

Learning Computing With Robots in C++

Edited by
Deepak Kumar



Institute for Personal Robots in Education
www.roboteducation.org

This text is provided as is with no guarantees of any kind.

Permission is hereby granted to reproduce facsimiles of this text in any form for distribution provided an e-mail note is sent to the editor at dkumar@cs.brynmawr.edu describing the manner and intended purpose of distribution.

Permission is also granted to modify the content of this text in any manner seemed desirable by instructors. We recommend that all derivations from this text be distributed through IPRE in the same manner as this text. To obtain the permission we require the same process as laid out for distribution above.

We strongly recommend distributing the text and its modifications with a not for profit motive.

We are making this text available in the manner described above so as to enable the widest possible accessibility of the content as well as the ideas illustrated in this text. We hope that users will respect and honor and promote such practices with their own works.

We have taken great care to attribute all images used in this text to their original sources. This text is intended for educational purposes and such use constitutes fair use. Where feasible, we have obtained permissions to use copyrighted images. Please contact the editor in case there are any pending issues in this regard. All images not explicitly attributed in the text are created by members of the IPRE team.

This version was produced by Bruce MacLennan in January 2011 based on the Fall '08 Python text, revised June 2009.

Institute for Personal Robotics in Education
Fall 2008, revised January 17, 2011.

Contributors

IPRE was created to explore the use of personal robots in education with primary funding from Microsoft Research, Georgia Institute of Technology, and Bryn Mawr College. This text would not have been possible without the hardware, software, and course materials developed by members of the IPRE team.

Ben Axelrod, Georgia Institute of Technology
Tucker Balch, Georgia Institute of Technology
Douglas Blank, Bryn Mawr College
Natasha Eilbert, Bryn Mawr College
Ashley Gavin, Bryn Mawr College
Gaurav Gupta, Georgia Institute of Technology
Mansi Gupta, Bryn Mawr College
Mark Guzdial, Georgia Institute of Technology
Jared Jackson, Microsoft Research
Ben Johnson, Georgia Institute of Technology
Deepak Kumar, Bryn Mawr College
Bruce MacLennan, University of Tennessee, Knoxville
Marwa Nur Muhammad, Bryn Mawr College
Keith O'Hara, Georgia Institute of Technology
Shikha Prashad, Bryn Mawr College
Richard Roberts, Georgia Institute of Technology
Jay Summet, Georgia Institute of Technology
Monica Sweat, Georgia Institute of Technology
Stewart Tansley, Microsoft Research
Daniel Walker, Georgia Institute of Technology



Contents

Preface	v
Chapter 1	
The World of Robots	1
Chapter 2	
Personal Robots	19
Chapter 3	
Building Robot Brains	41
Chapter 4	
Sensing From Within	71
Chapter 5	
Sensing The World	101

Chapter 6	
Insect-Like Behaviors	139
Chapter 7	
Behavior Control	169
Chapter 8	
Sights & Sounds	201
Chapter 9	
Image Processing & Perception	231
Chapter 10	
Artificial Intelligence	261
Chapter 11	
Computers & Computation	289
Chapter 12	
Fast, Cheap & Out of Control	327
C++ Overview	339
Myro Overview	351

Preface

What is a computer? A personal computer?
What is a robot? A personal robot?
What is computing? What is personal computing?

This book will introduce you to the world of computers, robots, and computing. Computers, as we know them today have been around less than 75 years. To our young audience for whom this book is designed, that may seem like an eternity. While computers are all around us, the world of computing is just in its infancy. It is a baby!

In this book, you will learn that computing is no more about computers than astronomy is about telescopes. In that respect, the roots of computing could be traced back more than 2000 years to Euclid in 300 BC. Could it be that computing is actually as old as or even older than some of the other disciplines? Prior to the development of modern computers, the word *computing* mostly implied numerical calculation. Today, computing encompasses all forms of electronic manipulation of objects: numbers, text, sound, music, photographs, movies, genes, motors, etc., and even ideas! Every day we are discovering new things that can be computed.

Robots are not that much older than computers. We will learn that, contrary to popular belief, a robot is any mechanism that is guided by automated control.

A robot can be controlled as a result of computation. In fact, in today's robots computation is a key component.

Advances in the past few decades have brought about many changes in the way we use and live with technology. Computers have become smaller, cheaper, faster, and more personal. The same is also true for robots. In this book we have created a confluence of these two technologies to use as a basis for introducing the world of computing to students.

The word *personal* is key to this new confluence. Computers and robots are no longer the realm of large corporate offices and industrial manufacturing scenarios. They have become personal in many ways: they help you write your term papers, store and organize your photographs, your music collection, your recipes, and keep you in touch with your friends and family. Many people might even dismiss the computer as an appliance: like a toaster oven, or a car. However, as you will learn in this book, a computer is far more versatile than most appliances. It can be used as an extension of your mind. We were not kidding above when we mentioned that computers can be used to process ideas. It is in this form, that a computer becomes an enormously powerful device regardless of your field of interest. How personal you can make a computer or a robot is entirely up to you and your abilities to control these devices. This is where this book comes in. You will learn the basics of how a computer and a robot is controlled and how you can use these ideas to personalize your computers even further, limited only by your imagination.

How to use this book

We have strived in creating this book to make the learning experience engaging, entertaining, and enlightening. This book comes with a robot kit and you will make extensive use of your personal robot as you go through the book. As a part of your learning process, you will be bringing all of the material in this book to life in the form of computer programs and robot behaviors. You will create new and interesting behaviors for your robots. In order to get the most out of this book, you will need the following:

- This book.
- A robot. This book uses the Scribbler robot.
- A Fluke Lancet dongle.
- A Bluetooth enabled computer. If your computer does not have built-in Bluetooth, you can purchase a USB Bluetooth dongle to add that functionality. You can use any of the following operating systems: MACOS (Version 10.4 or later), Windows (XP or Vista), or Linux.
- Software: the C++ version of Myro. This book is based on C++, but other editions use other languages, such as Python.
- Some accessories: A basic gamepad controller, Sharpie permanent color markers, a flashlight, spare AA batteries, poster or any other large drawing paper, a carrying case, some craft materials for decorating the robot, etc.

You can order the robot, the Fluke dongle, the Bluetooth dongle, as well as a software CD, and the Python-based book from Georgia Robotics Inc. (georgiarobotics.com). For more information on additional sources, documentation (including a copy of this book), and latest updates visit the IPRE web site at roboteducation.org.

While our approach is tightly integrated with the use of a personal robot we will also take frequent excursions into other forms of computing: numbers, text, sounds, music, images, etc. The thing to remember is that all the concepts of computing introduced in this book can be used to do computing on all kinds of objects.

To get the most out of this book we would recommend reading each chapter first and then bringing all the ideas discussed in the chapter to life by actually creating computer programs and robot behaviors. Before proceeding to the next chapter, try and do most of the exercises at the end of each chapter and think about further applications that come to mind. Try out a few of those and then proceed to the next chapter. We would also like to encourage you to work in an environment where you can share your work and ideas with others around you. After all, you are going to be creating several robot behaviors. Find the opportunity to demonstrate your robot's newly programmed

capabilities to your friends, classmates, and even the wider world (through web sites like youtube.com). Make learning from this book a social activity and you will find yourself learning more. The bottom-line keyword here is *doing*. Computing is an abstract concept, but the only way to grasp it is by *doing it*. Like that famous Nike commercial, *Just do it!*

What you will learn

By using this book as described above you will learn a number of fundamental concepts in computing and robotics. We have written the text for entry-level students into computing. Other than being able to use a computer to do something, we do not expect our readers to have any additional background in computing. Some of the material in the book does rely on basic high school level mathematics. It is ideally suited for a college-level introductory computing course which is widely recognized as CS1 in the computing education community. By the time a reader works through the entire book they will have covered the following topics:

- Fundamentals of computing: problem solving, program design, basic syntax and semantics of the programming language C++; the role of names; data values: numbers, strings, and booleans; data structures: vectors, lists, structures, dictionaries, objects; functions, parameters, basic scoping rules, assignment, expressions, control (sequential, conditional and iteration constructs), modules and libraries (APIs); Simple input-output, file I/O, and other forms of I/O devices; testing and error checking; interactive design.
- Program development process: using an IDE, interactive style computing, debugging, and documentation.
- Basic robotics: robot sensing and control; sensor types, synchronous and asynchronous operations; control paradigms: direct or reactive, Braitenberg style, subsumption style, neural network; behavior design; real-time/embedded control.
- Media computation: numbers, sound and music, image: generation, processing, and understanding; gamepad controllers and game playing.
- Social implications of computing and robotics.

- Computation: history and essentials of algorithms; efficiency; limits of computing; other issues and open problems in computing.

For instructors, it should be clear that the above set of concepts provide sufficient coverage of the CS1 topics recommended by the latest ACM/IEEE Curricular guidelines¹ as well as the guidelines developed for liberal arts degrees in computer science². In many respects the content of the book goes well beyond the recommended CS1 guidelines. It does not explicitly mention recursion or sorting algorithms. We have deliberately chosen to omit these topics. Sorting of lists and other sequential structures is available in standard C++ libraries and students will use these to do sorting. We often find that students discover recursion in the course of doing exercises from the text. Rather than provide a separate chapter or a section on recursion, we have left it as a topic to be introduced by the instructor if she/he so chooses. We do use recursion in example programs in Chapter 10. It would also be fairly trivial for instructors to include class discussions of sorting algorithms if they feel so inclined.

Further, for instructors, the coverage and sequence of topics in this book will vastly differ from that found in most traditional CS1 texts. We have deliberately taken a *contextual approach* to introduce ideas. The nature of the context, *robots and computers*, has driven the nature, ordering, and sequencing of topics. We have found that this is a more engaging, exciting, and enlightening framework. We hope that you will be able to share this excitement while you are teaching a course based on this material. Further details on how to use this book and additional materials is provided in the Instructor's manual accompanying this text. Much of this work is an outcome of research conducted under the banner of IPRE specifically for developing a fresh new approach to CS1. We invite you to refer to our project web site (roboteducation.org) and read some of the papers and articles that elaborate much more on the underlying research and its goals.

¹ Computing Curricula 2001, ACM Journal on Educational Resources in Computing (JERIC) Vol. 1, No. 3, Fall 2001.

² A Model Curriculum for a Liberal Arts Degree in Computer Science, LACS, ACM JERIC, Vol. 7, No. 2, June 2007.

To students, we have taken efforts to make the material engaging and exciting for you. We hope that we have succeeded somewhat in this. This is an open source development project, so we hope that you will take the time to provide valuable feedback, as well as ideas and content for further improving this book. We do hope that you will exploit all the fun learning opportunities provided in this book.



1

The World of Robots

*I wouldn't ever want them to be brought back to Earth. We built them for Mars, and Mars is where they should stay. But **Spirit** and **Opportunity** have become more than just machines to me. The rovers are our surrogates, our robotic precursors to a world that, as humans, we're still not quite ready to visit.*

-: Steve Squyres in **Roving Mars**, Hyperion, 2005.


Opposite page: Mars Rover.

Photo courtesy of NASA/JPL-Caltech

The picture on the opposite page is among one of thousands sent back by *Spirit* and *Opportunity* from the surface of Mars. It goes without saying that it will probably be several years, decades or even more, before a human ever sets foot on Mars. The rovers *Spirit* and *Opportunity* landed on Mars in January 2004 as robot geologists whose mission was to analyze rocks and soils of the red planet in search for clues to past presence of water on the planet. The two robots were expected to last about 90 days. Years later, they were still exploring the planet's surface and sending invaluable geological and pictorial data from the planet.

Back on Earth, the same month as the rovers landing on Mars, the Tumbleweed robot rover traveled 40 miles across Antarctica's polar plateau transmitting local meteorological data back to its base station via satellite. Besides surviving adverse conditions on Mars and Antarctica, robots are slowly becoming household consumer items. Take for instance, the Roomba from iRobot Corporation. Introduced in 2002, several million Roombas have been sold to vacuum and clean floors.

One commonality among the robots mentioned above is that they are all designed for very specific tasks: analyze rocks and soils on the surface of Mars, meteorology on the polar cap, or vacuuming a room. Yet, the core of robot technology is almost as easy to use as computers. In this course you will be given a personal robot of your own. Through this personal robot, you will learn to give it instructions to do a variety of tasks. Like the robots mentioned above your robot is also a rover. However, unlike the robots above, your personal robot does not come pre-programmed to do any specific task. It has certain basic capabilities (that you will learn about) and it can be programmed to make use of its capabilities to do various tasks. We hope that the process of learning about the capabilities of your robot and making it do different things will be exciting and fun for you. In this chapter, we introduce you to the world of robots and then introduce you to your own personal robot and some of its capabilities.



The rim of Victoria Crater on Mars.

The rover *Opportunity* has been superimposed on the crater rim to show scale.

Photo courtesy of JPL/NASA/Cornell University, October, 2006.

What is a robot?

The Merriam-Webster Online Dictionary gives the following definitions of the word *robot*:

1. a machine that looks like a human being and performs various complex acts (as walking or talking) of a human being; also a similar but fictional machine whose lack of capacity for human emotions is often emphasized; and also an efficient insensitive person who functions automatically
2. a device that automatically performs complicated often repetitive tasks
3. a mechanism guided by automatic controls

In today's world, the first two definitions will probably be considered archaic (the third interpretation in the first definition notwithstanding). Robots were originally conceived as human-like entities, real or fictional, devoid of emotions, that performed tasks that were repetitive or full of drudgery. Today's robots come in all kinds of shapes and sizes and take on all kinds of tasks (see below for some examples). While many robots are put to use for repetitive or dull tasks (including the Roomba; unless you enjoy the therapeutic side effects of vacuuming :-), robots today are capable of doing a lot more than implied by the first two definitions above. Even in fictional

robots the lack of emotional capacity seems to have been overcome (see for instance Steven Spielberg's movie, *Artificial Intelligence*).



