we discuss the if statement and Boolean logic. If you look back at the rule for types you see that both int and boolean fit under the rule *PrimitiveType*. The *PrimitiveType* rule we are using for this chapter also allows two other types: double and char.

The double type is used to represent numbers, like int, only we use the double type when the number has a fractional part. The double type is also good for representing very big numbers for which you would normally use scientific notation. You might wonder why you should ever use the int type if the double type can store more general numbers and has a larger range. There are two answers to this. First, math done on the double type is inexact. If you go further in Computer Science you will likely learn the details of how the double type works and why it is inexact, but the basic idea is that it only keeps a certain number of digits of precision and after that it rounds. So if you are doing something that needs to have exact math, you should do it with an int type. Probably more important to you at this point is the second reason. There are parts of the Java language that require use use of ints. Also, if something wants a double and you pass it an int everything works fine because every possible int can be represented by a double. However, if you try to use a double in a place that wants an int, like the setLocation(int x,int y) method in Actor, you will get an error. You can test this out by trying to call setLocation(3.3,5.6) and then compiling. If you have a value that is a double and you really need to use it as an int you have to do something called a type cast. To do a type cast in Java, simply put the type you want in parentheses in front of the expression you want to cast. Here is a somewhat contrived example of that.

```
double x=6.7;
double y=4.3;
setLocation((int)x,(int)y);
```

Note that when you cast a double to an int it doesn't round, it truncates. Any fractional value is simply thrown away and you get the whole number part. So the location in this example will be set to (6, 4), not (7, 4) which is what you would expect if the values were rounded.

The char type is the only other primitive type that we will deal with in this book. As the name implies, the char type is used to represent character data. Actually, it represents a single character. We put characters in our Java code by putting the character we want inside of single quotes as is shown in this sample code.

```
char a='A';
```

The *Type* rule could also become a *ClassType* instead of a *PrimitiveType*. We have seen this option used as well. The randomLeaves method declared the variable leaf that was of type Leaf. Every

class defines a type in Java, whether it is one that we write, or one that someone has written. The `Leaf` type is an example of something we wrote. We will, over time, also make variables that are of some of the types defined in the Greenfoot API. We will also use a number of different classes from the standard Java libraries. One of those types is so fundamental that it is worth introducing here. It is the `String` type. The `String` type represents text and is quite general across programming languages. A `String` is made up of a bunch of characters. When you want to make a `String` in a Java program you simply put the text that you want inside of double quotes.

```
String s="This is a test.";
```

Note that the type `String` starts with a capital letter. This is because it is the name of a class that someone else wrote and they were following the standard Java naming scheme. You can type whatever you want inside the double quotes to make a string. The one exception to this is that you can't put a newline inside of them. If you want a string with a newline, you should put `\n` inside of the string.

The `String` type also allows something that no other class type does. You can use `+` with strings and the two strings will be concatenated. This was added to the language because it is such a common operation. We will use strings in the next chapter when we deal with keyboard input.

2.4. Modulo and Integer Division

The code we wrote earlier for our person had it walk to the right. Eventually it ran into the edge of the world and didn't do anything else. This is because Greenfoot doesn't allow actors to be placed outside of the bounds of the world. It would be nice to have the person be able to do something that isn't just random jumping around and doesn't end with it simply running into the edge of the world either. There are many options we could pick, but for now we will pick a fairly simple option. We want to have the person walk back and forth across the screen. It should take five steps to the right, then five steps to the left, and repeat that pattern.

In the next chapter we will learn how we could do something like this with conditional execution. That is likely how most programmers would choose to solve a problem like this. Since we don't know how to do that yet we will do it with basic math and some problem solving skills. Specifically, we are going to use the capabilities of integer division and the modulo operator.

Have you stopped to wonder what happens if you divide two integers? The operations of addition, subtraction, and multiplication are all closed on the integers. That is to say that if we add, subtract, or multiply integers, the result is always an integer. So if n and m are integers, we know that $n+m$, $n-m$, and $n*m$ will also be integers⁴. This isn't normally true for division. Most of the time in math n/m will not be an integer unless n happens to be a multiple of m . If you do division in Java with doubles you will get roughly the result that you expect. However, if you do it with integers the result you get will also be an integer. Specifically, it gives you the integer quotient.

To understand this, think way back to 4th or 5th grade before you knew decimal notation. Back then you would write the answer to $13/5$ as 2 remainder 3. Integer division in Java gives you the 2 part of this. There is another operator that you might have noticed in the *Expression* rule that gives us the remainder. It is the modulo operator and it is denoted by the `%` in Java. If you write `13%5` in Java you get 3, the remainder from your elementary school math.

Combining integer division and modulo can allow us to do a number of fun things. In this case, we want to use it to make our `Person` walk back and forth across the screen. We have already seen that

⁴ It should be noted at this point that the `int` type won't always behave the way you expect integers to. This is because the `int` type has finite bounds. If you start dealing with numbers much larger than 2 billion, the value will "overflow" and become negative. Similarly, if you have numbers smaller than -2 billion they can wrap around to be positive.

moving to the right can be achieved by doing something like `x=getX()+1`. Similarly, walking to the left can be accomplished with `x=getX()-1`. We can think of this as `x=getX()+dir` where `dir` is an `int` with a value of either `-1` or `1`. So all we have to do is come up with some math that allows us to have `dir` be `1` the first five times we act, `-1` the next five, and repeat.

In order to do this, we have to know what step we are on. We have to be counting our steps. This can be done in the same way that the Wombat counted the number of leaves it had eaten. We need to add a private member variable to the `Person`. We might call it `steps`.

```
private int steps=0;
```

Inside of the `act` method we need to have `steps` count up. We can accomplish this by putting a line like the following at the end of `act`.

```
steps=steps+1;
```

With these two lines, our person will now keep track of how many steps it has taken. Now all we have to do is come up with some math to go from number of steps to either `1` or `-1`. We wanted to do this based on whether we are in steps `0-4`, `5-9`, `10-14`, `15-19`, etc.

Look at those groupings. What is true about them, particularly in regards to division? The first group is `n` such that `n/5` is `0`. The next group would all be `1` when divided by five, etc. So `steps/5` tells us something about which direction we should be heading. When `steps/5` is even we want to go to the right. When it is odd, we want to go to the left. So how do we figure out if it is even or odd? That is a great use for the modulo operator. Given an integer `n`, `n%2` is `0` when `n` is even and `1` when `n` is odd. So the expression `(steps/5)%2` will be `0` when we want to go to the right and `1` when we want to go to the left.

The problem now is that we don't want `0` and `1`, we want `1` and `-1`. That can be fixed with a fairly simple set of operation. If `n` is `0` or `1` we can do `(n*2)-1` and we get `-1` or `1`. That is the opposite of what we want so we simply subtract instead of adding. This gives us final code for `act` that looks like the following.

```
public void act()
{
    int dir=((steps/5)%2)*2-1;
    int x=getX()-dir;
    int y=getY();
    setLocation(x,y);
    steps=steps+1;
}
```

Put this into your `Person` class, compile the code, and click on the `Run` button. If everything is typed in correctly your person should walk back and forth across the screen, just the way we wanted it to.

Chapter Rules

Class:

```
public class Name [extends ClassType] {
    [Constructors]
    [Methods]
```

```
    [MemberData]
}
```

Method:

```
public Type name([Type1 name1[,Type2 name2, ...]]) {
    [Statements]
}
```

Type:

PrimitiveType | *ClassType*

PrimitiveType:

boolean | char | int | double | void

The type void can only be used as the return type of a method, not for a variable declaration.

ClassType:

This is the name of a class. It can be one that you have written or one that is in an library. Class names typically start with a capital letter and have each new word capitalized. They follow the normal rules for Java names.

Name:

Java is fairly flexible about naming. Valid Java names have to start with a letter or underscore followed by any number of letters, underscores, or numbers.

Constructor:

```
public ClassName([Type1 name1[,Type2 name2, ...]]) {
    [Statements]
}
```

Statement:

VariableDeclaration | *Assignment* | *MethodCall* | *ReturnStatement* | *CodeBlock*

MethodCall:

```
[objectExpression .]methodName([expression1 [, expression2, ...]]); |
ClassName.methodName([expression1 [, expression2, ...]]);
```

The objectExpression is any expression whose type is an object type. If no objectExpression is given, the method is called on the current object. The methodName must be a valid name of a method on that type. The expression is often just the name of a variable, but it can itself be a method call that returns an object.

The second format is used for static methods. For these we don't need an object, we just put the name of the class in front of the dot. The Greenfoot class has only static methods in it so you call them with a format like Greenfoot.getRandomNumber(3);.

VariableDeclaration:

Type name[=expression];

This declares a variable of the given type with the specified name. You can think of variables as boxes that can hold a value. The box is labeled with the name we give it and it can only hold things of the proper type. What you can have for type is described after all the statement types are introduced.

A variable can only be used in a certain “scope”. The scope of a variable runs from the point at which it was declared down to the close curly brace that ends the code block it is declared in. Consider the following example method.

```
public void foo() {
    int a=5;
    {
        int b=8;
        // Point 1
    }
    // Point 2
}
```

At point 1 you could use either a or b. However, the scope of b is only inside the inner curly braces so you can't use b at point 2, only a.

Assignment:

name =expression;

This stores the value of an expression in the variable with the given name. The variable must be in scope. So it either has to be declared in the current method above this line or it should be a class level variable. In order for this to work, the types need to match.

Here are some example variable declarations and assignments that could be made from them.

```
int a;
int b;
double d;    // This type will be discussed later in the chapter.
String s;    // This type will be discussed later in the chapter.
```

```
a=6+7;
b=a-4;
d=2*3.14159;
s="This is a test.";
```

Expression:

An expression in Java is anything that has a type and value. Below is a list of options.

- Variable name – The name of a variable evaluates to the value currently stored in that variable.
- Numeric literal – Numbers are expressions. You can represent integer values or decimal fractions as you normally would. Very large or very small numbers are written using scientific

notation of the form 1.234e56 which is the number 1.234×10^{56} , a rather large number indeed.

- String literal – String literals can be any sequence of characters enclosed in double quotes. For example, “Hi mom.” is a string literal.
- null – This is the value for any object type when we don't have a valid object.
- Method call – A method call for any method that returns something other than void is an expression.
- Operator expression – Java supports different mathematical and logical operators. Most will join two different expressions. Some operate on one expression and one uses three expressions. This is a list that includes most of the ones we are likely to use.
 - A+B : Standard addition for numeric values. If either A or B is a String the plus sign will do string concatenation.
 - A-B : Subtraction for numeric values.
 - A*B : Multiplication for numeric values.
 - A/B : Division for numeric values.
 - A%B : Modulo operator that returns the remainder after division for integer values.
 - A=B : This is the assignment we saw above. It is also an expression which takes on the value that is being stored. In this case, A needs to be the name of a variable while B can be any expression of the right type.
 - new *ClassName*(args) : The new operator is used to make objects. This will make a new object of type *ClassName* using the constructor that is matched by the arguments passed in.

Member Data:

```
private Type name[=expression];
```

This defines a class level variable with the specified name and type. The initial value doesn't have to be specified.

Constructors and methods have a sequence of statements inside of curly braces. As a general rule, curly braces are used in Java to group lines of code. When a method is called in Java, the statements in that method are executed one at a time in order from the top of the method to the bottom. At this point we will only consider five types of statements in Java. So any place that you can have a statement, you can have one of these five. More statement types will be introduced later on.

Here is a simple little example of a class. This one doesn't do much, but it demonstrates the three things that you can put inside of a class.

```
public class SimpleClass {
    public SimpleClass(int v) {
        value=v;
    }
    public int getValue() {
        return value;
    }
    private int value;
}
```

ReturnStatement:

```
return [expression];
```

This statement stops the execution of a method and returns the value of the specified expression. If the method has a void return type, then there isn't an expression. Otherwise there must be an expression and the expression type must match the return type of the method or be convertible to it.