For our purposes, the third definition is more abstract and perhaps more appropriate. A *robot* is a mechanism or an artificial entity that can be guided by automatic controls. The last part of the definition, *guided by automatic controls*, is what we will focus on in this course. That is, given a mechanism capable of such guidance, what is involved in creating its controls?

A Short History of Robots

Modern robots were initially conceived as industrial robots designed to assist in automated manufacturing tasks. The first commercial robot company, Unimation, was created nearly 50 years ago. As the use of robots in industrial manufacturing grew, people also started experimenting with other uses of robots. Earlier industrial robots were mainly large arms that were attached to a fixed base. However, with the development of mobile robots people started to find uses for them in other domains. For instance, in exploring hazardous environments ranging from radioactive sites, volcanoes, finding and destroying mines, military surveillance, etc. We started this chapter by introducing you to two Mars rover robots. The first ever planetary rover landed on Mars in 1997. Increasingly in the last decade or so robots have ventured into newer and more exciting areas like medicine (Google: *robotic surgery*, *robot wheelchair*, etc.), toys and entertainment (Google: *Pleo*, *SONY Aibo*, *LEGO Mindstorms*, etc.), and even education (Google: *IPRE*). Some of the most exciting developments in robotics are still in


Today, it is hard to imagine life without a web search engine. While there are several search engines available, the one provided by Google Inc. has become synonymous with web searching. So much so that people use it as a common phrase: "Google it!"

You may have your own personal preference for a search engine. Go ahead use it and search for the items suggested here.

research stages where, for example, in Artificial Intelligence research people are trying to develop intelligent robots and also using robots to understand and explore models of human intelligence. Here we have provided some pointers (do the searches mentioned above) for examples of various robots and their uses. There are numerous web sites where you can look up more about the history of robots. We will leave that as an exercise.

Robots and Computers

In the last few decades computers have become increasingly ubiquitous. Most likely you are reading this sentence on a computer. If you're reading this text online, the text itself is coming to you from another computer (located somewhere on the western banks of the Delaware River in south-eastern parts of the state of Pennsylvania in the United States of America). On its journey from the computer in Pennsylvania to your computer, the text has probably traveled through several computers (several dozen if you are outside the state of Pennsylvania!). What makes this journey of this text almost instantaneous is the presence of communication networks over which the internet and the World Wide Web operate. Advances in the technologies of wireless communication networks make it possible to access the internet from nearly any place on the planet. The reason that you are sitting in front of a computer and learning about robots is primarily because of the advent of these technologies. While robots are not quite as ubiquitous as computers, they are not too far behind. In fact, it is precisely the advances in computers and communications technologies that have made it feasible for you to become more familiar with the world of robots.



A Postage stamp titled *World of Invention (The Internet)* was issued by UK's Royal Mail on March 1, 2007 honoring the development of the World Wide Web.

The relationship between robots and computers is the basis for the use of the phrase *automatic controls* in describing a robot. Automatically controlling a robot almost always implies that there is a computer involved. So, in the process of learning about and playing with robots you will also uncover the world of computers. Your robot has a computer embedded in it. You will be controlling the robot through your computer. Moreover, you will do this over a wireless communication technology called *bluetooth*. Initially, for our purposes, learning to automatically control a robot is going to be synonymous with learning to control a computer. This will become more obvious as we proceed through these lessons.

Automating control involves specifying, in advance, the set of tasks the robot or the computer is to perform. This is called *programming*. Programming involves the use of a *programming language*. Today, there are more programming languages than human languages! Perhaps you have heard of some of them: Java, C, Python, etc. In this course, we will do all our robot programming in the programming language *C++* (pronounced *see-plus-plus*). C++ is widely used in computer science and engineering for all kinds of programming.

While we are talking about computers and languages, we should also mention the *Myro* (for *My robot*) software system. Myro was developed by us to simplify the programming of robots. Myro provides a small set of robot commands that extend the C++ language. This makes it easy, as you will see, to specify automatic controls for robots.

A Robot of Your Own: The Scribbler

The scribbler robot, shown here is also a rover. It can move about in its environment. The wheels, and its other functions, can be controlled through a computer via a wireless interface. Your laboratory assistants will provide you with a Scribbler and the required components to enable wireless



The Scribbler Robot

communication. Once configured, you will be able to control the robot's movements (and all other features) through the computer. Besides moving, your robot can also play sounds (beeps) and, with the help of a pen inserted in its pen port, it can draw a line wherever it goes (hence its name, *Scribbler*). The robot can move forward, backward, turn, spin, or perform any combination of these movements giving it adequate functionality to travel anywhere on the surface of an environment. Besides roving, the Scribbler robot can also sense certain features of its environment. For example, it is capable of sensing a wall or an obstacle, or a line on the floor. We will discuss the Scribbler's sensing capabilities later.

Do This

The first few activities show you how you to set up the computer and the robot and will help you get familiarized with your Scribbler. This will involve the following four activities:

1. First things first: Setup Myro
2. Name your robot
3. Drive your robot around
4. Explore a little further

You may need the assistance of your instructor for the first activity to ensure that you know how to set up and use your robot for the remainder of the text.

Dear Student:


Every chapter in this book will include several hands-on activities. These are short learning exercises designed carefully to explore some of the concepts presented in the chapter. It is important that you do all of the activities in the chapter before moving on to the next chapter.

We would also recommend trying out some (or all) of the exercises suggested at the end to gain a better understanding.

1. First things first: Setup Myro

At the time you received your robot, its software and hardware was configured for use. The software we will be using for controlling the robot is called, **Myro** (for **My Robot**) which works in conjunction with the C++ language. In this, the first exercise, we will start the robot and the software and ensure that the software is able to successfully communicate with the robot through your computer. If the C++ version of Myro has not been installed on your computer, you should follow your lab instructor's directions to install it.

In a typical session, you will write a C++ program to connect to the robot through the Myro library, and then to control the robot through it. We have set up the system so that all communication between the computer and the robot occurs wirelessly over a Bluetooth connection. Bluetooth technology is a common wireless communication technology that enables electronic devices to talk to each other over short distances. For example, Bluetooth is most commonly used in cell phones to enable wireless communication between a cell phone (which may be in your pocket) and your wireless headset. This kind of communication requires two physical devices that serve as receivers and



The Fluke Dongle adds Bluetooth and other capabilities to the Scribbler.



The Scribbler robot with the Fluke Dongle.

transmitters. In the Scribbler kit you received, there is a pair of these Bluetooth devices: one connects into the scribbler (Fluke Dongle) and the other into the USB port of your computer. If your computer has a built-in Bluetooth capability, you may not need the one that goes into your computer. Go ahead and make sure that these devices are plugged in, the robot is turned on, and so is your computer. Open up the Bluetooth configuration facility on your computer to get the port number being used for communication. On Windows this will be something like COM12 (the actual number will vary). Detailed instructions for setting up the Bluetooth connection came with your robot. Make sure you have established the connection and obtained the port number before proceeding.

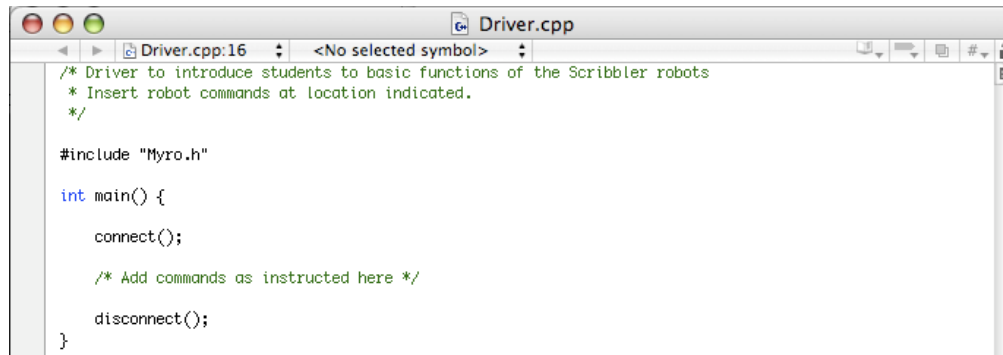
2. Name your robot

In this exercise, we will connect to the robot and make it do something simple, like make it beep. Then, we will give the robot a name to personalize it. These tasks can be performed by using the following steps:

1. Edit the driver
2. Connect to the robot
3. Make the robot beep
4. Give the robot a name

Since this is your very first experience with using robots, we will provide detailed instructions to accomplish the task outlined above.

1. Edit the driver: We are providing some sample C++ code, called `Driver.cpp`, for programming the robot; your instructor will tell you where you can find it. You should copy this file into a folder where you plan to store all your robot programs. Once done, navigate to that folder and open it with the editor you will use for C++ programs (e.g., vi). A window such as the following window should appear on your computer screen:



```
Driver.cpp:16 <No selected symbol>
/* Driver to introduce students to basic functions of the Scribbler robots
 * Insert robot commands at location indicated.
 */

#include "Myro.h"

int main() {
    connect();

    /* Add commands as instructed here */

    disconnect();
}
```

The next step is to use Myro to connect to the robot.

2. Connect to the robot: The exact procedure for connecting to the robot will vary depending on the system you are using (e.g., OS X, linux, Windows, or Vista) and how it has been set up. Therefore you should consult the documentation for the Myro system you are using or get some advice from your instructor.

Look at this line, which appears in the preceding figure:

```
#include "Myro.h"
```

This informs the C++ preprocessor to include the contents of the file `Myro.h`, which contains the C++ interface to the Myro library. This `#include` directive is something you will use each time you want to control the robot. When you run your program, it starts obeying commands at the beginning of your main program, indicated by the line `int main() {`. (You will learn the precise meaning of these and the other cryptic expressions in this program soon enough!) First, your control program connects to the robot by using the command `connect()`. When your program executes this command, the computer attempts to communicate with the robot. If this is successful, the program responds with `Connected to Robot`. You will understand the meaning of the other commands in `Driver.cpp` as you learn more about C++, but now it's time to make your robot do something. First, let us give it a command to make a beep so that we know that we are in control of the robot:

3. Make the robot beep: Insert the following line in your `Driver.cpp` program at the location indicated:

```
robot.beep(1, 880);
```

The command above directs the robot to make a beeping sound at 880 Hertz for 1 second. Go ahead and try it. Your robot will beep for 1 second at 880 Hz. Go ahead and try the following variations to hear different beeps:

```
robot.beep(0.5, 880);  
robot.beep(0.5, 500);  
etc.
```

So now, you should realize that you are in control of the robot. By issuing simple commands like the ones above, you can make the robot perform different behaviors. Now, we can learn the command to give the robot a new name.

4. Give the robot a name: Suppose we wanted to name the robot `Shrek`. To do this, all you have to do is give it the following command:

```
robot.setName("Shrek");
```

Whatever name you decide to give your robot, you can specify it in the command above replacing the words `Shrek`. From now on, that will be the name of the robot. How do we know this for a fact? Go ahead and try asking it its name by putting the following in your file:

```
cout << "Hello, my name is " << robot.getName() << "!\n";
```

(Don't worry now about all the cryptic characters; you will learn about them eventually!) If the robot's name is `Shrek` it will print this on the terminal:

```
Hello, my name is Shrek!
```

As in the preceding figure, when your driver program is done controlling the robot, it should disconnect from it:

```
disconnect();
```

Congratulations! You have now completed the first Exercise and you are well on your way to more fun and exciting things with your robot. Before we move on, it would be a good idea to review what you just did. Each program controlling a robot begins by including the files to interface the C++/Myro library and to connect to the robot. From then on, you can issue any command to the robot. At the end, include the command to disconnect from the robot.

The Myro library contains dozens of commands that enable various kinds of robot behaviors. In the next few weeks we will be learning several robot commands and learning how to use them to program complex robot behaviors. One thing to remember at this juncture is that all commands are being issued in the C++ language. Thus, as you learn more about your robot and its behaviors, you will also be learning the C++ language.

One characteristic of programming languages (like C++) is that they have a very strict way of typing commands. That is, and you may already have experienced this above, the language is very precise about what you type and how you type it. Every parenthesis, quotation mark, and upper and lower case letter that makes up a command has to be typed exactly as described. While the rules are strict, luckily there aren't too many of them. Soon you will get comfortable with this syntax and it will become second nature. The precision in syntax is required so that the computer can determine exactly one interpretation for the command resulting in desired action. For this reason, computer languages are often distinguished from human languages by describing them as *formal languages* (as opposed to *natural languages* that are used by humans).

3. Drive the robot around

In this exercise,* we will introduce you to a way of making the robot move about in its environment manually controlled by a game pad device (see picture to the right). As above, place the robot on an open floor, turn the robot on, and edit your driver to connect to the robot. You may already have this from Exercise 2 above. Also, plug the game pad controller into an available USB port of your computer. At the prompt, enter the following command:

```
robot.gamepad();
```



The game pad controller.

In response to this command, you will get some help text printed on the console window showing what would happen if you pressed various game pad buttons. If you look in the picture of the game pad controller above, you will notice that it has eight (8) blue buttons (numbered 1 through 8 in the picture), and an axis controller (the big blue swivel button on the left). The axis controller can be used to move the robot around. Go ahead and try it. Pressing each of the numbered buttons will result in different behaviors, some will make the robot beep, some will make the computer speak or say things. Button#1 will result in the robot taking a picture of whatever it is currently seeing through its camera and display it on your computer screen. Button#8 will quit from the game pad control mode.