CodeBlock:

```
{ [statements] }
```

Any place a single statement is expected in Java, it can be replaced by multiple statements inside of curly braces. They don't all have to be on the same line as I have typed it here.

Exercises

1. Have it start with 20 houses instead of 5.
2. Spin the house.
3. Walk in a box shape.
4. Walk randomly.
5. Walk in a circle with sin and cos.
6. Make the person reproduce.
7. Have person remove itself at certain point.
8. Combine some of the above.

Chapter 3 - Conditionals

In the last chapter we saw how to program. We looked through the code for the Wombat and related classes. We even added our own functionality to the Person class to make it so that the person would move in various ways. We were somewhat limited in what we could do though. There was a fundamental reason for this. Given what we know, when we call a method, it executes one line after the other in order and always executes every line once. We were able to make some reasonable interesting behavior using different math functions, but at a certain level we were limited. We can add a lot more power to our programs if we have the ability to choose whether certain lines are executed or not. This power is called conditional execution and it is what we are going to learn about in this chapter.

3.1. The if Statement

The most basic conditional statement, and one that you will find in pretty much every programming language, is the if statement. It reads very much like plain English. If something is true, then do something. The rule for the if statement is shown below along with a modified definition for *Statement* that includes the if statement and the switch statement that will be introduced later in this chapter. There are two things to note about this rule. First, the else is optional. If you don't specify an else, when the condition is false execution will simply go to the statement after the if. Second, this is a common place where the *CodeBlock* rule is used. The *statement1* and *statement2* elements in this rule can be single statements, but it is very common to use a block of code in these situations. This means that you typically see the if statement followed by a curly brace, a set of indented statements, and a close curly brace. In fact, in this text we will always use code blocks with if statements, even when we only have a single statement in the curly braces and the braces aren't required.

Statement:

VariableDeclaration | *Assignment* | *MethodCall* | *ReturnStatement* | *CodeBlock* | *IfStatement* | *SwitchStatement*

IfStatement:

if(*booleanExpression*) *statement1* [else *statement2*]

This is the most basic conditional statement. If the boolean expression is true, *statement1* will execute. If it is false either nothing happens or *statement2* is executed if the else clause is used. A common usage might look like the following:

```
if(a==5) {  
    b=25;  
}
```

There is one other thing in this rule that might not make sense yet: the *booleanExpression* in the parentheses after the if statement. You know what the boolean type is and you know what an expression is. This notation indicates that it is an expression with the type boolean. What types of expressions have the boolean type? If an expression is a method call it will be boolean if the method it calls returns the boolean type. An example of this is the `canMove()` method that was in the Wombat class. The keywords `true` and `false` are also valid boolean expressions though they aren't very

interesting to use.

Java also has a set of operators that return boolean types. The first set of these that we want to talk about are numerical comparisons. These include the operators `==`, `<`, `>`, `<=`, `>=`, and `!=`. The first of these operators is a check for equality. Note that there are two equal signs. A single equal sign denotes assignment in Java, as we saw in the last chapter. If you want to compare two values to see if they are the same, you put two equal signs. In most uses, if you only put one you will get an error when you compile because the type won't be boolean and the if statement requires a boolean. The next four operators are fairly self-explanatory. They are true if the first value is less than, greater than, less than or equal to, or greater than or equal to the second value. The last operator, `!=`, stands for not equal and will be true if the two values are different to false if they are the same.

These few operators, along with the if statement give us the ability to do some things. To start with, we will revisit a problem that we have already solved before. It is the problem of making our person walk back and forth. In the last chapter we did this with integer division and modulo arithmetic. This is a perfectly valid way of doing it, but it is an approach that most people would find somewhat confusing. We can get the same result with conditional execution using the if statement.

Let's think through how we want to do this. As we saw in the last solution, the trick is that we have to keep track of our steps and change what direction we are going in at the proper time. Now that we have an if statement we can state what we want to do in English in such a way that it is almost directly convertible to Java. If we have done five steps in one direction, we should turn around and go the other direction.

As with our earlier solution, we want to keep a variable, `steps`, that counts steps for us. Unlike the earlier solution, this time the variable won't keep the total number of steps we have taken. That is because we don't really care how many steps total we have taken. We only care about the number of steps that have been taken since the last time we turned around. We will also move the variable `dir` to the class level because now we need to remember what direction we are going from one call of `act` to the next. The relevant code from `Person` for this solution is shown below.

```
private int steps=0;
private int dir=1;

public void act()
{
    if(steps==5) {
        dir=-1*dir;
        steps=0;
    }
    int x=getX()+dir;
    int y=getY();
    setLocation(x,y);
    steps=steps+1;
}
```

This code should be mostly self-explanatory. In the `act` method we first check if we have taken five steps. If we have, we do two things. We turn around and we reset our count of how many steps we have taken. The rest of the code is basically the same as what we saw before. The only minor trick here is multiplying `dir` by `-1` in order to turn around. You don't have to do it this way. The `act` method could be done like this instead:

```
public void act()
```

```

{
    if(steps==5) {
        if(dir==1) {
            dir=-1;
        } else {
            dir=1;
        }
        steps=0;
    }
    int x=getX()+dir;
    int y=getY();
    setLocation(x,y);
    steps=steps+1;
}

```

This solution employs an additional if statement with an else instead of doing the little multiplication math trick. Both are correct and they run exactly the same way in this code. The first one is a bit shorter, but if you find the second one easier to read and understand, you could certainly use it.

This version of the code also shows an interesting feature of the if statement: they can be nested. If you look back at the rule for the if statement, it simply says

Here are some additions to the expression. We won't repeat all the operators from the last chapter, we will only specify some additional operators that are relevant for our needs here. We just discussed the first six. The others will be discussed later in the chapter with the exception of the last one which appears in the next chapter.

Expression:

An expression in Java is anything that has a type and value. Below is a list of options.

- Operator expression – Java supports different mathematical and logical operators. Most will join two different expressions. Some operate on one expression and one uses three expressions. This is a list that includes most of the ones we are likely to use.
 - $A==B$: This returns true if A is the same as B. For object types use `.equals` instead.
 - $A<B$: Returns true if A is less than B. Works for numeric types.
 - $A>B$: Returns true if A is greater than B. Works for numeric types.
 - $A<=B$: Returns true if A is less or equal to than B. Works for numeric types.
 - $A>=B$: Returns true if A is greater than or equal to B. Works for numeric types.
 - $A!=B$: Returns true if A not equal to B.
 - $A \parallel B$: This boolean operator is true if A or B is true. It is inclusive or so it is true if both are true.
 - $A \&\& B$: The boolean operator returns true only if both A and B are true.
 - $!A$: This boolean operator is the not operator. It returns true if A is false and false if A is true.
 - $A?B:C$: This is the ternary operator. A should be a boolean. If A is true, the value will be B, otherwise the value will be C.
 - $A \text{ instanceof } C$: Here A is an expression with an object type and C is the name of a class. It tell you if A is of the type C.

3.2. Boolean Logic

In the last section we introduced some operators that return boolean values and allow us to build some simple boolean expressions. There are times when we need more complex boolean expressions that combine multiple other boolean expressions. There are operators that can do this for us. The use of these operators falls under the realm of boolean logic. It is very much like arithmetic in the two valued system of the boolean type.

We will look at three of these boolean operators: and, or, and not. To illustrate their usage, let's consider some simple examples. Say that we want to give our person an integer value for an age and have it behave differently depending on the value of age and some random chances. The behaviors will be something like combinations of things you made people do in the exercises for the last chapter. When their age is less than 24 or greater than 45 the person should move around randomly. When their age is between 25 and 30, inclusive, they should have a 20% chance of reproducing. For an age of 70 or greater the person should remove itself from the world. At all other ages the person just stays where it is.

We don't care so much about the code we would use to make the different actions happen. We care more about the boolean logic that goes into selecting which one we would do. Let's start with the first option. We are supposed to take that option if age is less than 24 or greater than 45. Putting some math notation into that we could write it as $\text{age} < 24 \text{ or } \text{age} > 45$. All we need is a way to express or. In Java we express or with two pipes, `||`. So our condition in Java will read as `age < 24 || age > 45`.

There are some things to note about the `||` operator. It is a short circuit inclusive or. The fact that it is inclusive is different from the normal usage of the word or in English, which is exclusive. When you say A or B in English it generally means that something is true if either A is true or B is true, but not if both are true. The inclusive or is true unless both operands are false. So it is true when both of the arguments are true.

The `||` operator is also a short circuit operator. In the case of or, this means that if the first operand is true, the second operand isn't even evaluated. That is because `true || B` is always true, regardless of what B is. This makes the operator more efficient and can be useful for other reasons as well.

For the second condition, we want an expression that is true if age is between 25 and 30, inclusive. In math class you might have written that as $25 \leq \text{age} \leq 30$. That won't work in Java. To understand why, you have to understand something about what the compiler sees. You might remember from back in algebra that for operators of the same precedence, you normally do them from left to right⁵. Java does this same thing. So if you type in `25 <= age <= 30` java sees $(25 \leq \text{age}) \leq 30$. To make it clear why this is a problem, pick a value for age, let's say 26. Because 26 is greater than 25, the expression in parentheses evaluates to true. So then Java sees `true <= 30`. That expression doesn't make sense. You are comparing a number to a boolean and the Java compiler will give you an error telling you that this is nonsense.

To see how we should do this, we can reword our English description of the expression a bit. Instead of saying that age is between 25 and 30, inclusive, we can say that age must be greater than or equal to 25 and be less than or equal to 30. Using the math operators that would be $25 \leq \text{age}$ and $\text{age} \leq 30$. We express the and in Java with the `&&` operator. So the condition we would want in our if statement is `25 <= age && age <= 30`.

The `&&` operator is a short circuit logical and. It operates on two boolean values and only returns true if both arguments are true. If either is false, or both are false, it will return false. The fact that it is short circuit means if the first argument is false, it doesn't bother evaluating the second argument. That is because `false && b` is going to be false no matter what b is. This property of the

⁵ Exponentiation is an exception to that rule as it should be done right to left, but since there is no exponentiation operator in Java that doesn't concern us here.

&& operator comes in very handy in certain situations. As we will see later on, there are some situations where an expression could cause an error. In that case you can put that expression second and have the first expression be true only if the second expression is safe to evaluate.

For this second condition we only wanted to take the particular action 20% of the time. There are two ways that we could do this. The first is to nest another if statement inside of the first one and have it check a random number. That would give us an if like the following:

```
if(25<=age && age<=30) {
    if(Greenfoot.getRandomNumber(100)<20) {
        // reproduce
    }
}
```

The alternative to this is to simply put the random number check in the outer if and combine it with the other checks using another &&. That would look like this:

```
if(25<=age && age<=30 && Greenfoot.getRandomNumber(100)<20) {
    // reproduce
}
```

This version is shorter, and assuming that there isn't something that we should do the other 80% of the time, this is probably the way that most people would write this code.

The last check is a simple check of age against the value 70. If we put all of this together we get the following code to represent our logic.

```
if(age<24 || age>45) {
    // walk randomly
} else if(25<=age && age<=30 && Greenfoot.getRandomNumber(100)<20) {
    // reproduce
} else if(age>=70) {
    // remove self from world
}
```

Because all of the conditions in this example are mutually exclusive we have used a common shortcut for nesting if statements in the else clause of other if statements. If the options weren't mutually exclusive, we would simply eliminate the else clauses all together and have each if start one a new line.

The expressions we used in this example all included either an || or an &&. None of them mixed the two. When you have expressions that use both && and ||, it often matters which one is evaluated first. You are familiar with the fact that in math, we always do multiplication before we do addition. So $2+3*5$ is 17, not 25. For boolean operators, && takes precedence over || and will happen first. If you ever have a hard time remembering that though, simply use parentheses to make it clear what operations are happening first.

There is another boolean operator in the Java language: not. The logical not operator is represented by an exclamation point, !, in the Java language. It is a simple operator that flips a true value to false and a false value to true. It should precede the boolean expression that it is operating on. The ! operator is higher precedence than && or ||. In fact, it is higher precedence than multiplication. As such, you will often want to put the expression you are negating with ! inside of parentheses.

3.3. Keyboard Input

Having conditional statements gives us the capability to add some interesting functionality to our Person class. In particular, we can make our person respond to keyboard input. If you look in the API at the Greenfoot class you will find that there the introductory information talks about keyboard input and there are two methods that deal with getting information from the keyboard: `String getKey()`, and `boolean isKeyDown(String)`. The first method will return a `String` to you that represents the most recently pressed key. The second method takes the key you are asking about and returns `true` or `false` to tell you whether it is currently being held down. This is the method that we will work with first.

Our goal is to make it so that the Person moves around responding to keyboard input. For the first cut, we'll just make it so that the movement is controlled using the arrow keys. If you read the API you saw that we represent all of the keys as strings. In the case of the arrow keys, they are "up", "down", "left", and "right". So we can pass these arguments into `isKeyDown` and Greenfoot will tell us whether or not that key is pressed. We can write four `if` statements where each one has a call to `isKeyDown` with a different argument. In the body of those `if` statements we could `setLocations`, but we will alter variables that we declare before the `if` statements. So assuming we use the variables `x` and `y` again, and initialize them to the current location of the person, then in the `if` statement that checks if the right key is pressed we will make `x` one larger. Doing this for all four directions gives us code like the following for the `act` method.

```
public void act() {
    int x=getX();
    int y=getY();
    if(Greenfoot.isKeyDown("left")) {
        x=x-1;
    }
    if(Greenfoot.isKeyDown("right")) {
        x=x+1;
    }
    if(Greenfoot.isKeyDown("up")) {
        y=y-1;
    }
    if(Greenfoot.isKeyDown("down")) {
        y=y+1;
    }
    setLocation(x,y);
}
```

You should enter this code into your Person class, compile it, and click the Run button. The people will just sit there unless you press an arrow key. What happens when you press an arrow key? Why? Is this what we wanted?

What you should have seen was that when you press an arrow key, both of the people in the world go moving in the same direction. When you release the arrow key, they both stop. Technically this is what we asked for. We wanted to make it so that the Person class was controlled by the arrow keys and we succeeded in that. Now that we have it though, we realize that this is a bit dumb. After all, what is the point of having both people do exactly the same thing, unless perhaps we were trying to make some type of line dancing program. For most programs it might be nice to have the capability to make it so that different keys control the different people. So maybe one person is controlled with the arrow keys and the other is controlled with 'w', 'a', 's', and 'd'. If we want to make this happen, the first thing that we should do is try to make our code more general.

In the `act` method shown above, we have hard coded in string literals, the stuff inside of double

quotes, to work with the arrow keys. What we should really have done is to make it so that the values we are checking are member variables. That way we could change the control keys by simply changing the values of those variables. Our first step is to make this generalization. We will change the code some without actually adding any functionality. This type of change, where you alter how a program does something without altering what it does, is called refactoring and it can be a very helpful thing to do when we want to add functionality. The purpose of refactoring before we actually enhance the functionality is that we can test the refactored code to make sure it still does what it did before, then after we have verified that it still works we can add the new functionality. This is a different style of breaking a large problem down. Instead of breaking a problem down into code, we are breaking down the way we code. Our refactored code might look something like this:

```
private String leftKey="left";
private String rightKey="right";
private String upKey="up";
private String downKey="down";

public void act() {
    int x=getX();
    int y=getY();
    if(Greenfoot.isKeyDown(leftKey)) {
        x=x-1;
    }
    if(Greenfoot.isKeyDown(rightKey)) {
        x=x+1;
    }
    if(Greenfoot.isKeyDown(upKey)) {
        y=y-1;
    }
    if(Greenfoot.isKeyDown(downKey)) {
        y=y+1;
    }
    setLocation(x,y);
}
```

If you compile and run this code you will see that it does exactly the same thing that the code had done before. So why did we do it? It sets us up nicely for the next step where we make it so different instances of Person can use different keys.

To accomplish this we need to add a constructor to the Person class. Actually, we will add two constructors. We'll see why in just a second. The first constructor we will add is one that sets the values of leftKey, rightKey, upKey, and downKey. Remember that constructors are like special methods that have the same name as the class and lack a return type. They are used specifically to set up values for new objects and are called when we use new. The constructor that we want to add looks like this:

```
public Person(String left,String right,String up,String down) {
    leftKey=left;
    rightKey=right;
    upKey=up;
    downKey=down;
}
```

If you put this code into your class and click on Compile All you will get an interesting error. Part of what is interesting is that the error isn't in Person. You will get an error in the City class. The error will read something like, "cannot find symbol – constructor Person()" and it will be on one of the lines calling addObject that adds a new person. What this error is telling us is that we need a constructor that doesn't take any arguments. You might wonder why it is that we didn't have one before and it worked, but it doesn't work now. A constructor with no arguments is often called the default constructor. As it happens, Java will add one for you by default if you don't have any other constructors. However, it only adds it by default if you don't have other constructors. So early Java provided a blank default constructor for us. Now that we have written a constructor, we have to make a constructor with no arguments explicitly. That means adding the following lines of code.

```
public Person() {  
}
```

With these in you should be able to compile. If you run the program it behaves exactly as it did before. Both the the people respond to the arrow keys. This is because both people are using the default constructor and keeping the initial values for the different strings.

In order to make it so that one of the people responds to different keys we need to edit code in the City class where the two different people are created and added to the world. There are two lines in this class that look like the following:

```
addObject(new Person(),2,3);  
addObject(new Person(),4,7);
```

We only need to alter one of these because it is fine if one of the people is controlled by the arrow keys, we just need the other one to be controlled by different keys. What we need to do is pass in four different strings in the parentheses after new Person so that those will be set as the controls for that player. After doing that, those two lines might look like the following.

```
addObject(new Person("a","d","w","s"),2,3);  
addObject(new Person(),4,7);
```

If you compile this code and run it you will find that now only one of the people moves when you use the arrow keys. The other responds to the letters we have told it to use. You could try adding yet a third person and giving it a different set of keys.

The Greenfoot class had another method for getting keyboard input called getKey. If you look at this method in the API you see that it returns a String. That String represents the most recent key pressed since the last time this method was called. If no key has been pressed it returns null. This null is something new to us. It is quite significant in Java and worth discussing. If you remember the rule for *Type* there were two options. The *PrimitiveType* we have covered fairly completely. The *ClassType* we talked about a bit and we said previously that String was one of these types. It turns out that there are other differences between these classes of types. We said earlier that you could picture a variable as something like a box and that box can hold a value of the specified type. This is exactly the picture you should have in your head for primitive types. For class types it isn't exactly correct. When we declare a variable of type String, for example, we don't get a box that holds strings, we get a box that can hold references to strings. Having a reference to an object is like having a person's address. It isn't quite the same as having them, but you know where to find them when you need them. Because these references store references and not actual values, they are often called reference types.

The fact that variables of class types store references allows a new possibility that can't happen with primitive types. A variable of a reference type can have an invalid reference. An integer variable is always an integer. It could be zero, but it can't be nothing. The same is true for other primitive types. Reference types, on the other hand, are allowed to store a reference that doesn't point to anything. The name of this invalid reference is null. The `getKey` method does what many methods will do that return a reference type. If there isn't a good value to return, it returns null. Just to see what it does, let's modify the `act` method in `Person` so that it uses `getKey` instead of `isKeyDown`. A first cut at this modification might look like the following.

```
public void act() {
    int x=getX();
    int y=getY();
    if(Greenfoot.getKey()==leftKey) {
        x=x-1;
    }
    if(Greenfoot.getKey()==rightKey) {
        x=x+1;
    }
    if(Greenfoot.getKey()==upKey) {
        y=y-1;
    }
    if(Greenfoot.getKey()==downKey) {
        y=y+1;
    }
    setLocation(x,y);
}
```

This is basically the same logic we had before, but now instead of checking the value of `isKeyDown` for each key, it is checking if the return of `getKey` is the same as the keys we are looking for.

It turns out that this rather naïve approach doesn't work. There are several reasons why. You should put in this code and run it to see what happens. The answer is nothing. No matter what key you press, the people don't move. This points us to the first thing that we have done wrong. The `==` operator is a great choice when working with ints. However, it typically doesn't do what you want for reference types. This is because `==` tells you if two references are the same, not whether the things that they reference are equal. If you think of a reference as storing the address that you can find something at, `==` checks to see if two things have the same address. That often isn't what you want to check for. Most of the time you want to know if the two things living at the referenced addresses are the same. To that you should call a method named `equals`. Every object in Java has this method, including strings. So the first change we need to make is to replace `==` with calls to the `equals` method. That makes the code look like the following.

```
public void act() {
    int x=getX();
    int y=getY();
    if(Greenfoot.getKey().equals(leftKey)) {
        x=x-1;
    }
    if(Greenfoot.getKey().equals(rightKey)) {
        x=x+1;
    }
    if(Greenfoot.getKey().equals(upKey)) {
```

```

        y=y-1;
    }
    if(Greenfoot.getKey().equals(downKey)) {
        y=y+1;
    }
    setLocation(x,y);
}

```

If you enter this code and run it now you get a different behavior. Instead of doing nothing, this version crashes. It brings up another window with a few lines of red text that begins by telling you that you have a `java.lang.NullPointerException`. This means that you tried to call a method on a null reference. Remember that null is an invalid reference. Because it doesn't reference anything, you can't do anything with it. When you try, you get this exception.

So the question is, where did we try to do something with null? If you look at the second line of red text, it will tell you. It says something like the following.

```

at Person.act(Person.java:33)

```

The part in parenthesis tells you that this happened in the file `Person.java`, where you were writing the `Person` class. In this case, it was on line 33. Your line number might be different. When you go look at that line you will see that it is the first if statement in the `act` method. So where is null coming from? It is being returned by `Greenfoot.getKey()`. After all, that is what the API told us it would do if no key had been pressed since the last time this was called. Our code crashed the first time it was called because no key had been pressed at all. We can fix that by adding checks to see if we have a null. However, if you read closely what `getKey` says it returns, and you run through the code that we wrote, you will see that we have another problem as well. The `getKey` method consumes key presses. If we call it multiple times in the `act` method, the first call might get a non-null value, but all the other calls will certainly get null. To fix that problem we need to add another variable. That variable will have the type `String` and will be set with the get of `getKey`. Then instead of calling `getKey` repeatedly, we use the variable. That gives us code that looks like this.

```

public void act() {
    int x=getX();
    int y=getY();
    String key=Greenfoot.getKey();
    if(key.equals(leftKey)) {
        x=x-1;
    }
    if(key.equals(rightKey)) {
        x=x+1;
    }
    if(key.equals(upKey)) {
        y=y-1;
    }
    if(key.equals(downKey)) {
        y=y+1;
    }
    setLocation(x,y);
}

```

This still hasn't fixed the problem with null when no key has been pressed though. To do that

we will use a little boolean logic. Consider the first if statement. The condition says that it is true if key equals the leftKey string for that person. We want to add the additional requirement that key not be null. So we want to move if key is not null AND key equals leftKey. The following code shows that implemented in all the if statements.

```
public void act() {
    int x=getX();
    int y=getY();
    String key=Greenfoot.getKey();
    if(key!=null && key.equals(leftKey)) {
        x=x-1;
    }
    if(key!=null && key.equals(rightKey)) {
        x=x+1;
    }
    if(key!=null && key.equals(upKey)) {
        y=y-1;
    }
    if(key!=null && key.equals(downKey)) {
        y=y+1;
    }
    setLocation(x,y);
}
```

Note that the order of the two checks on either side of the && matters. This is where the fact that && is a short circuit operation comes into play. If key is null the first condition is false and the second condition is not evaluated. That is critical, because if key is null, the second condition is an error and we don't want it to happen. Feel free to put the key!=null checks on the right side of the && and see what happens.