Spend some time experimenting with the various game pad control features. See how well you can navigate the robot to go to various places, or follow a wall, or go round and round something (like yourself!). You may also place the robot on a big piece of paper, insert a Sharpie pen in its pen port and then

* The current C++/Myro interface does not support the game pad.

move it around to observe its scribbling. Can you scribble your name (or initials)? Try a pattern or other shapes.

Without creating a program, this is an effective remote way of controlling the movements of your robot. The next exercise asks you to try and issue commands to the robot to move.

4. Explore a little further

OK, now you are on your own. Edit `Driver.cpp` to connect to the robot, and give commands to move forward, backward, turn left and right, and spin. Use the commands: `forward(SPEED)`, `backward(SPEED)`, `turnLeft(SPEED)`, `turnRight(SPEED)`, and `rotate(SPEED)`. `SPEED` can be any number between `-1.0` and `1.0`. For example, to make the robot go forward at half speed, use:

```
robot.forward(0.5);
```

These and all other robot commands are detailed in the Myro Reference Manual. This would be a good time to review the descriptions of all the commands introduced in this section.

Myro Review

```
#include "Myro.h"
```

This preprocessor directive inserts into your program all the robot commands available in the C++/Myro library. We will use this whenever we intend to write programs that use the robot.

```
connect();
```

This command establishes a wireless communication connection with the robot and reports when the connection has been made.


```
disconnect();
```

This command disconnects the wireless connection to your robot.

```
robot.bEEP (<TIME>, <FREQUENCY>);
```

Makes the robot beep for <TIME> seconds at frequency specified by <FREQUENCY>.

```
robot.getName();
```

Returns the name of the robot.

```
robot.setName (<NEW NAME>);
```

Sets the name of the robot to <NEW NAME>. The new name should be enclosed in double quotes, no spaces, and not more than 16 characters long. For example: setName ("Bender").

```
robot.gamepad();
```

Enables manual control of several robot functions and can be used to move the robot around.*

C++ Review

Note: All commands you type (including the Myro commands listed above) are essentially C++ commands. Later, in this section we will list those commands that are a part of the C++ language.

Exercises

1. Where does the word *robot* come from? Explore the etymology of the words *robot* and *robotics* and write a short paper about it.

* The game pad is not supported by the current C++/Myro interface.

2. What are Asimov's Laws of robotics? Write a viewpoint essay on them.
3. Look up the Wikipedia entry on robots, as well as the section in AI Topics (see links above). Write a short essay on the current state of robots.
4. Write a short essay on a robot (real or fictional) of your choice. Based on what you have learned from your readings, evaluate its capabilities.
5. *Spirit* and *Opportunity* were not the first rovers to land on Mars. On July 4, 1997, the *Mars Pathfinder* landed on Mars with a payload that included the *Sojourner* rover. The United States Postal Service issued the stamp shown here to commemorate the landing. This is perhaps the first ever real robot to appear on a postage stamp! Find out what you can about the *Mars Pathfinder* mission and compare the *Sojourner* rover with *Spirit* and *Opportunity*.
6. Through the exercises, you have experienced a subset of capabilities of the Scribbler robot. Reflect/write about the physical capabilities of the Scribbler and the kinds of tasks you could make it perform.
7. Insert a pen (provided in your kit) in the robot's pen port. Place the robot on a surface where it is OK to write/draw. Drive the robot around with the game pad controller. It will scribble on the paper as it moves. Observe its scribbles by moving it forward and then backwards. Does it trace its path exactly? Why or why not?
8. Using the game pad operation,* make your robot Scribble your name on the floor. You may find this difficult for several reasons. Try to make the robot write your initials instead. Also, see if you can guide the robot to draw a five point star. This task is in some sense not too different from controlling a robot to perform surgery. Research the capabilities of today's surgical robots and write a short paper about it.




Facsimile of the Mars
Pathfinder Postage Stamp

* The game pad is not supported by the current C++/Myro interface.

9. Using the game pad controller* draw the Bluetooth logo (see picture) using a pen inserted in the Scribbler robot. Do a web search for Harald Blåtand and read more about the runic alphabets.

Harald Blåtand Gormson

What's in a name?



The Bluetooth logo is derived from *runic* alphabet letters H and B juxtaposed together. HB for Harald Blåtand a Scandinavian King (from the 10th century AD) who was legendary in uniting Denmark and Norway. The wireless technology we use today is named in his honor (Blåtand means “Bluetooth”) because the technology itself was developed by Ericsson, a Scandinavian company. The technology is designed to unite computers and telecomm devices. Bluetooth devices are most commonly found in cell phones. We’re using it here to communicate between your robot and the computer.

Further Reading

1. Wikipedia entry on Robots (<http://en.wikipedia.org/wiki/Robot>)
2. AI Topics: Robots from the American Association for Artificial Intelligence (AAAI) (<http://www.aaai.org/AITopics/html/robots.html>)
3. Social Robots are robots that interact with and learn from people around them. Here is an interview with Cynthia Breazeal who heads the Robotic Life Group at MIT's Media Lab. (<http://www.pbs.org/saf/1510/features/breazeal.htm>)
4. Visit the online Robot Hall of Fame and find out more about the real and fictional robots that have been inducted into it. (<http://www.robthalloffame.org/>)

* The game pad is not supported by the current C++/Myro interface.



2

Personal Robots

Every Pleo is autonomous. Yes, each one begins life as a newly-hatched baby Camarasaurus, but that's where predictability ends and individuality begins. Like any creature, Pleo feels hunger and fatigue - offset by powerful urges to explore and be nurtured. He'll graze, nap and toddle about on his own -when he feels like it! Pleo dinosaur can change his mind and his mood, just as you do.

From: www.pleoworld.com

Opposite page: Pleo robots
Photo courtesy of UGOBE Inc.

Most people associate the personal computer (aka the PC) revolution with the 1980's but the idea of a personal computer has been around almost as long as computers themselves. Today, on most college campuses, there are more personal computers than people. The goal of One Laptop Per Child (OLPC) Project is to “provide children around the world with new opportunities to explore, experiment, and express themselves” (see www.laptop.org). Personal robots, similarly, were conceived several decades ago. However, the personal robot ‘revolution’ is still in its infancy. The picture on the previous page shows the Pleo robots that are designed to emulate behaviors of an infant Camarasaurus. The Pleos are marketed mainly as toys or as mechatronic “pets”. Robots these days are being used in a variety of situations to perform a diverse range of tasks: like mowing a lawn; vacuuming or scrubbing a floor; entertainment; as companions for elders; etc. The range of applications for robots today is limited only by our imagination! As an example, scientists in Japan have developed a baby seal robot (shown on the opposite page) that is being used for therapeutic purposes for nursing home patients.

Your Scribbler robot is your personal robot. In this case it is being used as an educational robot to learn about robots and computing. As you have already seen, your Scribbler is a rover, a robot that moves around. Such robots have become more prevalent in the last few years and represent a new dimension of robot applications. Roaming robots have been used for mail delivery in large offices and as vacuum cleaners in homes. Robots vary in the ways in which they move about: they can roll about like small vehicles (like the lawn mower, Roomba, Scribbler, etc.), or even ambulate on two, three, or more legs (e.g. Pleo). The Scribbler robot moves on three wheels, two of which are powered. In this chapter, we will get to know the Scribbler in some more detail and also learn about how to use its commands to control its behavior.

The Scribbler Robot: Movements

In the last chapter you were able to use the Scribbler robot through Myro to carry out simple movements. You were able to start the Myro software, connect to the robot, and then were able to make it beep, give it a name, and move it around using a joystick. By inserting a pen in the pen port, the

scribbler is able to trace its path of movements on a piece of paper placed on the ground. It would be a good idea to review all of these tasks to refresh your memory before proceeding to look at some more details about controlling the Scribbler.



The Paro Baby Seal Robot.
Photo courtesy of National
Institute of Advanced
Industrial Science and
Technology, Japan (paro.jp).

If you hold the Scribbler in your hand and take a look at it, you will notice that it has three wheels. Two of its wheels (the big ones on either side) are powered by motors. Go ahead turn the wheels and you will feel the resistance of the motors. The third wheel (in the back) is a free wheel that is there for support only. All the movements the Scribbler performs are controlled through the two motor-driven wheels. In Myro, there are several commands to control the movements of the robot. The command that directly controls the two motors is the `robot.motors` command:

```
robot.motors(LEFT, RIGHT);
```

In the command above, `LEFT` and `RIGHT` can be any value in the range `[-1.0...1.0]` and these values control the left and right motors, respectively. Specifying a negative value moves the motors/wheels backwards and positive values move it forward. Thus, the command:

```
robot.motors(1.0, 1.0);
```

will cause the robot to move forward at full speed, and the command:

```
robot.motors(0.0, 1.0);
```

will cause the left motor to stop and the right motor to move forward at full speed resulting in the robot turning left. Thus by giving a combination of left and right motor values, you can control the robot's movements. Myro has also

provided a set of often used movement commands that are easier to remember and use. Some of them are listed below:

```
forward(SPEED)
backward(SPEED)
turnLeft(SPEED)
turnRight(SPEED)
stop()
```

Another version of these commands takes a second argument, an amount of time in seconds:

```
forward(SPEED, SECONDS)
backward(SPEED, SECONDS)
turnLeft(SPEED, SECONDS)
turnRight(SPEED, SECONDS)
```

Providing a number for SECONDS in the commands above specifies how long that command will be carried out. For example, if you wanted to make your robot traverse a square path, you could issue the following sequence of commands:

```
robot.forward(1, 1);
robot.turnLeft(1, .3);
robot.forward(1, 1);
robot.turnLeft(1, .3);
robot.forward(1, 1);
robot.turnLeft(1, .3);
robot.forward(1, 1);
robot.turnLeft(1, .3);
```

of course, whether you get a square or not will depend on how much the robot turns in 0.3 seconds. There is no direct way to ask the robot to turn exactly 90 degrees, or to move a certain specified distance (say, 2 ½ feet). We will return to this later.

You can also use the following movement commands to translate (i.e. move forward or backward), or rotate (turn right or left):


```
translate(SPEED)
rotate(SPEED)
```

Additionally, you can specify, in a single command, the amount of translation and rotation you wish use:

```
move(TRANSLATE_SPEED, ROTATE_SPEED)
```

In all of these commands, `SPEED` can be a value between `[-1.0...1.0]`.

You can probably tell from the above list that there are a number of redundant commands (i.e. several commands can be specified to result in the same movement). This is by design. You can pick and choose the set of movement commands that appear most convenient to you. It would be a good idea at this point to try out these commands on your robot.

Do This: Edit your driver program to connect to the robot and to try out the following movement commands on your Scribbler:

First make sure you have sufficient room in front of the robot (place it on the floor with a few feet of open space in front of it).

```
robot.motors(1, 1);
wait(5);
robot.motors(0, 0);
```

(The `wait` command waits for a specified number of seconds.) Observe the behavior of the robot. Specifically, notice if it does (or doesn't) move in a straight line after issuing the first command. You can make the robot carry out the same behavior by issuing the following commands:

```
robot.move(1.0, 0.0);
wait(5);
robot.stop();
```

Go ahead and try these. The behavior should be exactly the same. Next, try making the robot go backwards using any of the following commands:

```
robot.motors(-1, -1);  
robot.move(-1, 0);  
robot.backward(1);
```

Again, notice the behavior closely. In rovers, precise movement, like moving in a straight line, is difficult to achieve. This is because two independent motors control the robot's movements. In order to move the robot forward or backward in a straight line, the two motors would have to issue the exact same amount of power to both wheels. While this technically feasible, there are several other factors than can contribute to a mismatch of wheel rotation. For example, slight differences in the mounting of the wheels, different resistance from the floor on either side, etc. This is not necessarily a bad or undesirable thing in these kinds of robots.

Under similar circumstances even people are unable to move in a precise straight line. To illustrate this point, you can try the experiment shown on right.

For most people, the above experiment will result in a variable movement. Unless you really concentrate hard on walking in a straight line, you are most likely to display similar variability as your Scribbler. Walking in a straight line requires constant feedback and adjustment, something humans are quite adept at doing. This is hard for robots to do. Luckily, roving does not require such precise moments anyway.

Do humans walk straight?

Find a long empty hallway and make sure you have a friend with you to help with this. Stand in the center of the hallway and mark your spot. Looking straight ahead, walk about 10-15 paces without looking at the floor. Stop, mark your spot and see if you walked in a straight line.

Next, go back to the original starting spot and do the same exercise with your eyes closed. Make sure your friend is there to warn you in case you are about to run into an object or a wall. Again, note your spot and see if you walked in a straight line.

Do This: Review all of the other movement commands listed above and try them out on your Scribbler. Again, note the behavior of the robot from each of these commands.

Defining New Commands

Trying out simple commands by modifying `Driver.cpp` is a nice way to get to know your robot's basic features. We will continue to use this each time we want to try out something new. However, making a robot carry out more complex behaviors requires several series of commands. Having to type these over and over in a control program can get tedious. C++ provides a convenient way to package a series of commands into a brand new command called a *function*. For example, if we wanted the Scribbler to move forward and then move backward (like a yoyo), we can define a new command (function) called `yoyo` as follows:

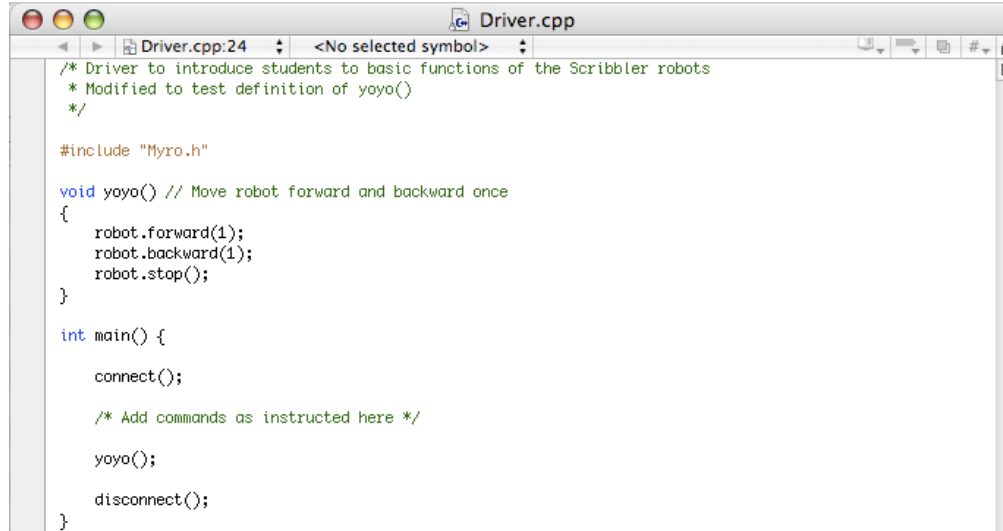
```
void yoyo() // Move robot forward and backward once
{
    robot.forward(1);
    robot.backward(1);
    robot.stop();
}
```

The first line defines the *name* of the new command/function to be `yoyo`. (We will explain later the exact significance of the word `void`, which indicates that we are defining a command.) The left and right curly braces (`{}`) mark the beginning and end of the function's *body*. The lines in between are slightly indented and contain the commands that make up the `yoyo` behavior. That is, to act like a yoyo, move forward and then backward and then stop. The indentation is important and is used in all programming languages to improve readability. We will have more to say about this later.

Put the definition of `yoyo` into `Driver.cpp` before the definition of `main`. After the new command has been defined, you can try it by entering the command into the usual place in `main` as shown below:

```
yoyo();
```

After you have done this editing, your `Driver.cpp` file should look like the following figure:



```
Driver.cpp
Driver.cpp:24 <No selected symbol>
/* Driver to introduce students to basic functions of the Scribbler robots
 * Modified to test definition of yoyo()
 */

#include "Myro.h"

void yoyo() // Move robot forward and backward once
{
    robot.forward(1);
    robot.backward(1);
    robot.stop();
}

int main() {
    connect();

    /* Add commands as instructed here */

    yoyo();

    disconnect();
}
```

Do This: If you have your Scribbler ready, go ahead and try out the new definition above by modifying your driver to include the definition and command `yoyo()` as shown above.

Observe the robot's behavior when you run the control program. You may need to run the program several times. The robot momentarily moves and then stops. If you look closely, you will notice that it does move forward and backwards.

In C++, you can define new commands (functions) by using the syntax as shown above. Note also that defining a new command doesn't mean that the commands inside it get carried out. You have to explicitly issue the new command to do this. This is useful because it gives you the ability to use the command over and over again (for example, writing `yoyo()` more than once). Issuing the new function like this in C++ is called, *invocation*. Upon invocation, all the commands that make up the function's definition are executed in the sequence in which they are listed in the definition.

How can we make the robot's `yoyo` behavior more pronounced? That is, make it move forward for, say 1 second, and then backwards for 1 second, and then stop? You can use the `SECONDS` option in `forward` and `backward` movement commands as shown below:

```
void yoyo() /* Move robot
forward and backward for 1
sec. */
{
    robot.forward(1,1);
    robot.backward(1,1);
}
```

The same behavior can also be accomplished by using the command `wait`, which is used as shown below:

```
wait(SECONDS);
```

where `SECONDS` specifies the amount of time the robot waits before moving on to the next command. (Note that `wait` is not a command *to the robot*, and therefore it is not preceded by “`robot.`”.) In effect, the robot continues to do whatever it had been asked to do just prior to the `wait` command for the amount of time specified in the `wait` command. That is, if the robot was asked to move forward and then asked to wait for 1 second, it will move forward for 1 second before applying the command that follows the `wait`. Here is the complete definition of `yoyo` that uses the `wait` command:

```
void yoyo() // Move robot forward and backward for 1 sec.
{
    robot.forward(1);
    wait(1);
    robot.backward(1);
    wait(1);
    robot.stop();
}
```

Scribbler Tip:

Remember that your Scribbler runs on batteries and with time they will get drained. When the batteries start to run low, the Scribbler may exhibit erratic movements. Eventually it stops responding. When the batteries start to run low, the Scribbler's red LED light starts to blink. This is your signal to replace the batteries.

Do This: Go ahead and try out the new definitions exactly as above and issue the command to the scribbler. What do you observe? In both cases you should see the robot move forward for 1 second followed by a backward movement for 1 second and then stop.

Adding Parameters to Commands

Take a look at the definition of the `yoyo` function above and you will notice the use of parentheses, `()`, both when defining the function as well as when using it. You have also used other functions earlier with parentheses in them and probably can guess their purpose. Commands or functions can specify certain *parameters* (or values) by placing them within parentheses. For example, all of the movement commands, with the exception of `stop` have one or more numbers that you specify to indicate the speed of the movement. The number of seconds you want the robot to wait can be specified as a parameter in the invocation of the `wait` command. Similarly, you could have chosen to specify the speed of the forward and backward movement in the `yoyo` command, or the amount of time to wait. Below, we show three definitions of the `yoyo` command that make use of parameters:

```
void yoyo1(double speed)
{
    robot.forward(speed, 1);
    robot.backward(speed, 1);
}

void yoyo2(double waitTime)
{
    robot.forward(1, waitTime);
    robot.backward(1, waitTime);
}

void yoyo3(double speed, double waitTime)
{
    robot.forward(speed, waitTime);
    robot.backward(speed, waitTime);
}
```

In the first definition, `yoyo1`, we specify the speed of the forward or backward movement as a parameter. The word `double` means that the following parameter can be a fractional number, such as 0.5. (The exact significance of the word `double` will not concern us now.) Using this definition, you can control the speed of movement with each invocation. For example, if you wanted to move at half speed, you can issue the command:

```
yoyo1(0.5);
```

Similarly, in the definition of `yoyo2` we have parameterized the wait time. In the last case, we have parameterized both speed and wait time. For example, if we wanted the robot to move at half speed and for 1 ½ seconds each time, we would use the command:

```
yoyo3(0.5, 1.5);
```

This way, we can customize individual commands with different values resulting in different variations on the yoyo behavior. Notice in all of the definitions above that we did not have to use the `stop()` command at all. Why?

Saving New Commands in Files

As you can imagine, while working with different behaviors for the robot, you are likely to end up with a large collection of new functions. It would make sense then that you do not have to type in the definitions over and over again. C++ enables you to define new functions and store them in files in a folder on your computer, which can then be easily used over and over again. Let us illustrate this by defining two behaviors: a parameterized yoyo behavior and a wiggle behavior that makes the robot wiggle left and right. The two definitions are given below:

```
// File: moves.h
// Purpose: Two useful robot commands to try out as a file.

void yoyo(double speed, double waitTime)
{
    robot.forward(speed);
    wait(waitTime);
    robot.backward(speed);
    wait(waitTime);
    robot.stop();
}

void wiggle(double speed, double waitTime)
{
    robot.rotate(-speed);
    wait(waitTime);
    robot.rotate(speed);
    wait(waitTime);
    robot.stop();
}
```

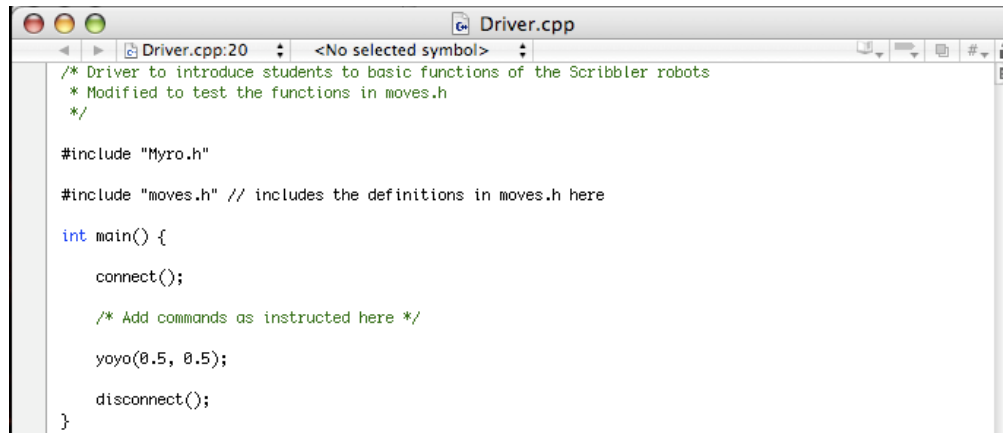
All lines beginning with the “//” characters are called comments. These are simply annotations that help us understand and document programs in C++. You can place these comments anywhere, including right after a command. The // signs clearly mark the beginning of the comment and anything following it on that line is not interpreted as a command by the computer. C++ also permits multi-line comments, which extend from the characters “/*” to the next occurrence of “*/”. You have already seen several examples of this style of comments. Comments are quite useful and we will make liberal use of them in all our programs.

Do This: To store the `yoyo` and `wiggle` behaviors in a file, you can use your program editor to enter the text containing the two definitions and then save them in a file (let’s call it `moves.h`) in your Myro folder (same place you have the `Driver.cpp` file). All files containing C++ definitions of this kind will end with the filename extension `.h` (for *header* file) and you should make sure they are always saved in the same folder as the `Driver.cpp` file. This will make it easy for you as well as the C++ compiler to locate your files when you use them.

Once you have created the file, there is a simple way you can use it. In `Driver.cpp`, just enter the directive:

```
#include "moves.h"
```

in front of the definition of `main`, where the two new commands will be used. For example, the following shows how to use the `yoyo` function after including the `moves` module:



```
Driver.cpp
Driver.cpp:20 <No selected symbol>
/* Driver to introduce students to basic functions of the Scribbler robots
 * Modified to test the functions in moves.h
 */
#include "Myro.h"
#include "moves.h" // includes the definitions in moves.h here
int main() {
    connect();

    /* Add commands as instructed here */
    yoyo(0.5, 0.5);
    disconnect();
}
```

As you can see from above, accessing the commands defined in a file is similar to accessing the capabilities of `Myro` in the `Myro.h` file. This is a nice feature of C++. In C++, you are encouraged to extend the capabilities of any system by defining your own functions, storing them in files and then using them by including them. Thus including definitions from the `moves.h` file is no different than including definitions from the `Myro.h` file. The directive

```
#include "Myro.h"
```

includes everything in the `Myro.h` file, just as though you had typed it at that place in the driver program. Everything defined in the `Myro.h` library is listed and documented in the C++ `Myro` Reference Manual. The nice thing that this facility provides is that you can now define your own set of commands that

extend the basic commands available in Myro to customize the behavior of your robot. We will be making use of this over and over again in this course.

Functions as Building Blocks

Now that you have learned how to define new commands using existing ones, it is time to discuss a little more C++. The basic syntax for defining a C++ function takes the form:

```
void <FUNCTION NAME> (<PARAMETERS>)
{
    <SOMETHING>
    ...
    <SOMETHING>
}
```

That is, to define a new function, start by using the word `void` followed by the name of the function (`<FUNCTION NAME>`) followed by `<PARAMETERS>` enclosed in parenthesis followed by a left curly brace (`{`). (You will learn later that function definitions do not have to begin with `void`.) This line is followed by the commands that make up the function definition (`<SOMETHING> ... <SOMETHING>`). A right curly brace (`}`) completes the definition. Usually each command is placed on a separate line, and all lines that make up the definition should be indented (aligned) the same amount. The number of spaces that make up the indentation is not that important as long as they are all the same. This may seem a bit awkward and too restricting at first, but you will soon see the value of it, for it makes the definition(s) more readable. For example, look at the following definitions for the `yoyo` function:

```
void yoyo(double speed, double waitTime) {
    robot.forward(speed);
    wait(waitTime);
    robot.backward(speed);
    wait(waitTime);
    robot.stop(); }
```

```
void yoyo(double speed, double waitTime)
{ robot.forward(speed); wait(waitTime);
  robot.backward(speed); wait(waitTime);
  robot.stop(); }
```

Some program editors and *integrated design environments* (IDEs) help you in making your indentations consistent by automatically indenting the next line like the previous.

Another feature built into some IDEs and program editors, which improves readability of C++ programs, is the use of color highlighting. Notice in the above examples (where we use screen shots from an IDE) that pieces of your program appear in different colors. For example, the word `void` in a function definition appears in blue, and the name of your function, `yoyo`, appears in black. Other colors are also used in different situations, look out for them. This particular IDE displays all C++ words (like `void`) in blue and all names defined by you (like `yoyo`) in black. What do you think the other colors indicate?

The idea of defining new functions by using existing functions is very powerful and central to computing. By defining the function `yoyo` as a new function using the existing functions (`forward`, `backward`, `wait`, `stop`) you have *abstracted* a new behavior for your robot. You can define further higher-level functions that use `yoyo` if you want. Thus, functions serve as basic building blocks in defining various robot behaviors, much like the idea of using building blocks to build bigger structures. As an example, consider defining a new behavior for your robot: one that makes it behave like a yoyo twice, followed by wiggling twice. You can do this by defining a new function as follows:

```
void dance()
{
  yoyo(0.5, 0.5);
  yoyo(0.5, 0.5);
  wiggle(0.5, 1);
  wiggle(0.5, 1);
}
```

Do This: Go ahead and add the `dance` function to your `moves.h` file. Put it after the definitions of `yoyo` and `wiggle`, since in C++ we have to define things before we use them. Try the `dance` command on the robot. Now you have a very simple behavior that makes the robot do a little shuffle dance.