This code should compile and run. However, it still doesn't work as well as what we had before using isKeyDown. When you run this code you will find that only one of the people moves. Why is that? It is because one person always has its act method called before the other. The one consumes the key press and the second person always has a key value of null. At this point you can probably see why we used isKeyDown to start with. It and getKey have very different behaviors and should be used in very different types of situations. At this point you can go ahead and revert to the code using isKeyDown as we will be coming back to that in a few sections.

3.3. The switch Statement

The if statement isn't the only conditional statement in Java. There is also the much less used switch statement. Technically, the if statement can perform any conditional execution that you want. There is no need for alternatives. However, the syntax of the if statement can be a bit verbose for certain types of situations where there are lots of possible cases, especially if those cases are all checking against the value of a single expression. The switch statement is used for exactly those situations.

The idea of the switch statement is that you need to perform one of many possible sets of instructions and which one you execute depends on the value of an integer expression⁶. The rule for the switch statement is shown below. The way switch works is that it evaluates the expression in

⁶ Technically you can do this on the char type as well because characters are just numbers to the computer. What makes them characters is how they are treated.

parentheses, then runs through the different cases until it find the one that matches. If none match, it will do the default case.

SwitchStatement:

```
switch(intExpression) {
case intValue1:
    [statements]
    break;
case intValue2:
    [statements]
    break;
case intValue3:
    [statements]
    break;
...
default:
}
```

The switch statement allows you to branch to one of many options based on the value of an integer expression. The values don't have to start at 1 as I have shown here. They can be whatever you want. Break statements are also optional, but they are almost always needed. Break statements can also be used in loops to terminate the loop, but I won't ever do that.

One thing to note about the expression in the switch statement is that it can't be a String. In fact, it can't be any reference type⁷ or a double. For this reason, the switch statement wouldn't have been useful in the keyboard input examples we have looked at. An example that could occur in this book would be like the following.

```
public void moveDir() {
    int x=getX();
    int y=getY();
    switch(dir) {
    case 0:
        x=x+1;
        break;
    case 1:
        y=y+1;
        break;
    case 2:
        x=x-1;
        break;
    case 3:
        y=y-1;
        break;
    default:
    }
    setLocation(x,y);
}
```

⁷ Beginning with Java 5 there is a new reference type called an enum that we will not discuss in this book. They can be used with switch statements.

```
}
```

The idea here is that our class has a member called `dir` that is an `int`. Instead of having four `if` statements to check the value of `dir`, we can use a single `switch` statement with four cases in it. The `Wombat` class in the previous chapter had this type of functionality in it, but it was not implemented with `switch` statements. For the most part the `switch` statement is mentioned here for completeness just in case you ever run across code that uses it. However, most uses of the `switch` statement can be done in better ways in Java so it will not be used extensively in this book.

3.4. Ternary Operator

One of the limitations of the `if` statement is that it is a statement. This is great when you want to do something depending on a condition. However, there are times when you would like the value of an expression to depend on some condition. The `if` statement can't do this. It isn't an expression and can't be used in places where you want a value. To see what the problem is here, consider the following code.

```
int b;  
if(a<5) {  
    b=34;  
} else {  
    b=42;  
}
```

All this code does is assign a different value to `b` based on the value of `a`. If the `if` statement could be used as an expression we could type in something more like this:

```
b=if(a<5) 34 else 42;
```

This isn't valid Java and if you try to enter it into a program you will get errors when you try to compile. It is worth noting that there are other languages where this would work. Because it is so useful, there is a way of doing this in Java. It just uses more symbols instead of English words.

The way we do this is with the ternary operator. It is an operator that takes three arguments. The syntax was given in the rule for expressions earlier. The first operand is the condition, just like what you put in the `if` statement. After that is a question mark which is followed by an expression that should be the value if the condition is true. That is followed by a colon and then the expression if the value is false. So our example above could be typed into Java with the following syntax.

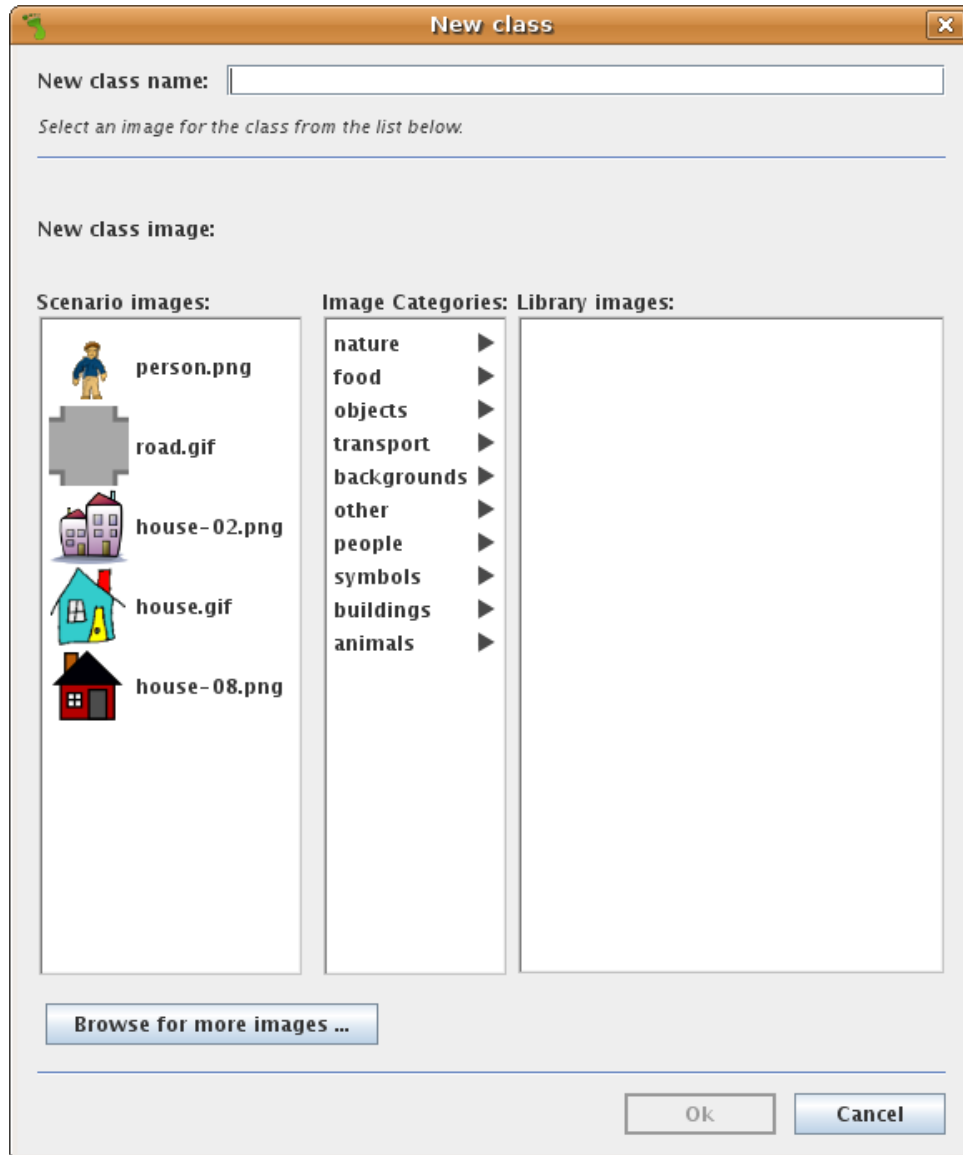
```
b=(a<5)?34:42;
```

The problem with the ternary operator is that it can be hard to read. People commonly put in extra parentheses to make things more clear, but even then their meaning isn't always straightforward. For this reason you won't see ternary operators used all that much. They are far less common than the `if` statement, but when they are useful they can greatly reduce the length of code and so they do occur occasionally.

3.5. Our First Game (An Example)

Now it is time to extend our `CityScape` code so that it really has some functionality. Instead of just having the people move around, we'd like to make this into a game. The two people will compete

against one another in this game. The idea is to have them race to pick up items. Whoever collects the most by the time they are all gone wins. Picking up the houses doesn't make much sense. We will use them in a different way later. That means that we need to introduce a new type of actor into our world, something that it makes sense for the people to pick up. To do this, right click on the Actor box that is above House and Person. This will bring up a menu with two options. The second one is "New Subclass ...". When you select it, the following dialog box will come up.



We will call our new class "Food" and under food select the apple1 image. Once you have done this and clicked "Ok", a new class should appear under Actor with the name Food.

Like the Leaf in the PSPGWombats scenario, the Food class isn't going to do much of anything itself. The Person class will interact with it, much as the Wombat class interacted with the Leaf class. The class we want to edit is the Person class. We want to keep the code that we had earlier to do movement with `isKeyDown`, but we want to add onto that code that checks for whether we are standing on the same square as a food. If so, we want to pick it up and add it to our count.

All of the code that we want to write could be put into the act method. Making really big methods though is generally considered bad form. What is worse, it often makes the code harder to understand. What we would like to do is decompose this problem into some major pieces and then write each of those separately. The act method can call each of these pieces as needed.

To do this we have make new methods. We saw the rule for methods in the last chapter when looking at code, but we haven't actually had to write new methods ourselves. We have only edited methods that were already provided for us. Proper problem decomposition requires us to create new methods so that is what we will do here. We will do this in a top down manner. First we will edit the code in act and basically structure it so that calls other methods that don't yet exist. The idea being that once those other methods are written, the act method will work. For what we have described so far, the act method might look like the following.

```
public void act() {
    if(canPickUpFood()) {
        getFood();
    }
    keyboardMove();
}
```

Now all we have to do is write the three methods that this calls: canPickUpFood, getFood, and keyboardMove. We'll start with the last one because we have basically already written it. It should look something like this.

```
public void keyboardMove() {
    int x=getX();
    int y=getY();
    if(Greenfoot.isKeyDown(leftKey)) {
        x=x-1;
    }
    if(Greenfoot.isKeyDown(rightKey)) {
        x=x+1;
    }
    if(Greenfoot.isKeyDown(upKey)) {
        y=y-1;
    }
    if(Greenfoot.isKeyDown(downKey)) {
        y=y+1;
    }
    setLocation(x,y);
}
```

Notice that this is exactly what our act method looked like earlier. The only difference is that instead of having the method called act, it is now called keyboardMove. It is public since we are currently making all of our methods public. It has a return type of void because we don't need any information back from it. If you look at the act method, we didn't do anything with a return value so it makes sense that there shouldn't be one.

The canPickUpFood method does need to return something. This is because it is used in an if statement. Remember that the condition of an if statement has to be a boolean. So the return type of canPickUpFood needs to be boolean as well. Inside of the canPickUpFood method we need to answer the question of whether or not the person is standing on the same square as a food object. If you look

at the API for Actor you will see that there is a method with the name, `getOneIntersectingObject`. This method takes one argument of type `java.lang.Class`. This might sound scary, but it really isn't. We can easily get a Class object by putting `.class` after the name of the java class we want the Class of. If there isn't an intersecting object of that type, this method should return null. So if we pass in `Food.class`, it will give us a Food object if the person is on one. If the person isn't on one, it will give us null. So the `canPickUpFood` method, given the rules we are currently using might look like the following.

```
public boolean canPickUpFood() {
    return getOneIntersectingObject(Food.class) != null;
}
```

You might be tempted to write this method using an if statement like this.

```
public boolean canPickUpFood() {
    if(getOneIntersectingObject(Food.class) != null) {
        return true;
    } else {
        return false;
    }
}
```

These two versions do the same thing, but the first one is a lot shorter and potentially faster as well. As a general rule, when you have an if where all it does is return true and false in the conditions, simply return the condition instead. After all, the condition is a boolean and it has the value that we want.