Guided by Automated Controls

Earlier we agreed that a robot is a “mechanism guided by automated controls”. You can see that by defining functions that carry out more complex movements, you can create modules for many different kinds of behaviors. The modules make up the programs you write, and when they are invoked on the robot, the robot carries out the specified behavior. This is the beginning of being able to define automated controls for a robot. As you learn more about the robot’s capabilities and how to access them via functions, you can design and define many kinds of automated behaviors.

Summary

In this chapter, you have learned several commands that make a robot move in different ways. You also learned how to define new commands by defining new C++ functions. Functions serve as basic building blocks in computing and defining new and more complex robot behaviors. C++ has specific syntax rules for writing definitions. You also learned how to save all your function definitions in a file and then using them as a module by including it. While you have learned some very simple robot commands, you have also learned some important concepts in computing that enable the building of more complex behaviors. While the concepts themselves are simple enough, they represent a very powerful and fundamental mechanism employed in almost all software development. In later chapters, we will provide more details about writing functions and also how to structure parameters that customize

individual function invocations. Make sure you do some or all of the exercises in this chapter to review these concepts.

Myro Review

```
robot.backward(SPEED);
```

Move backwards at `SPEED` (value in the range -1.0...1.0).

```
robot.backward(SPEED, SECONDS);
```

Move backwards at `SPEED` (value in the range -1.0...1.0) for a time given in `SECONDS`, then stop.

```
robot.forward(SPEED);
```

Move forward at `SPEED` (value in the range -1.0..1.0).

```
robot.forward(SPEED, TIME);
```

Move forward at `SPEED` (value in the range -1.0...1.0) for a time given in seconds, then stop.

```
robot.motors(LEFT, RIGHT);
```

Turn the left motor at `LEFT` speed and right motor at `RIGHT` speed (value in the range -1.0...1.0).

```
robot.move(TRANSLATE, ROTATE);
```

Move at the `TRANSLATE` and `ROTATE` speeds (value in the range -1.0...1.0).

```
robot.rotate(SPEED);
```

Rotates at `SPEED` (value in the range -1.0...1.0). Negative values rotate right (clockwise) and positive values rotate left (counter-clockwise).

```
robot.stop();
```

Stops the robot.

```
robot.translate(SPEED);
```

Move in a straight line at `SPEED` (value in the range -1.0...1.0). Negative

values specify backward movement and positive values specify forward movement.

```
robot.turnLeft(SPEED);
```

Turn left at `SPEED` (value in the range -1.0..1.0)

```
robot.turnLeft(SPEED, SECONDS);
```

Turn left at `SPEED` (value in the range -1.0..1.0) for a time given in seconds, then stops.

```
robot.turnRight(SPEED);
```

Turn right at `SPEED` (value in the range -1.0..1.0)

```
robot.turnRight(SPEED, SECONDS);
```

Turn right at `SPEED` (value in the range -1.0..1.0) for a time given in seconds, then stops.

```
wait(TIME)
```

Pause for the given amount of `TIME` seconds. `TIME` can be a decimal number.

C++ Review

```
#include "FILE NAME"
```

Includes the contents of the file named `FILE NAME` just as though it had been typed at this location in the current file. The file is sought in the same directory as your other C++ programs.

```
void <FUNCTION NAME>(<PARAMETERS>)
```

```
{  
    <SOMETHING>  
    ...  
    <SOMETHING>  
}
```

Defines a new command/function named `<FUNCTION NAME>`. A function name should always begin with a letter and can be followed by any sequence of letters, numbers, or underscores (`_`), and not contain any spaces. Try to choose names that appropriately describe the command being defined.

Exercises

1. Compare the robot's movements in the commands `turnLeft(1)`, `turnRight(1)` and `rotate(1)` and `rotate(-1)`. Closely observe the robot's behavior and then also try the motor commands:

```
robot.motors(-0.5, 0.5);  
robot.motors(0.5, -0.5);  
robot.motors(0, 0.5);  
robot.motors(0.5, 0);
```


Do you notice any difference in the turning behaviors? The `rotate` commands make the robot turn with a radius equivalent to the width of the robot (distance between the two left and right wheels). The `turn` command causes the robot to spin in the same place.

2. Insert a pen in the scribbler's pen port and then issue it commands to go forward for 1 or more seconds and then backwards for the same amount. Does the robot travel the same distance? Does it traverse the same trajectory? Record your observations.

3. Measure the length of the line drawn by the robot in Exercise 2. Write a function `travel(DISTANCE)` to make the robot travel the given `DISTANCE`. You may use inches or centimeters as your units. Test the function on the robot a few times to see how accurate the line is.

4. Suppose you wanted to turn/spin your robot a given amount, say 90 degrees. Before you try this on your robot, do it yourself. That is, stand in one spot, draw a line dividing your two feet, and then turn 90 degrees. If you have no way of measuring, your turns will only be approximate. You can study the behavior of your robot similarly by issuing it turn/spin commands and making them wait a certain amount. Try and estimate the wait time required to turn 90 degrees (you will have to fix the speed) and write a function to turn that amount. Using this function, write a behavior for your robot to transcribe a square on the floor (you can insert a pen to see how the square turns out).

5. Generalize the wait time obtained in Exercise 3 and write a function called `degreeTurn (DEGREES)`. Each time it is called, it will make the robot turn the specified degrees. Use this function to write a set of instructions to draw a square.
6. Using the functions `travel` and `degreeTurn`, write a function to draw the Bluetooth logo (See Chapter 1, Exercise 9).
7. Choreograph a simple dance routine for your robot and define functions to carry it out. Make sure you divide the tasks into re-usable moves and as much as possible parameterize the moves so they can be used in customized ways in different steps. Use the building block idea to build more and more complex series of dance moves. Make sure the routine lasts for at least several seconds and it includes at least two repetitions of the entire sequence. You may also make use of the beep command you learned from the last section to incorporate some sounds in your choreography.
8. Record a video of your robot dance and then dub it with a soundtrack of your choosing. Use whatever video editing software is accessible to you. Post the video online on sites like YouTube to share with friends.
9. Lawn mower robots and even vacuuming robots can use specific *choreographed* movements to ensure that they provide full coverage of the area to be serviced. Assuming that the area to be mowed or cleaned is rectangular and without any obstructions, can you design a behavior for your Scribbler to provide full coverage of the area? Describe it in writing. [Hint: Think about how you would mow/vacuum yourself.]



3

Building Robot Brains

What a splendid head, yet no brain.
Aesop (620 BC-560 BC)

Opposite page: Home Simpson's Brain
Photo courtesy of The Simpson's Trivia (www.simpsonstrivia.com)

If you think of your robot as a creature that acts in the world, then by programming it, you are essentially building the creature's brain. The power of computers lies in the fact that the same computer or the robot can be supplied a different program or brain to make it behave like a different creature. For example, a program like Firefox or Explorer makes your computer behave like a web browser. But switching to your Media Player, the computer behaves as a DVD or a CD player. Similarly, your robot will behave differently depending upon the instructions in the program that you have requested to run on it. In this chapter we will learn about the structure of C++ programs and how you can organize different robot behaviors as programs.

The world of robots and computers, as you have seen so far is intricately connected. You have been using a computer to connect to your robot and then controlling it by giving it commands. Most of the commands you have used so far come from the Myro library which is specially written for easily controlling robots. The programming language we are using to do the robot programming is C++. C++ is a general purpose programming language. By that we mean that one can use C++ to write software to control the computer or another device like a robot through that computer. Thus, by learning to write robot programs you are also learning how to program computers. Our journey into the world of robots is therefore intricately tied up with the world of computers and computing. We will continue to interweave concepts related to robots and computers throughout this journey. In this chapter, we will learn more about robot and computer programs and their structure.

Basic Structure of a Robot Brain

The basic structure of a C++ program (or a robot brain) is shown below:

```
int main() {  
    <do something>  
    <do something>  
    ...  
}
```

This is essentially the same as defining a new function. In fact, in C++ the *main program*, where execution begins, is always a function called `main`. In general, the structure of your robot programs will be as shown below (we have provided line numbers so we can refer to them):

```
Line 1: #include "Myro.h"
Line 2: <any other includes>

Line 3: <function definitions>
Line 4: int main() {
Line 5:     connect();
Line 6:     <do something>
Line 7:     <do something>
Line 8:     ...

Line 9:     disconnect();
Line 10: } // end of main program
```

Every robot brain program will begin with the first line (Line 1). This, as you have already seen, includes the Myro library and other declarations for communicating with the robot. In case you are using any other libraries, you will then include them (this is shown in Line 2). This is followed by the definitions of functions (Line 3), and then the definition of the function `main`. The first line of `main` (Line 5) includes the commands to connect to the robot. Finally, the last line of the function body (Line 9) includes the commands to disconnect from it. When you run this program it will begin execution on Line 5 and finish after executing Line 9. In order to illustrate this, let us write a robot program that makes it do a short dance using the `yoyo` and `wiggle` movements defined in the last chapter.

```
/*
 * File: dance.cpp
 * Purpose: A simple dance routine
 */

// First include standard declarations

#include "Myro.h"
```

```
// Define the new functions...

void yoyo(double speed, double waitTime) {
    robot.forward(speed, waitTime);
    robot.backward(speed, waitTime);
}

void wiggle(double speed, double waitTime) {
    robot.motors(-speed, speed);
    wait(waitTime);
    robot.motors(speed, -speed);
    wait(waitTime);
    robot.stop();
}

// The main dance program
int main() {
    connect();
    cout << "Running the dance routine..." << endl;
    yoyo(0.5, 0.5);
    wiggle(0.5, 0.5);
    yoyo(1, 1);
    wiggle(1, 1);
    cout << "...Done" << endl;
    disconnect();
}
```

We have used a new C++ command in the definition of the `main` function: the `cout <<` command. This command will print out the text enclosed in double quotes (") when you run the program. This is followed by `<< endl`, which ends the output line, so that subsequent output goes on a new line. This program is not much different from the `dance` function defined in the previous chapter except we are using a spin motion to wiggle. However, instead of naming the function `dance` we are calling it `main`. This is necessary in C++ to identify the main program in a program file.

Do This: In order to run this program on the robot, you can start your program editor or IDE, enter the program in it, save it as a file (`dance.cpp`) and then compile it according to the procedure on your computer system.

When you run the program you will notice that the robot carries out the dance routine specified in the `main` program. Also notice the two messages printed in your computer's window. These are the results of the output command (`cout <<`). This is a very useful command in C++ and can be used to output essentially anything you ask it to. While you are in this session, go ahead and change the output command to the following:

```
Speak("Running the dance routine");
```

`speak` is a Myro command that enables speech output from your computer. Go ahead and change the other output command also to the `speak` command and try your program. Once done, enter some other `speak` commands in your program. For example:

```
speak("Dude! Pardon me, would you have any Grey Poupon?");
```

The speech facility is built into most computers these days. Later we will see how you can find out what other voices are available and also how to change to them.

Speaking C++

We have launched you into the world of computers and robots without really giving you a formal introduction to the C++ language. In this section, we provide more details about the language. What you know about C++ so far is that it is needed to control the robot. The robot commands you type are integrated into C++ by way of the Myro library. C++ comes with several other useful libraries or “classes” that we will try and learn in this course. If you need to access the commands provided by a library, all you have to do is include them.

The libraries themselves are largely made up of sets of functions (they can contain other entities but more on that later). Functions provide the basic building blocks for any program. Typically, a programming language (and C++ is no exception) includes a set of pre-defined functions and a mechanism

for defining additional functions. In the case of C++, it is the function definition construct. You have already seen several examples of function definitions and indeed have written some of your own by now. When defining a new function, you have to give the new function a *name*. Names are a critical component of programming and C++ has rules about what forms a name.

What's in a name?

A name in C++ must begin with either an alphabetic letter (a-z or A-Z) or the underscore (i.e. `_`) and can be followed by any sequence of letters, digits, or underscore letters. For example,

```
iRobot  
myRobot  
jitterBug  
jitterBug2  
my2cents  
my_2_cents
```

are all examples of valid C++ names. Additionally, another important part of the syntax of names is that C++ is *case sensitive*. That is the names `myRobot` and `MyRobot` and `myrobot` are distinct names as far as C++ is concerned. Once you name something a particular way, you have to consistently use that exact case and spelling from then on. Well, so much about the syntax of names, the bigger question you may be asking is *what kinds of things can (or should) be named?*

So far, you have seen that names can be used to represent functions. That is, what a robot does each time you use a function name (like `yoyo`) is specified in the definition of that function. Thus, by giving functions a name you have a way of defining new functions. Names can also be used to represent other things in a program. For instance, you may want to represent a quantity, like speed or time by a name. In fact, you did so in defining the function `yoyo`, which is also shown below:


```
void yoyo(double speed, double waitTime) {
    robot.forward(speed, waitTime);
    robot.backward(speed, waitTime);
}
```

Functions can take parameters that help customize what they do. In the above example, you can issue the following two commands:

```
yoyo(0.8, 2.5);
yoyo(0.3, 1.5);
```

The first command is asking to perform the `yoyo` behavior at speed 0.8 for 2.5 seconds where as the second one is specifying 0.3 and 1.5 for speed and time, respectively. Thus, by parameterizing the function with those two values, you are able to produce similar but varying outcomes. This idea is similar to the idea of mathematical functions: $\text{sine}(x)$ for example, computes the sine of whatever value you supply for x . However, there has to be a way of defining the function in the first place that makes it independent of specific parameter values. That is where names come in. In the definition of the function `yoyo` you have named two parameters (the order you list them is important): `speed` and `waitTime`. By declaring each to be a `double` parameter, you have said that they are expected to be fractional quantities. (All C++ parameters must have a declared *type*, such as `double`, which specifies the allowed values of the parameter.) Then you have used those names to specify the behavior that makes up that function. That is the commands `forward`, and `backward` use the names `speed` and `waitTime` to specify whatever the speed and wait times are included in the function invocation. Thus, the names `speed` and `waitTime` represent or designate specific values in this C++ program.

Names in C++ can represent functions as well as values. What names you use is entirely up to you. It is a good idea to pick names that are easy to read, type, and also appropriately designate the entity they represent. What name you pick to designate a function or value in your program is very important, for you. For example, it would make sense if you named a function `turnRight` so that when invoked, the robot turned right. It would not make any sense if the

robot actually turned left instead, or worse yet, did the equivalent of the yoyo dance. But maintaining this kind of semantic consistency is entirely up to you.

Values

In the last section we saw that names can designate functions as well as values. While the importance of naming functions may be obvious to you by now, designating values by names is an even more important feature of programming. By naming values, we can create names that represent specific values, like the speed of a robot, or the average high temperature in the month of December on top of the Materhorn in Switzerland, or the current value of the Dow Jones Stock Index, or the name of your robot, etc. Names that designate values are also called *variables*. C++ provides a simple mechanism for defining variables, specifying their types and giving them initial values:

```
double speed = 0.75;
double aveHighTemp = 37.0;
double DowIndex = 12548.30;
string myFavoriteRobot = "C3PO";
```

These examples illustrate two of the many kinds of values in C++, namely *numbers* and *strings* (anything enclosed in double-quotes, "). The above are examples of *variable definitions* in C++. This is the basic syntax of a variable definition:

```
<type> <variable name> = <expression>;
```

Variables are so called because their values can be changed, which is accomplished with an assignment statement. For example,

```
speed = 0.9;
```

changes the value of `speed` to 0.9. (We do not mention the variable's type, `double`, because it was already declared in the definition of `speed`. In C++

the type of a variable cannot be changed once it is declared.) The exact syntax of an assignment statement is given below:

```
<variable name> = <expression>;
```

You should read the above statement as: *Let the variable named by <variable name> be assigned the value that is the result of calculating the expression <expression>.* So what is an <expression>? Here are some examples:

```
5
5 + 3
3 * 4
3.2 + 4.7
10 / 2
```

The simplest expression you can type is a number (as shown above). A number evaluates to itself. That is, a 5 is a 5, as it should be! And $5 + 3$ is 8. As you can see, addition (+), subtraction (-), multiplication (*), and division (/) can be used on numbers to form expressions that involve numbers.

You may have also noticed that numbers can be written as whole numbers (3, 5, 10, 1655673, etc) or with decimal points (3.2, 0.5, etc) in them. C++ (and most computer languages) distinguishes between them. Whole numbers are called *integers* and those with decimal points in them are called *floating point* numbers. In C++ the most commonly used floating point numbers are called `double`s, as you have seen. While the arithmetic operations are defined on both kinds of numbers, there are some differences you should be aware of. Look at the outputs produced by the following examples (suitably embedded in a main program):

```
cout << 10.0/3.0 << endl;
```

```
3.3333333333333335
```

```
cout << 10/3 << endl;
```

```
3
```

```
cout << 1/2 << endl;
```

```
0
```

```
cout << 1.0/2 << endl;
```

```
0.5
```

When you divide a floating point number by another floating point number, you get a floating point result. However, when you divide an integer by another integer, you get an integer result. Thus, in the examples above, you get the result 3.3333333333333335 when you divide 10.0 by 3.0, but you get 3 when you divide 10 by 3. Knowing this, the result of dividing 1 by 2 (see above) is zero (0) should not be surprising. That is, while the division operation looks the same (/), it treats integers differently than floating point values. However, if at least one of the numbers in an arithmetic operation is a floating point number, C++ will give you a floating point result (see last example above). You should keep this in mind. More on numbers later, before we get back to robots, let us quickly introduce you to strings.

Computers came to be called so because they excelled in doing calculations. However, these days, computers are capable of manipulating any kind of entity: text, images, sounds, etc. Text is made of letters or *characters* and strings are simply sequences of characters. C++ requires that strings be written enclosed in quotes ("I am a string"). If you want a string to include quotes, they must be preceded by the escape character \ ("I use the \"escape\" character"). The escape character is also used to include a backslash in a string ("The \\ character is called \"backslash\"") and for several other purposes. Treating a string as a value is a powerful feature of C++. C++ also provides some operations on strings using which you can write some useful string expressions, but to use them you have to put

```
#include <string>
using namespace std;
```

at the beginning of your program. The `#include` directive gets you the definition of the type `string`, and the `using` declaration makes it a little more convenient to use (as will be explained later). Here are some examples:

```
string mySchool = "Bryn Mawr College";
string yourSchool = "Georgia Institute of Technology";
cout << mySchool << endl;
cout << yourSchool << endl;
cout << mySchool << " " << yourSchool << endl;
cout << yourSchool+mySchool << endl;
cout << yourSchool + ", " + mySchool << endl;
```

This is the output you would get:

```
Bryn Mawr College
Georgia Institute of Technology
Bryn Mawr College Georgia Institute of Technology
Georgia Institute of TechnologyBryn Mawr College
Georgia Institute of Technology, Bryn Mawr College
```

The operation `+` is defined on strings and it results in concatenating the two strings. The output command `cout <<` is followed by one or more C++ expressions, separated by `<<`. It evaluates all the expressions and prints out the results on the screen. As you have also seen before, this is a convenient way to print out results or messages from your program.

A Calculating Program

Ok, set your robot aside for just a few more minutes. You have now also learned enough C++ to write programs that perform simple, yet interesting, calculations. Here is a simple problem:

On January 1, 2008 the population of the world was estimated at approximately 6.650 billion people. It is predicted that at current rates of population growth, we will have over 9 billion people by the year 2050. A gross estimate of population growth puts the annual increase at +1.14% (it has been as high as +2.2% in the past). Given this data, can you estimate by how much the world's population will increase in this year (2008)? Also, by how much will it increase each day?

In order to answer the questions, all you have to do is compute 1.14% of 6.650 billion to get the increase in population this year. If you divide that number by 366 (the number of days in 2008) you will get average daily increase. You can just use a calculator to do these simple calculations, as shown below:

$$6650000000 \times 1.14 / 100 = 75810000.0$$

$$75810000 / 365 = 207131.1475409836$$

That is, in this year there will be an increase of 75.81 million in the world's population which implies an average daily increase of over 207 thousand people). So now you know the answer!

You can also use C++ to do it, so let us try and write a program to do the above calculations. A program to do the calculation is obviously going to be a bit of overkill. Why do all the extra work when we already know the answer? Small steps are needed to get to higher places. So let's indulge and see how you would write a C++ program to do this. Below, we give you one version:

```
/*
 * File: worldPop.cpp
 * Purpose:
 *     Estimate the world population growth in a year and
 *     also per day.
 *     Given that on January 1, 2008 the world's population
 *     was estimated at 6,650,000,000 and the estimated
 *     growth is at the rate of +1.14%
 */

#include <iostream>
using namespace std;

int main() {
    double population = 6650000000.0;
    double growthRate = 1.14 / 100.0;

    double growthInOneYear = population * growthRate;
    double growthInADay = growthInOneYear / 365;
```

```

cout << "World population on January 1, 2008 is "
    << population << endl;
cout << "By Jan. 1, 2009, it will grow by "
    << growthInOneYear << endl;
cout << "An average daily increase of " << growthInADay
    << endl;

return 0;
}

```

The program follows the same structure and conventions we discussed above, but there are a couple things worth mentioning. Since this program does not control the robot, we omitted `#include "Myro.h"`. However, you can see that we have added these lines:

```

#include <iostream>
using namespace std;

```

These allow us to use `cout <<` for output. We didn't need these lines in our other programs because they were included as part of `Myro.h`. We have defined variables with names `population`, and `growthRate` to designate the values given in the problem description. Although `population` is not a fractional quantity, we declared it `double`, which has to be used for large numbers (above about 2 billion). We also defined the variables `growthInOneYear` and `growthInADay` and used them to designate the results of the calculations. First, in the `main` program we define the values given, followed by

The Energy Problem



The root cause of world energy problems is growing world population and energy consumption per capita.

How many people can the earth support? Most experts estimate the limit for long-term sustainability to be between 4 and 16 billion.

From: *Science, Policy & The Pursuit of Sustainability*, Edited by Bent, Orr, and Baker. Illus. by Shetter. Island Press, 2002.

performing the calculation. Finally, we use the output commands to print out the result of the computations. Notice that our main program ends with:

```
return 0;
```

This is the normal way to exit a C++ main program and to signal that it completed normally. The reason you haven't seen it before is that the exit operation was performed by `disconnect()`, which we cannot use in this program.

Do This: Enter the program, and run it (just as you would run your robot programs) and observe the results. Voila! You are now well on your way to also learning the basic techniques in computing! In this simple program, we did not feel the need to define any functions. But this was a trivial program. However, it should serve to convince you that writing programs to do computation is essentially the same as controlling a robot.

Using Input

The program we wrote above uses specific values of the world's population and the rate of growth. Thus, this program solves only one specific problem for the given values. What if we wanted to calculate the results for a different growth rate? Or even a different estimate of the population? What if we wanted to try out the program for varying quantities of both? Such a program would be much more useful and could be used over and over again. Notice that the program begins by assigning specific values to the two variables:

```
double population = 6650000000.0;
double growthRate = 1.14/100.0;
```

One thing you could do is simply modify those two lines to reflect the different values. However, typical programs are much more complicated than this one and it may require a number of different values for solving a problem. When programs get larger, it is not a good idea to modify them for every specific problem instance but it is desirable to make them more useful for all problem instances. One way you can achieve this is by using the *input*

facilities of C++. All computer programs typically take some input, do some computation (or something), and then produce some output. C++ has a simple input command (defined in `<iostream>`) that can be used to rewrite the program above as follows:

```
/* File: worldPop.cpp
 * Purpose:
 *     Estimate the world population growth in a year and
 *     also per day.
 *     Given that on January 1, 2008 the world's population
 *     was estimated at 6,650,000,000 and the estimated
 *     growth is at the rate of +1.14%
 */

#include <iostream>
using namespace std;

int main() {
    double population, growthRate, growthInOneYear,
           growthInADay;
    // print out the preamble
    cout << "This program computes population growth figures."
          << endl;
    // Input the values
    cout << "Enter current world population: ";
    cin >> population;
    cout << "Enter the growth rate: ";
    cin >> growthRate;
    growthRate = growthRate / 100.0;
    // Compute the results
    growthInOneYear = population * growthRate;
    growthInADay = growthInOneYear / 365;
    // output results
    cout << "World population today is " << population
          << endl;
    cout << "In one year, it will grow by " << growthInOneYear
          << endl;
    cout << "An average daily increase of " << growthInADay
          << endl;
    return 0;
}
```

Read the program above carefully. Look at the first line of main:

```
double population, growthRate, growthInOneYear,  
    growthInADay;
```

This is an abbreviation of the following four variable declarations:

```
double population;  
double growthRate;  
double growthInOneYear;  
double growthInADay;
```

This defines the four variables, but does not give them any initial value, since we will be reading their values from input or calculating them from input values. Notice also that we have added additional comments as well as output statements. This improves the overall readability as well as the interaction of this program. Notice also the use of the `cin >>` statements above. The basic pattern for requesting input is shown below:

```
cout << <some prompt string>;  
cin >> <variable name>;
```

When the output command is executed, the computer will print out the string as a prompt. Notice that we have omitted the usual `<< endl`. This means that the output will not go to a new line, and so the input typed by the user will go on the same line as the prompt. The input command (`cin >>`) waits for the user to enter a value compatible with the variable (e.g., a number). The user can enter whatever value in response to the prompt and then hit the `RETURN` or `ENTER` key. That value is then assigned to the variable `<variable name>`. Now, look at the use of the `cin >>` command in the program above. With this modification, we now have a more general program which can be run again and again. Below, we show two sample runs:

```

*worldPop
This program computes population growth figures.
Enter current world population: 6650000000
Enter the growth rate: 1.14
World population today is 6.65e+09
In one year, it will grow by 7.581e+07
An average daily increase of 207699
*worldPop
This program computes population growth figures.
Enter current world population: 6725810000
Enter the growth rate: 2.2
World population today is 6.72581e+09
In one year, it will grow by 1.47968e+08
An average daily increase of 405391
*|

```

Notice how you can re-run the program by just typing the name of the executable file. There are other ways of obtaining input in C++. We will see those a little later. Also notice the curious way some of the numbers are printed: $6.65e+09$ is computerese for the scientific notation 6.65×10^9 , that is, 6,650,000,000. C++ allows you to control the format in which numbers and other things are printed, but we won't worry about that just yet.

Robot Brains

Writing programs to control your robot is therefore no different from writing a program to perform a computation. They both follow the same basic structure. The only difference is that all robot programs you will write will make use of the Myro library. There will be several robot programs that will require you to obtain input from the user (see exercises below). You can then make use of the `cin >>` command as described above.

One characteristic that will distinguish robot programs from those that just do computations is in the amount of time it will take to run a program. Typically, a program that only performs some computation will terminate as soon as the computation is completed. However, it will be the case that most of the time your robot program will require the robot to perform some task over and over again. Here then, is an interesting question to ask:

Question How much time would it take for a vacuuming robot to vacuum a 16ft X 12ft room?