Last we will write the `getFood` method. This method should take the food out of the world and increment a counter that keeps track of how many food objects the person has collected. Removing an object from the world can be done with the `removeObject` method in the World class. Since Person is an Actor, not a World, we will have to use the `getWorld` method of Actor as we did before when we wanted to call `addObject`. We also need to pass it the food object that we want to remove. We can get this object using `getOneIntersectingObject`, just like we did in the `canPickUpFood` method. The new code might look like the following.

```
private int foodCollected=0;

public void getFood() {
    foodCollected=foodCollected+1;
    getWorld().removeObject(getOneIntersectingObject(Food.class));
}
```

Go ahead and put in these different methods and compile the code. You can manually add a few pieces of food to allow you to test and make sure that the code is working. Once you have added the food, click Run and move the people on to some of the food. It should disappear. Once you have gotten rid of all the food, click Pause and right click on one of the people and select the inspect option from the menu. This will bring up information on the person including the value of `foodCollected`. Using this you can see which of the people had eaten more of the food. You can go into the City class and add a few lines to the constructor there so that every time you compile or reset, some food will be added to the world, perhaps at random locations. We will talk about a better way to do that in the next chapter.

Once you put in that code you actually have something of a game. It is a very simple game

though, and we haven't used the houses at all. One way we could fix both of these problems is to make it so that a person can only carry a certain number of food at a given time and has to go to a house to drop off food so that more can be collected. The number of food a person can carry should be fairly small. It is interesting to note that if it were always 1, things could be coded a bit differently. We'll point out how as we edit the code.

To make this modification, we should start with the act method again. We start here because we want to see at a high level what pieces the problem can be broken into. There are a lot of different options for the modified act method. This is one of them.

```
public void act() {
    if(canPickUpFood()) {
        getFood();
    }
    if(canDropFood()) {
        dropFood();
    }
    keyboardMove();
}
```

The approach here is to simply add a check for whether we are able to drop the food and if so, to drop it. To implement these methods we will have to add some member data that stores how much food the person is currently carrying. This value will also be used in modified versions of canPickUpFood and getFood. Note that if the player is allowed to hold no more than one food, this member data can be a boolean telling us whether or not the person is holding food. If it can be larger than one though, an int is needed. This code shows that member data, the two new methods, and the modified versions of the methods we had before.

```
private int foodCarried=0;

public boolean canPickUpFood() {
    return foodCarried<2 && getOneIntersectingObject(Food.class)!=null;
}

public void getFood() {
    foodCollected=foodCollected+1;
    foodCarried=foodCarried+1;
    getWorld().removeObject(getOneIntersectingObject(Food.class));
}

public boolean canDropFood() {
    return getOneIntersectingObject(House.class)!=null;
}

public void dropFood() {
    foodCarried=0;
}
```

This version of the code allows the player to carry two pieces of food at any one time. Let's run through each piece to see what is happening. First is the declaration of a new value that the Person remembers called foodCarried. After that, the canPickUpFood method has been altered so that the person not only has to be on food, but can't be carrying more than one food. The constant 2 here is

what sets how many food objects a person can pick up before having to run to a house. Having random constants like this in your code is generally frowned upon. Because this one only appears in one place in the code, and because doing the other approaches properly requires introducing some new keywords that we won't bring in until later, we will leave it as just the number 2 for now.

The `getFood` method has one new line in it that increments the number of food items being carried. The two new methods are both quite simple. The first one returns a boolean to let us know if we are standing on a house and the second one sets that amount of food being carried to zero. Put in this code and run it. Make sure it works for you and that you understand what it is doing and how it works.

3.6. The Mandelbrot Set (An Example)

Before we leave this chapter on conditionals, there is one more little example that we will do. In many ways this example is much simpler than our last one. The reason for doing it is that it is a very different type of example. We are going to have Greenfoot actors form the Mandelbrot set for us. The Mandelbrot set is probably the most famous of all fractals. It is a set of points in the complex number plane and the fractal is named after Benoît Mandelbrot.

Any time you solve a problem, whether it is on the computer or elsewhere, the first step is to figure out exactly what problem you are solving. In computer science this step is commonly called analysis. In this case, we want to make a program that will show us a graphical representation of the Mandelbrot set. This set is defined to be all points in the complex plane such that the iteration of a particular equation stays bound. The equation is $z_{i+1}=z_i^2+c$, where c is the point we are considering in the plane and $z_0=0$.

Note that the z and c values are complex numbers. For this reason it is probably worth doing a quick math refresher. Most of the time we work with real numbers. These are the numbers on the standard number line. You can do almost any type of mathematical operation you want on real numbers. There are exceptions though. The main exception related to complex numbers is that the square root operation is not closed over the reals. That is to say that there are some real numbers that when you take the square root, you don't get back a real number. In particular, the square root of any negative number can't be a real number, because any real number squared is a positive real⁸. To fix this, we have imaginary numbers. The number i is defined to be the square root of -1 . So the square root of -4 is $2i$ and so forth. The imaginary numbers can be laid out on a number line, just like the real numbers.

Complex numbers are simply the sum of a real number and an imaginary number. For example, $3+5i$ is a complex number. Complex numbers are often represented on a plane with the x-axis being the real number line and the y-axis being the imaginary number line. So our example of $3+5i$ would be located three units to the left and five units above the origin.

When we add complex numbers we simply add the real and imaginary parts independently. So $(3+5i)+(4-3i)=7+2i$. More generally, $(a+bi)+(c+di)=(a+c)+(b+d)i$. When we multiply you just use the standard FOIL rule that you learned in algebra. So $(a+bi)*(c+di)=ac+adi+bci-bd=(ac-bd)+(ad+bc)i$. The formula for the Mandelbrot set has us square one value, then add it to another so we will be doing both multiplication and addition.

To keep things simple, we will say that a point is not in the set if the magnitude of z_i ever becomes larger than 2. We will also only go for a very simple graphical representation. If you search on the internet for the Mandelbrot set you will see lots of pretty pictures. We are only going for a small

⁸ This problem is more general than just square roots. Most generally, it can be a problem for any exponent that isn't a whole number. Recall that a square root is the same as raising the value to the $\frac{1}{2}$ power.

black and white image.

Now that we know what problem we are solving, we can think about how to solve it. This stage is problem solving is called design in computer science. Like analysis, this step doesn't involve writing any code. During design you should just think about how you can solve the problem. You are thinking about what you would put into the code, but you aren't really writing any code. The real question is, what do we know how to do in Greenfoot? We know how to put actors in the world. We know how to move them. We know how to remove them. We have the ability to do some math and do conditional logic. So what if we made an Actor type called Point and made the image for Point just a single dot on the screen. Then we could fill the world with a grid of these points. Each point could know its c and z values. The act method for the point could do the Mandelbrot iteration formula and check the magnitude of the point. If it gets too big, it can remove itself from the world.

So that is our design. Load in the PSPGMandelbrot scenario and we will implement it. Some parts have already been implemented. The ComplexPlane class deals with setting up the grid of Points. For now we will just assume that code works. In the next chapter we will look more closely at it. What we want to do is implement the Point class. We know from our design that the point needs to remember the c and z values for it. Java doesn't have a type for complex numbers. Instead, we will store the real and imaginary parts independently in variables of type double. We use double because these aren't going to be whole numbers and the double type is as close as we get to real numbers. So our member data declarations might look something like this.

```
private double cr;  
private double ci;  
private double zr;  
private double zi;
```

The act method for Point needs to do the Mandelbrot iteration on the z values, then check to see if the magnitude of the z value is greater than 2. If it is greater than 2 then we want to remove this point from the world. Here is an act method that does exactly that.

```
public void act() {  
    double tmp=zr;  
    zr=zr*zr-zi*zi+cr;  
    zi=2*tmp*zi+ci;  
    double mag=zr*zr+zi*zi;  
    if(mag>4) {  
        getWorld().removeObject(this);  
        return;  
    }  
}
```

The first three lines do the math for the Mandelbrot. We have to introduce a temporary variable, called tmp, to store the old value of zr that we will need to calculate zi after we have changed the value of zr. After those lines is a line that calculates the square of the magnitude and checks to see if it is greater than 4. If it is, it removes this point from the world.

The line that removes the current point from the world includes something new that we haven't seen before. It is the keyword this. The this keyword acts like a variable that references the object that the method you are in was called on. It is implicitly used in a lot of the code we have already written. The act method could actually be rewritten with this explicitly put in before every reference to member data or methods. That would make it look like the following.

```

public void act() {
    double tmp=this.zr;
    this.zr=this.zr*this.zr-this.zi*this.zi+this.cr;
    this.zi=2*tmp*this.zi+this.ci;
    double mag=this.zr*this.zr+this.zi*this.zi;
    if(mag>4) {
        this.getWorld().removeObject(this);
        return;
    }
}

```

Some developers like this style as it makes it explicitly clear when you are referring to something at the class level. At you can see here though, it can make things far more verbose which is why we have not adopted that style.

If you put this code inside of the Point class and try to compile you will actually get an error in the ComplexPlane class. It is on the line that tries to make a new Point object. That is because we haven't included a constructor that allows the ComplexPlane class to set the values of cr and ci. That constructor might look like the following.

```

public Point(double real_0, double img_0) {
    cr=real_0;
    ci=img_0;
}

```

Once you have included that you should be able to compile the code. Click on the Act button a few times and see what happens.

The first thing that you are likely to notice is that this runs really slowly, even though the image that we are generating is rather small. That is because the solution we designed isn't exactly the best way to solve this problem. It has a lot of overhead due to the fact that every little point in the picture is represented by a Point object. In fact, that is why the image isn't any bigger. If we made it much bigger, you would likely get errors saying that Java had run out of memory. Once we have learned more different things that way can do in Greenfoot we can come back to this problem and write more efficient solutions.

Chapter Rules

Statement:

VariableDeclaration | *Assignment* | *MethodCall* | *ReturnStatement* | *CodeBlock* | *IfStatement* | *SwitchStatement*

IfStatement:

if(booleanExpression) statement1 [*else statement2*]

This is the most basic conditional statement. If the boolean expression is true, statement1 will execute. If it is false either nothing happens or statement2 is executed if the else clause is used. A common usage might look like the following:

```

if(a==5) {

```

```
        b=25;
    }
```

SwitchStatement:

```
switch(intExpression) {
case intValue1:
    [statements]
    break;
case intValue2:
    [statements]
    break;
case intValue3:
    [statements]
    break;
...
default:
}
```

The switch statement allows you to branch to one of many options based on the value of an integer expression. The values don't have to start at 1 as I have shown here. They can be whatever you want. Break statements are also optional, but they are almost always needed. Break statements can also be used in loops to terminate the loop, but I won't ever do that.

Expression:

An expression in Java is anything that has a type and value. Below is a list of options.

- Operator expression – Java supports different mathematical and logical operators. Most will join two different expressions. Some operate on one expression and one uses three expressions. This is a list that includes most of the ones we are likely to use.
 - $A==B$: This returns true if A is the same as B. For object types use `.equals` instead.
 - $A<B$: Returns true if A is less than B. Works for numeric types.
 - $A>B$: Returns true if A is greater than B. Works for numeric types.
 - $A<=B$: Returns true if A is less or equal to than B. Works for numeric types.
 - $A>=B$: Returns true if A is greater than or equal to B. Works for numeric types.
 - $A!=B$: Returns true if A not equal to B.
 - $A || B$: This boolean operator is true if A or B is true. It is inclusive or so it is true if both are true.
 - $A \&\& B$: The boolean operator returns true only if both A and B are true.
 - $!A$: This boolean operator is the not operator. It returns true if A is false and false if A is true.
 - $A?B:C$: This is the ternary operator. A should be a boolean. If A is true, the value will be B, otherwise the value will be C.
 - $A \text{ instanceof } C$: Here A is an expression with an object type and C is the name of a class. It tell you if A is of the type C.