Seemingly trivial question but if you think about it a little more, you may reveal some deeper issues. If the room does not have any obstacles in it (i.e. an empty room), the robot may plan to vacuum the room by starting from one corner and then going along the entire length of the long wall, then turning around slightly away from the wall, and traveling to the other end. In this manner, it will ultimately reach the other side of the room in a systematic way and then it could stop when it reaches the last corner. This is similar to the way one would mow a flat oblong lawn, or even harvest a field of crop, or receive an ice hockey rink using a Zamboni machine. To answer the question posed above all you have to do is calculate the total distance travelled and the average speed of the vacuum robot and use the two to compute the estimated time it would take. However, what if the room has furniture and other objects in it?

You might try and modify the approach for vacuuming outlined above but then there would be no guarantee that the floor would be completely vacuumed. You might be tempted to redesign the vacuuming strategy to allow for random movements and then estimate (based on average speed of the robot) that after some generous amount of time, you can be assured that the room would be completely cleaned. It is well known (and we will see this more formally in a later chapter) that random movements over a long period of time do end up providing uniform and almost complete coverage. Inherently this also implies that the same spot may very well end up being vacuumed several times (which is not necessarily a bad thing!). This is similar to the thinking that a herd of sheep, if left grazing on a hill, will result, after a period of time, in a nearly uniform grass height (think of the beautiful hills in Wales).

On the more practical side, iRobot's Roomba robot uses a more advanced strategy (though it is time based) to ensure that it provides complete coverage. A more interesting (and important) question one could ask would be:

Question: How does a vacuuming robot know that it is done cleaning the room?

Most robots are programmed to either detect certain terminating situations or are run based on time. For example, run around for 60 minutes and then stop. Detecting situations is a little difficult and we will return to that in the next chapter.

So far, you have programmed very simple robot behaviors. Each behavior which is defined by a function, when invoked, makes the robot do something for a fixed amount of time. For example, the `yoyo` behavior from the last chapter when invoked as:

```
yoyo(0.5, 1);
```

would cause the robot to do something for about 2 seconds (1 second to go forward and then 1 second to move backward). In general, the time spent carrying out the `yoyo` behavior will depend upon the value of the second parameter supplied to the function. Thus if the invocation was:

```
yoyo(0.5, 5.5);
```

the robot would move for a total of 11 seconds. Similarly, the `dance` behavior defined in the previous chapter will last a total of six seconds. Thus, the total behavior of a robot is directly dependent upon the time it would take to execute all the commands that make up the behavior. Knowing how long a behavior will take can help in pre-programming the total amount of time the overall behavior could last. For example, if you wanted the robot to perform the dance moves for 60 seconds, you can repeat the `dance` behavior ten times. You can do this by simply issuing the `dance` command 10 times. But that gets tedious for us to have to repeat the same commands so many times. Computers are designed to do repetitious tasks. In fact, repetition is one of the key concepts in computing and all programming languages, including C++, provide simple ways to specify repetitions of all kinds.

Doing Repetition in C++

If you wanted to repeat the dance behavior 10 times, all you have to do is:

```
for (int i = 1; i <= 10; i++) {  
    dance();  
}
```

(Notice that, as is often the case, the right curly brace is *not* followed by a semicolon.) This is a new statement in C++: the `for`-statement. It is also called a *loop statement* or simply a *loop*. It is a way of repeating something a fixed number of times. The basic syntax of a `for`-loop in C++ is:

```
for (<initialization>; <continuation>; <incrementation>) {  
    <do something>  
    <do something>  
    ...  
}
```

The loop specification begins with the keyword `for` which is followed by three expressions separated by semicolons and all surrounded by parentheses. This line sets up the number of times the repetition will be repeated. What follows is a set of statements surrounded by braces, which are called a *block* that forms the *body* of the loop (stuff that is repeated).

The first expression is the `<initialization>`, which typically defines an integer variable and gives it an initial value, for example, `int i = 1`. The `<continuation>` specifies a condition, and the loop will continue to repeat so long as the condition is true. For example, in the previous example the continuation condition is `i <= 10`, and so the loop will continue to repeat so long as the value of variable `i` is less than or equal to 10. (“`<=`” is a common way of writing “less than or equal to” in programming languages.) The third expression is the `<incrementation>`, which is a command that is executed at the end of each pass through the loop. In this case it is `i++`, which is a C++ abbreviation for “increment `i`” (“increase the value of `i` by one”). (Since C++ is a successor to a programming language named “C,” you can probably guess

how it got its name!) So our example `for` loop works like this. First it does `int i = 1`, which declares an integer variable called `i` and gives it the initial value 1. Since the continuation condition `i <= 10` is satisfied, it executes the commands in the loop body, which in this case is just `dance()`. When it gets to the end of the loop body, it executes `i++`, which increases the value of `i` to 2. It checks the continuation condition, and since `i` is still less than 10, it executes the loop body a second time. This continues, with the loop body being executed 10 times. However, when it gets to the end of the tenth pass, the `i++` command increases `i` to 11, and so the continuation condition is not satisfied and the repetition stops. Thus `i` (which is called a *loop index variable*) is assigned successive values in the sequence 1, 2, ..., 10, and for each of those values, the statements in the body of the loop are executed.

Do This: Let us try this out on the robot. Modify the robot program from the start of this chapter to include the dance function and then write a main program to use the loop above.

```
/*
 * File: dance.cpp
 * Purpose: A simple dance routine
 */

// First include standard declarations
#include "Myro.h"

// Define the new functions...

void yoyo(double speed, double waitTime) {
    robot.forward(speed, waitTime);
    robot.backward(speed, waitTime);
}

void wiggle(double speed, double waitTime) {
    robot.motors(-speed, speed);
    wait(waitTime);
    robot.motors(speed, -speed);
    wait(waitTime);
    robot.stop();
}
```

```
void dance() {
    yoyo(0.5, 0.5);
    yoyo(0.5, 0.5);
    wiggle(0.5, 1);
    wiggle(0.5, 1);
}

// The main dance program
int main() {
    connect();
    cout << "Running the dance routine..." << endl;
    for (int danceStep = 1, danceStep <= 10, danceStep++) {
        dance();
    }
    cout << "...Done" << endl;
    disconnect();
}
```

Notice that we have used `danceStep` (a more meaningful name than `i`) to represent the loop index variable. When you run this program, the robot should perform the dance routine ten times. Modify the value specified in the `for` command to try out some more steps. If you end up specifying a really large value, remember that for each value of `danceStep` the robot will do something for 6 seconds. Thus, if you specified 100 repetitions, the robot will run for 10 minutes.

In addition to repeating by counting, you can also specify repetition using time. For example, if you wanted the robot (or the computer) to do something for 30 seconds. You can write the following command to specify a repetition based on time:

```
while (timeRemaining(10)) {
    <do something>
    <do something>
    ...
}
```


The above commands will be repeated for 10 seconds. Thus, if you wanted the computer to say “Doh!” for 5 seconds, you can write:

```
while (timeRemaining(5)) {
    speak("Doh!", 0);
}
```

In writing robot programs there will also be times when you just want the robot to keep doing its behaviors forever! While technically by *forever* we do mean eternity in reality what is likely to happen is either it runs out of batteries, or you decide to stop it (by aborting the program). The C++ command to specify this uses a different loop statement, called a `while`-loop that can be written as:

```
while (true) {
    <do something>
    <do something>
    ...
}
```

`true` is also a value in C++ (along with `false`) about which we will learn more a little later. For now, it would suffice for us to say that the above loop is specifying that the body of the loop be executed forever!

Do This: Modify the `dance.cpp` program to use each of the `while`-loops instead of the `for`-loop. In the last case (`while (true)`) remember to abort the program to stop the repetitions (and the robot).

As we mentioned above, repetition is one of the key concepts in computing. For example, we can use repetition to predict the world population in ten years by repeatedly computing the values for each year:

```
for (int year = 1; year <= 10; year++) {
    population = population * (1 + growthRate)]
}
```

That is, repeatedly add the increase in population, based on growth rate, ten times.

Do This: Modify the `worldPop.cpp` program to input the current population, growth rate, and the number of years to project ahead and compute the resulting total population. Run your program on several different values (Google: “world population growth” to get latest numbers). Can you estimate when the world population will become 9 billion?

Summary

This chapter introduced the basic structure of C++ (and robot) programs. We also learned about *names* and *values* in C++. Names can be used to designate functions and values. The latter are also called *variables*. C++ provides several different types of values: integers, floating point numbers, strings, and also boolean values (`true` and `false`). Most values have built-in operations (like addition, subtraction, etc.) that perform calculations on them. C++ provides simple built-in facilities for obtaining input from the user. All of these enable us to write not only robot programs but also programs that perform any kind of computation. Repetition is a central and perhaps the most useful concept in computing. In C++ you can specify repetition using either a `for`-loop or a `while`-loop. The latter are useful in writing general robot brain programs. In later chapters, we will learn how to write more sophisticated robot behaviors.

Myro Review

```
Speak(<something>);
```

The computer converts the text in `<something>` to speech and speaks it out. `<something>` is also simultaneously printed on the screen. Speech generation is done synchronously. That is, anything following the `Speak` command is done only after the entire thing is spoken.

```
Speak(<something>, 0);
```

The computer converts the text in `<something>` to speech and speaks it out.

<something> is also simultaneously printed on the screen. Speech generation is done asynchronously. That is, execution of subsequent commands can be done prior to the text being spoken.

```
timeRemaining(<seconds>)
```

This is used to specify timed repetitions in a while-loop (see below).

C++ Review

Values

Values in C++ include numbers (integers or floating point numbers) and strings. Each type of value can be used in an expression by itself or using a combination of operations defined for that type (for example, +, -, *, /, % for numbers). Strings are considered sequences of characters (or letters).

Names

A name in C++ must begin with either an alphabetic letter (a-z or A-Z) or the underscore (i.e. _) and can be followed by any sequence of letters, digits, or underscore letters.

```
cin >> <variable>;
```

Waits for the user to type a value compatible with the type of the variable. This value becomes the new value of the variable.

```
#include "Myro.h"
<any other includes>
<function definitions>
int main() {
    connect();
    <do something>
    <do something>
    ...

    disconnect();
} // end of main program
```

This is the basic structure of a robot control program in C++. Without the `#include` line and the `connect` and `disconnect` commands, it is the basic structure of all C++ programs.

```
cout << <expression> << <expression> << ... << <expression>;
```

This function prints out the values of the expressions (one or more) in the console window. If any of the expressions are `endl`, the output will go to a new line at that point.

```
<type> <variable name> = <expression>;
```

This is how C++ defines a variable of the specified type (e.g., `double`, `int`, `string`). The value generated by `<expression>` will become the initial value of `<variable name>`.

```
<variable name> = <expression>;
```

This is how C++ assigns values to variables. The value generated by `<expression>` will become the new value of `<variable name>`.

```
<expression> <= <expression>
```

This condition is true if the first expression is less than or equal to the second. Other relations include `>=` (greater or equal), `<`, `>`, `==` (equal) and `!=` (not equal).

```
<variable>++
```

Increments `<variable>`, that is, increases the value of `<variable>` by 1.

Repetition

```
for (<initialization>; <continuation>; <incrementation>) {  
    <do something>  
    <do something>  
    ...  
}
```

```
while (timeRemaining(<seconds>)) {
    <do something>
    <do something>
    ...
}

while (true) {
    <do something>
    <do something>
    ...
}
```

These are different ways of doing repetition in C++. The first version will initialize a variable and repeat the loop body, executing the `<incrementation>` at the end of each repetition, and continue repeating so long as the continuation condition remains true. The second version will carry out the body for `<seconds>` amount of time. `timeRemaining` is a Myro function (see above). The last version specifies an un-ending repetition.

Exercises

1. Write a C++ program to convert a temperature from degrees Celsius to degrees Fahrenheit. Here is a sample interaction with such a program:

```
Enter a temperature in degrees Celsius: 5.0
That is equivalent to 41.0 degrees Fahrenheit.
```

The formula to convert a temperature from Celsius to Fahrenheit is: $C/5 = (F - 32) / 9$, where C is the temperature in degrees Celsius and F is the temperature in degrees Fahrenheit.

2. Write a C++ program to convert a temperature from degrees Fahrenheit to degrees Celsius.

3. Write a program to convert a given amount of money in US dollars to an equivalent amount in Euros. Look up the current exchange rate on the web (see xe.com, for example).

4. Modify the version of the dance program above that uses a `for`-loop to use the following loop:

```
for (int danceStep = 0; danceStep < 10; danceStep++) {  
    dance();  
}
```

5. Run the world population program (any version from the chapter) and when it prompts for input, try entering the following and observe the behavior of the program. Also, given what you have learned in this chapter, try and explain the resulting behavior.

a. Use the values 9000000000, and 1.42 as input values as above. Except, when it asks for various values, enter them in any order. What happens?

b. For any of the values to be input, replace them with a string. For instance enter "Al Gore" when it prompts you for a number. What happens?

6. Rewrite your solution to Exercise 4 from the previous chapter to use the program structure described above.

7. You were introduced to the rules of naming in C++. You may have noticed that we have made extensive use of *mixed case* in naming some entities. For example, `waitTime`. There are several naming conventions used by programmers and that has led to an interesting culture in of itself. Look up the phrase *CamelCase controversy* in your favorite search engine to learn about naming conventions. For an interesting article on this, see The Semicolon Wars (www.americanscientist.org/issues/pub/the-semicolon-wars).