Exercises

1. Walk across world and wrap
2. Enter houses
3. Have people walk randomly. They make buildings if they come together and there isn't one there.
4. Have people walk randomly, but make it so that they never go onto the same square as a house or another person.

4 - Loops and Greenfoot Lists

There have been a number of different times already where we have come across situations where we need to repeat a certain action multiple times. For example, in the CityScape scenario we wanted to add multiple buildings. The way we accomplished this was to take the line that added one and cut and paste it multiple times. There are many reasons why this is a bad solution. In this chapter we will look at better solutions. We will also learn how to work with the List type which are used in a number of places by Greenfoot.

4.1. Need for Loops

Before we look at how we should do repetitive actions in Java let's go through the reasons why copy and pasting code is so bad. There are two reasons that are easy for the novice programmer to understand. First, if you have to do something a lot the copy and paste method isn't very efficient. In the case of the Mandelbrot program, it added 40,000 Point objects to the world. Each one of those had slightly different values. Doing that by hand would be insane. That's the type of thing computers are good at and humans aren't. Still, the copy and paste method could have been used as long as you were really patient. What about the situation where you don't know how many times you want to do something? Consider the following method signature for the City with our game.

```
public void addFruit(int howMany)
```

This would be for a method that is supposed to add a specified number of fruit. If we call it and pass in 5 it should add five fruit. If we call it and pass in 10 it should add ten fruit. How would you do this with cut and paste? It can be done if you put every addObject call inside of an if statement and have each one check for a larger input, but even that has limits because you probably won't cut and paste that more than a few times and in theory someone could call it with an input parameter of 1000.

So the cut and paste approach lacks flexibility and is extremely tedious. There is another fundamental problem though that is less obvious to the novice programmer. Every line of code you write is a line of code that could have an error or that might need to be modified later. This is part of why we like solutions to things that are both general and short. Imagine you use the cut and paste method and made 100 copies of some line of code. After doing this, you find that your original line had an error in it. Now instead of having one error, you have 100. Even if you were really careful and made sure the first line was correct, you might find later on that you need to do something slightly different. Instead of having one line to change, now you have 100.

It is something of a general rule of programming that if you find yourself copying and pasting code more than once or twice, you should probably sit back and think if there is a better way of doing whatever you are doing. This is because having multiple copies of code that do nearly the same thing almost always leads to problems down the line. For now we will use loops to help us avoid this. That only gets the simple cases though. Later in this book we will talk a little bit about a concept called abstraction and see how it can help get us around a bunch of other cases.

4.2. The for Loop

To start with, let's consider the simplest case, a situation where we want to do something a specified number of times. The method, addFruit, mentioned above is a great example to begin with. We know what statement we should execute in order to add a single Food object to a random location in the world. What we need is some construct that will allow us to have that statement execute a

specific number of times. The most common way to do that in Java is with a for loop. The for loop is actually very general, and allows us to do many different things. However, it is most commonly used to count. In this case, that is exactly what we want to do. The first cut at writing the addFruit method might look like this.

```
public void addFruit(int howMany) {
    for(int i=0; i<howMany; i++) {
        addObject(new Food(),Greenfoot.getRandomNumber(getWidth()),
            Greenfoot.getRandomNumber(getHeight()));
    }
}
```

Like the if statement, a for statement starts with the proper keyword and is then followed by stuff in parentheses. Unlike the if statement, the for statement has to have three different things in the parentheses, and they need to be separated by semicolons. When you are writing a counting loop it will often look like this. The first thing inside the for loop will declare a loop variable and give it an initial value. Then there is a semicolon followed by a boolean condition that tells it how far to count. That is followed by another semicolon and a statement that increments the counting variable. Note that `i++` means exactly the same thing as `i=i+1`. The `++` operator is called an increment operator. There is also a decrement operator, `--`. These operators are commonly used in Java as shortcuts. We will use them in for loops and any place else that we need to increment or decrement values through the rest of this book. This isn't required. If it makes more sense to you, you can certainly have the last part of the for loop written with `i=i+1`.

!!! Consider and advanced aside here that describes the difference between pre and post increment/decrement.

So what really happens when the addFruit method is called? The way the for loop works, when it is reached, the initialization statement is executed. In this case that means that the variable `i` is declared and set to 0. Then the condition is checked. Note that the condition is checked before the body of the loop ever executes. The for loop is a pre-check loop. In this type of loop, it is possible for the body to never execute. For example, if you call this method with an argument of -1, no food objects will be added to the world because the condition, `0<-1` is false.

If the condition is true, the body of the loop is executed. After the body of the loop, the iteration statement is executed and the condition is checked again. This will repeat as long as the condition is true. Once the condition is false, the loop stops and the program continues with the line below the loop. In this case, that means the method is done and it exits.

If you call this method and pass in 5 for `howMany`, note that `i` will start at 0, then go to 1, 2, 3, and 4. After the body is executed with `i==4`, `i` is incremented to 5 and the loop stops. It is common practice to have loops in Java start counting at zero, just like we saw that the `x` and `y` coordinates of the world start at 0. We will see later in this chapter and in the next chapter that other things in Java begin at zero as well. It would be equally valid to rewrite the for loop in this example like this.

```
for(int i=1; i<=howMany; i++) {
```

This version would count from 1 to 5 instead of 0 to 4. There are some situations where you might want to do this, but most of the time we will stick with loops that begin counting at 0. After all, they require one fewer keystrokes because they don't need the equal sign in the condition.

Now let's look at a slightly more complex example. Open up the ComplexPlane class that was in the Mandelbrot scenario. The constructor for that class added all the Point objects into the world. In that case, there were a fixed number of points so technically cut and paste could have been used. In practice though, the fact that there were 40,000 points being added and that each call needed to use slightly different values make it impractical. Here is the part of that method that contains the for loops.

```
double realMin=-1.5;
double realMax=0.5;
double imgMin=-1;
double imgMax=1;
for(int i=0; i<getWidth(); i++) {
    double real=realMin+i*(realMax-realMin)/getWidth();
    for(int j=0; j<getHeight(); j++) {
        double img=imgMin+j*(imgMax-imgMin)/getHeight();
        Point p=new Point(real,img);
        addObject(p,i,j);
        p.setImage(pointImage);
    }
}
```

You should automatically see that this code has two for loops in it and that they are nested one inside of the other. Just like with if statements, the body of a for statement can contain any types of statements. So the body of the for and contain others fors. It could also contain ifs, switches, or any other valid Java statement. The generality of statements provides Java, and most other programming languages, with a lot of flexibility and gives you as a programmer much greater expressivity.

While it is simple to say that because a loop is a statement, it can be nested inside of other loops, it is worth taking a second to loop at what really happens when loops are nested. We said above that this code adds 40,000 points to the world. How does that happen? Both width and height are only 200. So neither loop ever counts above 200. To understand what is going on, we need to trace through the code. We start with the declarations of four variables, then we hit the outer for loop and we get $i=0$. Since $0 < 200$ the body of the loop executes and another variables is declared and initialized. Then we hit the inner loop and it does $j=0$. Since $0 < 200$ the body of this loop executes. It executes 200 times in fact as j goes from 0 to 1 to 2 and on up to 199. It stops when j gets to 200. At that point the body of the outer loop is also done so it does $i++$ and we now have $i=1$. It then goes back through, sets $j=0$ and counts up to 200 again. So when we nest loops, every time that the outer loop goes through one iteration, the inner loop will execute all the way through. This tends to make it so that the body of the inner loop happens a multiplicative number of times. In this case, each loop counts to 200 and $200 * 200 = 40,000$. One last point about nesting loops is that the inner loops generally need to have different loop variable names. If you tried to change the j to an i , this code would not compile. It would tell you that you had declared i twice, which indeed you would have.

The purpose of this code was the lay out a grid of points both on the Greenfoot world and in the complex number plane. The i and j counting variables run through the grid of the Greenfoot world. The double variables $real$ and img are given values that correspond to the x and y axis values in the complex plane inside of the range that is specified by the four variables at the top. You can actually change the values of those variables to see different parts of the Mandelbrot set.

The math that goes into calculating $real$ and img is also worth taking a few moments to look at. It does a linear scaling. In both cases, the math looks something like this: $v = \min + i * (\max - \min) / \text{count}$. When $i=0$ in this formula, we have $v = \min + 0 * (\max - \min) / \text{count} = \min$. When $i = \text{count}$ we have $v = \min + \text{count} * (\max - \min) / \text{count} = \min + \max - \min = \max$. As the value of v goes up from 0 to count , the

value of *v* goes from *min* to *max*. This is a fairly common formula to see as linear interpolations like this are fairly common in mathematical settings.

In its full general form, the `for` statement follows the rule below. The first and last statements can actually be any single statement that you want. The middle is any expression that has the type `boolean`. You are also allowed to leave any of these elements blank. Leaving the middle one blank isn't generally a good idea because it will always evaluate to `true` and the loop will go on forever. We will come back to this general form of the `for` loop in a few sections to contrast it with one of the other loop types.

ForStatement:

```
for(initStatement; booleanExpression; iteratorStatement) statement
```

Before we move on to our next type of loop, there are a few more operators in Java that we should discuss because they are commonly used with loops. First, there is a set of assignment operators. We just saw that `++` and `--` combine doing an addition or a subtraction with an assignment. These always increment or decrement by one. Because it is common to do an operation on a variable and store the value back in the variable, Java provides shortcut operators to do this with all of their normal operators. All you have to do is follow the normal operator by an equal sign. So `+=` and `-=` increment and decrement by the value of the expression to the right of the operator. You can also use `*=`, `/=`, `%=`, or other operators that we haven't talked about at this point.

One last detail of the Java `for` loop is that you can use a comma to separate statements in the iterator statement. Consider the case where we want to increment two variables, *a* and *b*. If we were writing this normally inside of a method we could write the following as two statements.

```
a++;  
b++;
```

If we need to do this as the iterator of a `for` loop (and yes, this does happen occasionally) you can't put a semicolon between them because the syntax of the `for` loop only allows two semicolons. In that case, you can separate the two with a comma instead. Due to the challenge of coming up with examples of these features in something that is Greenfoot based, the following code prints something to screen. This method will print the first ten powers of 2.

```
public void printPowersOfTwo() {  
    for(int i=0,n=1; i<10; i++,n*=2) {  
        System.out.println("2^"+i+"="+n);  
    }  
}
```

Printing is something that we haven't seen previously, but it not only helps with examples, it can help a lot when you are writing code for Greenfoot and things aren't behaving the way you expect. Debugging, as we often call the process of finding and removing bugs, in a Greenfoot program can be challenging because the only feedback we normally get is in the movement of actors. Using the `inspect` option can show you other values, but that can get very tedious. If you include a print statement, Greenfoot will open another window and show you what you have printed out.

`System` is a standard class in Java. It has a member data called `out` that has a method in it called `println`. You can print basically any type of information that you want this way. In this case, we are printing a string. Note the way this works. We have declared two variables, *i* and *n*, that are both of

type `int`. Each time through the loop we want to increment `i` and double `n`. This demonstrates how we can do two things in the iterator and also demonstrates the usage of the `*=` operator. The expression `n*=2` is equivalent to `n=n*2`. You should add this method to one of your classes in `Greenfoot` and call it manually to see what happens.

4.3. Lists

Earlier we looked at the `Wombat` code in the `PSPGWombat` scenario and saw that the `Wombat` has fairly simple logic controlling how it moves. It winds up running around the edge of the screen and missing most of the leaves. You were able to manually make it so that the `wombat` would walk around and eat all the leaves. As we discussed though, the solutions you came up with likely only worked for the specific configuration of leaves that we started with. If the leaves were moved or more were added, some would likely be missed. What we want to do now is fix this problem and make it so that the code in the `Wombat` class will direct a `wombat` to eat all the leaves in the world, no matter the configuration. We also want to make it so that the way this is done is at least reasonably efficient in how many times the `wombat` has to act. We aren't going to go for a truly optimal solution. That could be a problem for later, but it is a hard enough problem that we aren't going to try to solve it now. For now we simply want a good solution even if it isn't the best.

The way we are going to go about doing this is to change what happens if the `wombat` isn't standing on a leaf in the `act` method. When the `wombat` is on a leaf, it will eat it, just like before. If not, it will find the nearest leaf and if it is walking the right direction, it will step forward, otherwise it will turn. So the logic looks very much like what was in the `Wombat` class earlier. We can write this method like this.

```
public void act() {
    if(foundLeaf()) {
        eatLeaf();
    }
    else {
        Leaf leaf=findClosestLeaf();
        moveToward(leaf);
    }
}
```

Now we have two smaller problems we have to solve: find the closest leaf and move toward the leaf.

We saw previously how we can check to see if we are standing on the same square as a leaf using the `getOneIntersectingObject()` method. Unfortunately, that method isn't going to help us in this case because it only tells us about objects very close to us. We need to know about objects further away. What we really need is a way to get hold of all of the leaves in the world. There isn't any method in `Actor` that will do this with us. However, even the way it was described implies that this is a method that should be found in `World` instead. Sure enough, there is a method called `getObjects()` in the `World` class. This method, like `getOneIntersectingObject()`, takes an argument of the `Class` for the type that you are interested in finding. What it returns is something we have never seen: `java.util.List`. We will call this just by the class name, `List`. The `java.util` part is a package name. Code in Java is organized into packages to help both people and the compiler find things. The `java.util` package contains a number of helpful classes, including different collections for storing things including `List`⁹.

The `List` class in Java works a lot like a list that you might make for something like groceries.

⁹ Technically `List` is an interface and not a class. Interfaces are similar to classes in that they define types and have methods in them. There are significant differences though that we will not worry about at this point.

Such a list stores a number of different things and allows you to interact with it in different ways. For the grocery list to be useful, you need to be able to do things like add stuff to it, remove things from it, and look through it. The following table was copied from the API page for List. This is part of the larger Java API that we won't look at in full until much later in the text. Many of these methods include things that we don't know about yet, but fortunately they don't matter to us at this point. A few of them give us the exact functionality we were talking about for our grocery list. The add() methods add things to the List. The clear() method removes everything from the list. The remove methods remove specific items from the list, the get() method lets us look through the list, and the size() method tells us how many things are on the list.

Method Summary	
boolean	add (E e) Appends the specified element to the end of this list (optional operation).
void	add (int index, E element) Inserts the specified element at the specified position in this list (optional operation).
boolean	addAll (Collection <? extends E > c) Appends all of the elements in the specified collection to the end of this list, in the order that they are returned by the specified collection's iterator (optional operation).
boolean	addAll (int index, Collection <? extends E > c) Inserts all of the elements in the specified collection into this list at the specified position (optional operation).
void	clear () Removes all of the elements from this list (optional operation).
boolean	contains (Object o) Returns true if this list contains the specified element.
boolean	containsAll (Collection <?> c) Returns true if this list contains all of the elements of the specified collection.
boolean	equals (Object o) Compares the specified object with this list for equality.
E	get (int index) Returns the element at the specified position in this list.
int	hashCode () Returns the hash code value for this list.
int	indexOf (Object o) Returns the index of the first occurrence of the specified element in this list, or -1 if this list does not contain the element.
boolean	isEmpty () Returns true if this list contains no elements.

Iterator <E>	iterator() Returns an iterator over the elements in this list in proper sequence.
int	lastIndexOf(Object o) Returns the index of the last occurrence of the specified element in this list, or -1 if this list does not contain the element.
ListIterator <E>	listIterator() Returns a list iterator over the elements in this list (in proper sequence).
ListIterator <E>	listIterator(int index) Returns a list iterator of the elements in this list (in proper sequence), starting at the specified position in this list.
E	remove(int index) Removes the element at the specified position in this list (optional operation).
boolean	remove(Object o) Removes the first occurrence of the specified element from this list, if it is present (optional operation).
boolean	removeAll(Collection<?> c) Removes from this list all of its elements that are contained in the specified collection (optional operation).
boolean	retainAll(Collection<?> c) Retains only the elements in this list that are contained in the specified collection (optional operation).
E	set(int index, E element) Replaces the element at the specified position in this list with the specified element (optional operation).
int	size() Returns the number of elements in this list.
List <E>	subList(int fromIndex, int toIndex) Returns a view of the portion of this list between the specified <code>fromIndex</code> , inclusive, and <code>toIndex</code> , exclusive.
Object []	toArray() Returns an array containing all of the elements in this list in proper sequence (from first to last element).
<T> T[]	toArray(T[] a) Returns an array containing all of the elements in this list in proper sequence (from first to last element); the runtime type of the returned array is that of the specified array.

In a number of these methods, including the add methods, you see E used as a type. This is the called a generic type and it is the type of what the List can store. When we declare a variable of type List, we can tell Java what type it stores by putting that type in angle braces after the name List. So if

we want a list of leaves we would declare the variable to have the type `List<Leaf>`. If you do this, then every place you see `E` in the API it will be treated as `Leaf`. So the `findClosestLeaf()` method can call `getWorld().getObjects(Leaf.class)` and get a list of leaves. Combining this with the variable declaration we get the following line of code to put in our method.

```
List<Leaf> leaves=getWorld().getObjects(Leaf.class);
```

If you add this into the code and try to compile, you will get an error message. This is because we haven't told the compiler where to go find the `List` class. We need to add a new import statement at the top of the file to do this. The following does what we need.

```
import java.util.List;
```

Now that we have a list of leaves, the question is, how do we figure out which one is closest? To answer this question imagine you are handed a list of cities and told to find the one that is closest to you. If you had a map, you might try looking at it to see which ones are closest, but that isn't a perfect method and that visual work is very hard level. We aren't going to try to program a system like that. Instead imagine that you have a way to look up the distance from where you are to any of the other cities. Now how could you use that to find the closest city? Basically, you are going to run down the list and check the distance to each city, then pick the smallest. That is still really high level. If you think about it, there are details that leaves out. As you go through the list, you are not just looking up the distance to each city in turn, you are also remembering which of the cities that you have looked at was closest, and perhaps the distance to that city so you don't have to look it up again.

With that we are just about close enough to make this into an algorithm. We start by having a memory of the closest city we have found so far and how far it is. Initially the city memory is blank because we haven't looked at any cities. Then we loop through all the cities on the list. If the city is the first one you have looked at, you it automatically the closest and you remember it and the distance to it. If it isn't the first, you compare the distance to the current city to the distance to the closest so far. If this one is closer, you forget the information on the old one and remember this one. When the loop is done, the city you are remembering is the closest one on the list. Let's convert that to code and see what it looks like.

```
public Leaf findClosestLeaf() {
    List<Leaf> leaves=getWorld().getObjects(Leaf.class);
    Leaf closest=null;
    int minDistance=1000;
    for(int i=0; i<leaves.size(); i++) {
        Leaf leaf=leaves.get(i);
        int dist=calcDistanceTo(leaf);
        if(closest==null || dist<minDistance) {
            closest=leaf;
            minDistance=dist;
        }
    }
    return closest;
}
```

We begin by getting the list of leaves and then making variables to store the closest leaf and the distance to it. That is followed by a loop that counts from zero up to the size of the list. Note that here

again we begin at zero and count up to one less than the number. We do this because that is how the items in a List are numbered. The `get()` method returns the first object in the list when we pass it 0.

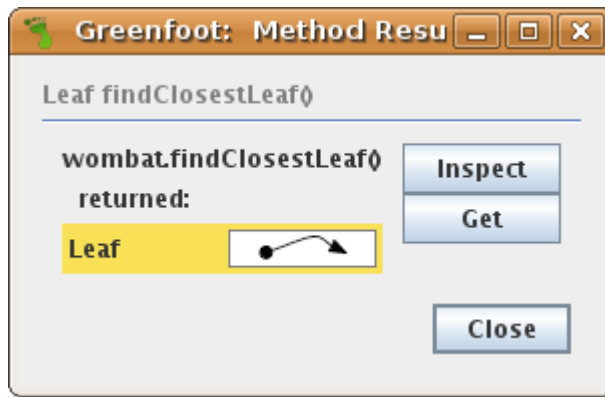
Inside of the loop the first line calls the `get()` method on the list to get the next item from the list. It then calls a method, `calcDistanceTo()`, that we haven't written yet, to calculate a distance and store that in the variable `dist`. This is followed by an if statement that check two things. The condition will be true if either this is the first object, in which case `closest` is still null, or the calculated distance is smaller than the minimum distance seen so far. If one of those is true, the current leaf and the distance to it are remembered. When the for loop is done, the value of `closest` is returned.

You should trace this method through to make certain you understand how it works. As a test of your understanding, consider what it would do if there were no leaves in the world. In that situation, the list won't have any elements in it and `leaves.size()` will return 0. What does the method return in that situation?

This method contained a call to `calcDistanceTo()` which was passed a Leaf object. There are many ways that this method could be written. By default you might consider using the Pythagorean theorem to return a normal Cartesian distance. This would work, but to some extent, it isn't a measure of the distance that our wombat will travel. That would be the distance if we moved on a straight line. We have so far only given our wombats the ability to move up, down, left, and right. That is something that we aren't going to change in this example. Indeed, Greenfoot locations have to be integers so making an Actor move on a smooth, straight line would take a bit more work. Given the style of movement, a better measure of distance is the separation in x plus the separation in y. That tells us how many times we will have to call `move()` in order to get from where we are to where the Leaf is. One thing to note is that the separations need to always be positive. If you just subtract one x value from the other, you will sometimes get a negative value. To fix this we need a way to take the absolute value. Java has a method that does this. It is called `abs()` and it is in the `Math` class. If we do this, the `calcDistanceTo()` method might look like the following.

```
public int calcDistanceTo(Leaf leaf) {  
    return Math.abs(getX()-leaf.getX())+Math.abs(getY()-leaf.getY());  
}
```

That finished off the `findClosestLeaf()` method. You should build a scenario and call it manually a few times just to make sure it doesn't crash. Greenfoot also provides two ways for you to tell if it really is returning the closest leaf. The display that pops up looks like the figure below. Because Leaf is a reference type, all Greenfoot shows us is an array saying that references something. To the right of this are two buttons: `Inspect` and `Get`. If you click on `Inspect`, it will bring up the `Inspection` window on the Leaf object that was returned. That will let you see the x and y coordinates which you can use to tell you if you got the right one. Alternately, if you click the `Get` button, when you move your mouse around the world, the leaf that was found will also move. Do this with a few different configurations so that you can make sure it works fairly generally.



Now all we have left to do is write the `moveToward()` method. When this was called in `act()` at the beginning of this section it was passed the `Leaf` object that we want to move toward. What we need to do in this method is make the `Wombat` step toward the leaf or turn, depending on whether it is facing a direction that will move it toward the leaf. This might be written in the following way.

```
public void moveToward(Leaf leaf) {
    if(facingGoodDirection(leaf)) {
        move();
    } else {
        turnLeft();
    }
}
```

Note that as we often do, we have included a call to a new method that doesn't yet exist, `facingGoodDirection()`. This method should return a boolean to us letting us know whether the direction we are currently facing will move us toward the leaf. In order to determine if the `Wombat` is facing a direction that will take it toward the leaf, we will use the `getXOffset()` and `getYOffset()` methods that were already in the class. These methods take a direction as an argument, and tell us how far we should move in x and y if we take a step. So when the direction has us heading up, `getXOffset()` returns 0 and `getYOffset()` returns -1.

The logic in `facingGoodDirection()` needs to check if the offset for our current direction moves us toward the leaf either in x or in y. As long as it is good for one axis it is appropriate to take a step that way. So how do we figure out if the offset will move us toward the leaf? There are many ways to do this and they have different complexities. The most obvious way is to write a bunch of if statements that do lots of checks comparing the offset values to the distance values. This way is also probably the most complex, the least efficient, and the more likely to be messed up. We will instead consider two alternate approaches.

The first one utilizes the fact that we have already written a method to find distance. Pretty much by definition a good step is one that decreases the distance. A bad step won't. So what we can do is make variables to remember the current position and distance, then move and calculate the modified distance. We then jump back and return whether the new distance was shorter than the old distance. This method might look like the following.

```
public boolean facingGoodDirection(Leaf leaf) {
    int cx=getX();
    int cy=getY();
    int cdist=calcDistanceTo(leaf);
```



```

    move();
    int ndist=calcDistanceTo(leaf);
    setLocation(cx,cy);
    return ndist<cdist;
}

```

The logic here is a little convoluted because we wind up moving and then going back to where we had been. This is because `facingGoodDirection()` isn't supposed to move us, it is only supposed to check if our direction is good. This method could be made nicer if `calcDistanceTo()` didn't always use our current position. If it allowed us to pass in a position, we could rewrite `facingGoodDirection()` in a more compact way. The following is a modified version of `calcDistanceTo()` that takes two additional arguments, an `x` and a `y`, and uses those in place of `getX()` and `getY()`.

```

public int calcDistanceTo(int x,int y,Leaf leaf) {
    return Math.abs(x-leaf.getX()+Math.abs(y-leaf.getY()));
}

```

Note that you can put this method in your class without erasing the old one. That is because Java allows methods to be **overloaded**. That means you can have two methods that have the same name as long as the arguments passed to them are different. If you call `calcDistanceTo()` and only pass it a `Leaf`, Java can easily tell that it should use the first version. If it is called with two ints and a `Leaf` Java will use this new version.

Using this modified `calcDistanceTo()` method we can go back a rewrite `facingGoodDirection` such that it doesn't actually move the `Wombat` and then pull it back. The modified version looks like this.

```

public boolean facingGoodDirection(Leaf leaf) {
    int cdist=calcDistanceTo(leaf);
    int ndist=calcDistanceTo(getX()+getXOffset(direction),getY()+getYOffset(direction),leaf);
    return ndist<cdist;
}

```

This version is a lot shorter and a lot easier to understand. Put this code into your `Wombat` and see what happens. The `wombat` should walk to every leaf, always going to the closest one, and eat each one when it gets to it. It works just fine until the `wombat` eats the last leaf. At that point you get an error. What type of error is it? Why is it happening?

The problem is that our `moveToward()` method assumes that it actually has a leaf. When there are no leaves left in the world, the `findClosestLeaf()` method will return `null` and that leads to a `NullPointerException`. To fix this, we should make it so that the `Wombat` doesn't try to move toward the leaf if no leaf was found. This can be done by making a simple addition to our `act()` method.

```

public void act() {
    if(foundLeaf()) {
        eatLeaf();
    }
    else {
        Leaf leaf=findClosestLeaf();
        if(leaf!=null) {
            moveToward(leaf);
        }
    }
}

```

```

    }
}

```

Let's consider one other method of writing `facingGoodDirection()`. This method isn't going to calculate distances. Instead, it is going to compare the direction to the leaf to the offset for the current direction. We are facing a good direction if the direction toward the leaf in x or y is the same as the offset along that axis. Checking if things are the same direction might seem like it would require a fair bit of logic, but it turns out that we can do it nicely with just a bit of math. If you have two numbers and you want to check if they are either both positive or both negative, simply multiply them. If the result is positive, the two match. If it is negative they don't. Using this, we can write `facingGoodDirection()` as follows.

```

public boolean facingGoodDirection(Leaf leaf) {
    int dx=leaf.getX()-getX();
    int dy=leaf.getY()-getY();
    return dx*getXOffset(direction)>0 || dy*getYOffset(direction)>0;
}

```

The last line could be modified to use addition instead of the logical or as shown here.

```

    return dx*getXOffset(direction)+dy*getYOffset(direction)>0;

```

This version is basically doing a dot product on two vectors to see if the angle between them is smaller than a right angle. This is actually the more general solution. The other solution could run into problems if we extended the `Wombat` so that it could move along diagonals.

There is one last improvement that we could make to this code. We specifically wrote all of our methods to deal with leaves. By using the `Leaf` type for things like the argument to `facingGoodDirection`, we have produced methods that work well for this particular problem, but which aren't fully general. Nearly everything we did in this section would work more generally with the `Actor` type. Indeed, `Leaf` is a subtype of `Actor` because it inherits from it, so if we altered the code to work with `Actor`, it would continue to work with `Leaf` as well. The following code shows the complete `Wombat` class with `Leaf` specific code generalized to work with `Actor`. In most places this just required changing types. We also change variable names to make it so that things make sense. The one exception is the `findClosestLeaf()` method which we have to add a parameter to so that we can tell it what type of `Actor` we are looking for.

```

public class Wombat extends Actor {
    private int direction;
    private int leavesEaten;

    public Wombat() {
        setDirection(0);
        leavesEaten = 0;
    }

    /**
     * Do whatever the wombat likes to to just now.
     */
    public void act() {

```

```

    if(foundLeaf()) {
        eatLeaf();
    }
    else {
        Actor leaf=findClosestActor(Leaf.class);
        if(leaf!=null) {
            moveToward(leaf);
        }
    }
}

/**
 * Check whether there is a leaf in the same cell as we are.
 */
public boolean foundLeaf() {
    Actor leaf = getOneObjectAtOffset(0, 0, Leaf.class);
    return leaf != null;
}

/**
 * Eat a leaf.
 */
public void eatLeaf() {
    Actor leaf = getOneObjectAtOffset(0, 0, Leaf.class);
    if(leaf != null) {
        // eat the leaf...
        getWorld().removeObject(leaf);
        leavesEaten = leavesEaten + 1;
    }
}

/**
 * Move one cell forward in the current direction.
 */
public void move() {
    if (canMove()) {
        setLocation(getX()+getXOffset(direction),getY()+getYOffset(direction));
    }
}

/**
 * Test if we can move forward. Return true if we can, false otherwise.
 */
public boolean canMove() {
    World myWorld = getWorld();
    int x = getX()+getXOffset(direction);
    int y = getY()+getYOffset(direction);
    // test for outside border
    if (x >= myWorld.getWidth() || y >= myWorld.getHeight()) {
        return false;
    }
    else if (x < 0 || y < 0) {
        return false;
    }
}

```

```

        return true;
    }

    /**
     * Turns towards the left.
     */
    public void turnLeft() {
        direction=direction-1;
        if(direction<0) {
            direction=3;
        }
        setDirection(direction);
    }

    /**
     * Sets the direction we're facing.
     */
    public void setDirection(int dir) {
        direction = dir;
        setRotation(90*direction);
    }

    /**
     * Tell how many leaves we have eaten.
     */
    public int getLeavesEaten() {
        return leavesEaten;
    }

    /**
     * Returns the x offset for the specified direction.
     * @param dir The direction we want an offset for.
     * @return The x offset.
     */
    public int getXOffset(int dir) {
        if(dir==0) {
            return 1;
        } else if(dir==2) {
            return -1;
        } else {
            return 0;
        }
    }

    /**
     * Returns the y offset for the specified direction.
     * @param dir The direction we want an offset for.
     * @return The y offset.
     */
    public int getYOffset(int dir) {
        if(dir==1) {
            return 1;
        } else if(dir==3) {
            return -1;
        } else {

```

```

        return 0;
    }
}

public Actor findClosestActor(Class c) {
    List<Actor> list=getWorld().getObjects(c);
    Actor closest=null;
    int minDistance=1000;
    for(int i=0; i<list.size(); i++) {
        Actor actor=list.get(i);
        int dist=calcDistanceTo(actor);
        if(closest==null || dist<minDistance) {
            closest=actor;
            minDistance=dist;
        }
    }
    return closest;
}

public int calcDistanceTo(Actor actor) {
    return Math.abs(getX()-actor.getX()+Math.abs(getY()-actor.getY()));
}

public int calcDistanceTo(int x,int y,Actor actor) {
    return Math.abs(x-actor.getX()+Math.abs(y-actor.getY()));
}

public void moveToward(Actor actor) {
    if(facingGoodDirection(actor)) {
        move();
    } else {
        turnLeft();
    }
}

public boolean facingGoodDirection(Actor actor) {
    int dx=actor.getX()-getX();
    int dy=actor.getY()-getY();
    return dx*getXOffset(direction)+dy*getYOffset(direction)>0;
}
}

```

4.4. The for-each Loop

In the findClosestActor() method we used a standard for loop to count through all the indexes in the list so that we could get out each element and do what we needed to with it. This is such a common activity that Java includes a second version of the for loop that does exactly this with a slightly shorter syntax. Not only is this syntax shorter, for some types of lists it can also be more efficient. This alternate style is called the for-each loop because it is used to loop through each element of a collection. The rule for the for-each loop is shown here.

ForEachStatement:

for(Type name:iterableExpression) statement

Inside of the parentheses you have what looks like a variable declaration followed by a colon and the thing you want to run through the elements of. The *Type* should be whatever type is stored in our list. The *name* is the variable name that we want the values to be called by inside the loop. After the colon can be any expression that evaluates to the type *Iterable*. For right now the only thing we know of that fits this description is a list.

Using this we could rewrite the `findClosestActor()` method as shown here.

```
public Actor findClosestActor(Class c) {
    List<Actor> list=getWorld().getObjects(c);
    Actor closest=null;
    int minDistance=1000;
    for(Actor actor:list) {
        int dist=calcDistanceTo(actor);
        if(closest==null || dist<minDistance) {
            closest=actor;
            minDistance=dist;
        }
    }
    return closest;
}
```

Notice that this shortens the code a little and can make it easier to read. Note that you are never required to use the for-each loop instead of the normal for loop. Indeed, there are times when you really need the normal for loop because you need the index value. However, when the for-each loop can be used it is often a superior choice.

4.5. The while Loop

The first use of a loop in this chapter was for the `addFruit` method. Indeed, we have seen many occasions in this book where we wanted to add a number of objects at random locations in the world. The for loop is great for doing this without having to copy and paste our statement and we were able to use random numbers to make it so that the objects go in random locations. There is just one problem with this. As you have probably seen, every so often the objects that are added are put at the same location.

The for loop is the most commonly used loop statement in Java, but it isn't the only one. There are two others. The first of which is the while loop. The rule for the while statement is given below. In many ways, while is the simplest loop in the Java language. It isn't used as much as for in large part because it is too simple. The while loop has a condition and a body. That is it, nothing else.

!!! smart adding of houses and food in the CityScape game so they don't overlap.

4.6. The do-while Loop

There is one other type of loop in Java

4.7. Recursion

!!! Recursion?

Chapter Rules

ForStatement:

for(*initStatement*; *booleanExpression*; *iteratorStatement*) *statement*

1. Cover for loop for repetition of a task.
2. Greenfoot lists and using loops to go through them.
3. For-each loop.
4. Real meaning of inheritance. Cover super and instanceof operator.

5. Arrays and Lists

5. Processing data
6. Using arrays to store direction for movement in Greenfoot.
7. Explain inheritance and polymorphism with lists.

6. File Processing

8. Exceptions (required for I/O)