8. Experiment with the `speak` function introduced in this chapter. Try giving it a number to speak (try both integers and floating point numbers). What is the largest integer value that it can speak? What happens when this limit is exceeded? Try to give the `speak` function a list of numbers, or strings, or both.

9. Write a C++ program that sings the ABC song: *ABCD...XYZ. Now I know my ABC's. Next time won't you sing with me?*



4

Sensing From Within

I see dead people.
Cole Sear (played by Haley Joel Osment) in *Sixth Sense*,
M. Night Shyamalan, 1999.

Opposite page: Candle Flame
Photo courtesy of Jon Sullivan (www.pdphoto.org)

Cole Sear in Shyamalan's *Sixth Sense* is not referring to dead bodies lying in front of him (for those who have not seen the movie). The five senses that most humans relate to are: touch, vision, balance, hearing, and taste or smell. In all cases our bodies have special sensory receptors that are placed on various parts of the body to enable sensing. For example the taste receptors are concentrated mostly on the tongue; the touch receptors are most sensitive on hands and the face and least on the back and on limbs although they are present all over the body, etc. Besides the difference in the physiology of each kind of receptors there are also different neuronal pathways and thereby sensing mechanisms built into our bodies. Functionally, we can say that each type of sensory system starts with the receptors which convert the thing they sense into electrochemical signals that are transmitted over neurons. Many of these pathways lead to the cerebral cortex in the brain where they are further *processed* (like, "Whoa, that jalapeno is hot!!"). The perceptual system of an organism refers to the set of sensory receptors, the neuronal pathways, and the processing of perceptual information in the brain. The brain is capable of combining sensory information from different receptors to create richer experiences than those facilitated by the individual receptors.

The perceptual system of any organism includes a set of *external* sensors (also called exteroceptors) and some *internal* sensing mechanisms (*interoceptors* or *proprioception*). Can you touch your belly button in the dark? This is because of proprioception. Your body's sensory system also keeps track of the internal state of your body parts, how they are oriented, etc.

Proprioception

Sensing from within

Get something really delicious to eat, like a cookie, or a piece of chocolate, or candy (whatever you fancy!). Hold it in your right hand, and let your right arm hang naturally on your side. Now close your eyes, real tight, and try to eat the thing you are holding. Piece of cake! (well, whatever you picked to eat :-)

Without fail, you were able to pick it up and bring it to your mouth, right?

Give yourself a *Snickers moment* and enjoy the treat.

Sensing is an essential component of being a robot and every robot comes built with internal as well as external sensors. It is not uncommon, for example, to find sensors that are capable of sensing light, temperature, touch, distance to another object, etc. An example of internal sensing in robots is the measurement of movement relative to the robot's internal frame of reference. Sometimes also called *dead reckoning*, it can be a useful sensing mechanism that you can use to design robot behaviors.

Robots employ electromechanical sensors and there are different types of devices available for sensing the same physical quantity. For example, one common sensor found on many robots is a *proximity* sensor. It detects the distance to an object or an obstacle. Proximity sensors can be made using different technologies: infrared light, sonar, or even laser. Depending upon the type of technology used, their accuracy, performance, as well as cost vary: infrared (IR) is the cheapest, and laser is on the expensive side. Lets us take a look at the perceptual system of your Scribbler robot starting with internal sensors.

Proprioception in the Scribbler

The Scribbler has three useful internal sensory mechanisms: *stall*, *time*, and *battery level*. When your program asks the robot to move it doesn't always imply that the robot is actually physically moving. It could be stuck against a wall, for example. The stall sensor in the Scribbler enables you to detect this. You have already seen how you can use time to control behaviors using the `timeRemaining` and `wait` functions. Also, for most movement commands, you can specify how long you want that movement to take place (for example `forward(1, 2.5)` means full-speed forward for 2.5 seconds). Finally, it is also possible to detect battery power level so that you can detect when it is time to change batteries in the robot.

Time

All computers come built-in with an internal clock. In fact, clocks are so essential to the computers we use today that without them we would not have

computers at all! Your Scribbler robot can use the computer's clock to sense time. It is with the help of this clock that we are able to use time in functions like `timeRemaining`, `wait`, and other movement commands. Just with these facilities it is possible to define interesting automated behaviors.

Do This: Design a robot program for the Scribbler to draw a square (say with sides of 6 inches). To accomplish this, you will have to experiment with the movements of the robot and correlate them with time. The two movements you have to pay attention to are the rate at which the robot moves, when it moves in a straight line; and the degree of turn with respect to time. You can write a function for each of these:

```
void travelStraight(double distance) {
    // Travel in a straight line for distance inches
    ...
}

void degreeTurn(double angle) {
    // Spin a total of angle degrees
}
```

That is, figure out by experimentation on your own robot (the results will vary from robot to robot) as to what the correlation is between the distance and the time for a given type of movement above and then use that to define the two functions above. For example, if a robot (hypothetical case) seems to travel at the rate of 25 inches/minute when you issue the command `translate(1.0)`, then to travel 6 inches you will have to translate for a total of $(6*60)/25$ seconds. Try moving your robot forward for varying amounts for time at the same fixed speed. For example try moving the robot forward at speed 0.5 for 3, 4, 5, 6 seconds. Record the distance travelled by the robot for each of those times. You will notice a lot of variation in the distance even for the same set of commands. You may want to average those. Given this data, you can estimate the average amount of time it takes to travel an inch. You can then define `travelStraight` as follows:

```
void travelStraight(double distance) {
    // set up your robot's speed
    double inchesPerSec = <Insert your robot's value here>;

    // Travel in a straight line for distance inches
    robot.forward(1, distance/inchesPerSec);
}
```

Similarly you can also determine the time required for turning a given number of degrees. Try turning the robot at the same speed for varying amounts of time. Experiment how long it takes the robot to turn 360 degrees, 720 degrees, etc. Again, average the data you collect to get the number of degrees per second. Once you have figured out the details use them to write the `degreeTurn` function. Then use the following main program:

```
int main() {
    // Transcribe a square of sides = 6 inches

    connect();
    for (int side = 1; side <= 4; side++) {
        travelStraight(6.0);
        degreeTurn(90.0);
    }
    speak("Look at the beautiful square I made.")
    disconnect();
}
```

Run this program several times. It is unlikely that you will get a perfect square each time. This has to do with the calculations you performed as well with the variation in the robot's motors. They are not precise. Also, it generally takes more power to move from a still stop than to keep moving. Since you have no way of controlling this, at best you can only approximate this type of behavior. Over time, you will also notice that the error will aggregate. This will become evident in doing the exercise below.

Do This: Building on the ideas from the previous exercise, we could further abstract the robot's drawing behavior so that we can ask it to draw any regular

polygon (given the number of sides and length of each side). Write the function:

```
void drawPolygon (int Sides, double Length) {
    // Draw a regular polygon with Sides number of sides
    // and each side of length Length.
}
```

Then, we can write a regular polygon drawing robot program as follows:

```
int main() {
    // Given the number of sides and the length of each side,
    // draw a regular polygon

    int nSides;           // number of polygon sides
    double sideLength;   // length of polygon side
    connect();

    // First, ask the user for the number of sides and
    // side length
    cout << "Given # of sides and side length I will draw"
         << endl;
    cout << "a polygon for you. " // string continues...
         << "Specify side length in inches." << endl;

    cout << "Enter # of sides in the polygon: ";
    cin >> nSides;
    cout << "Enter the length of each side: ";
    cin >> sideLength;

    // Draw the polygon
    drawPolygon(nSides, sidelength);

    speak("Look! I can draw.");
    disconnect();
}
```

To test the program, first try drawing a square of sides 6 inches as in the previous exercise. Then try a triangle, a pentagon, hexagon, etc. Try a polygon with 30 sides of length 0.5 inches. What happens when you give 1 as

the number of sides? What happens when you give zero (0) as the number of sides?

A Slight Detour: Random Walks

One way you can do interesting things with robot drawings is to inject some randomness in how long the robot does something. C++, and most programming languages, typically provide a library function for generating random numbers. Generating random numbers is an interesting process in itself but we will save that discussion for a later time. Random numbers are very useful in all kinds of computer applications, especially games and in simulating real life phenomena. For example, in estimating how many cars might be entering an already crowded highway in the peak of rush hour? Etc. In order to access the random number generating functions in C++ you have to include `cstdlib`, the C++ standard library:

```
#include <cstdlib>
using namespace std;
```

There are lots of features available in this library but we will restrict ourselves with just two, which are described below:

rand() Returns a random unsigned integer between 0 and `RAND_MAX` (inclusive).

RAND_MAX The maximum random number (at least 32767, but usually more than two billion).

Here are some sample outputs that demonstrate these two features:

```
cout << rand() << endl;
```

```
16807
```

```
cout << rand() << endl;
```

```
282475249
```

```
cout << rand() << endl;
```

```
1622650073
```

```
cout << rand() << endl;
```

```
984943658
```

```
cout << rand() << endl;
```

```
1144108930
```

Generally, we do not need random numbers in such a wide range. For example, we are more likely to need a random integer in the range 1 to 6 or in the range 0 to 9. This is easy to do with the C++ *modulus operator*, `%`, which gives the remainder after division. For example, `5 % 2` returns 1, `12 % 4` returns 0, and `17 % 10` returns 7. The following examples show the computation of random integers in the range 0 to 9 (inclusive):

```
cout << rand() % 10 << endl;
```

```
2
```

```
cout << rand() % 10 << endl;
```

```
4
```

```
cout << rand() % 10 << endl;
```

```
8
```

```
cout << rand() % 10 << endl;
```

```
3
```

```
cout << rand() % 10 << endl;
```

```
9
```

```
cout << rand() % 10 << endl;
```

```
0
```

Sometimes we want a fractional random number in some range. For example, to get a random number in the range 0.0 to 1.0 (inclusive) we can divide the number returned by `rand()` by `RAND_MAX`, but we must be careful because both these numbers are integers, and if we divide a smaller integer by a larger one, we will always get zero. As explained in Chapter 3, if you want to get a

fractional (floating point) value from a division, then at least one of its arguments must be floating point. One simple way to do this is to add 0.0 (a floating point number) to either the dividend or divisor:

```
cout << rand() / (0.0 + RAND_MAX) << endl;
```

```
0.519416
```

```
cout << rand() / (0.0 + RAND_MAX) << endl;
```

```
0.830965
```

Although this is a standard “trick,” which would be understood by most programmers, it is a bit misleading (Why would you bother to add zero to something?). Fortunately C++ has a more explicit and readable way of accomplishing the same thing. We can *cast* (explicitly convert) either the divisor or the dividend to the required type (`double`) in the following way:

```
cout << (double) rand() / RAND_MAX << endl;
```

```
0.0345721
```

```
cout << rand() / (double) RAND_MAX << endl;
```

```
0.0534616
```

As you can see, using the C++ standard library is easy enough, and similar to using the Myro library for robot commands. Given the `rand()` function, it is entirely up to you how you use it. Look at the program below:

```
int main() {
    int userInput;
    connect();

    // generate 10 random polygons
    for (int poly = 1; poly <= 10; poly++) {
        // generate a random polygon and draw it
        cout << "Place a new color in the pen port and then..."
            << endl;
        cout << "Enter any number: ";
        cin >> userInput;
        sides = 3 + rand() % 6; // random int in {3, ..., 8}
        size = 2 + rand() % 5; // random int in {2, ..., 6}
    }
}
```

```
        drawPolygon(sides, size);
    }

    // generate a random walk of 20 steps
    for (int step = 1; step <= 20; step++) {
        travelStraight((double) rand() / RAND_MAX);
        degreeTurn(rand() % 360);
    }
    disconnect();
}
```

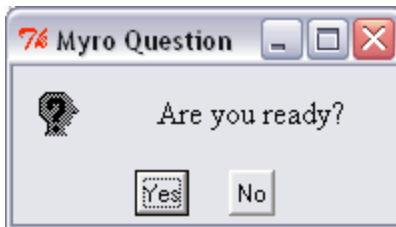
The first loop in the program draws 10 random polygons of sizes ranging from 3 to 8 sides and each side in the range 2 to 6 inches. The second loop carries out a random walk of 20 steps.

Asking Questions?¹

As you can see from above, it is easy to program various kinds of movements into the Scribbler. If there is a pen in the pen port, the Scribbler draws a path. Also in the example above, you can see that we can stop the program temporarily, pretend that we are taking some input and use that as an opportunity to change the pen and then go on. Above, we used the C++ output and input commands `cout <<` and `cin >>` to accomplish this. There is a better way to do this and it uses a function provided in the Myro library:

```
askQuestion("Are you ready?")
```

When this function is executed, a dialog window pops up as shown below:



¹ This section may not reflect the final C++/Myro API.

When you press your mouse on any of the choices (Yes/No), the window disappears and the function returns the name of the key selected by the user as a string. That is, if in the above window you pressed the Yes key, the function will return the value "Yes":

```
cout << askQuestion("Are you ready?") << endl;
```



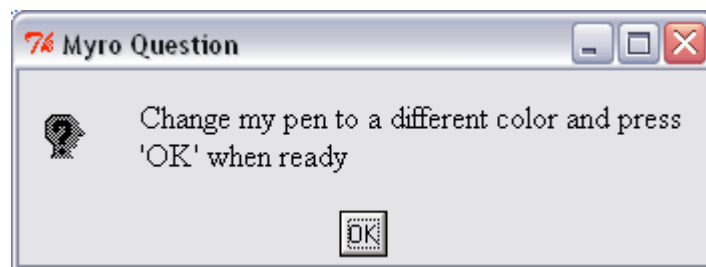
Yes

The `askQuestion` function can be used in the program above as follows:

```
string userInput = askQuestion ("Change my pen to a different"  
    " color and press 'Yes' when ready.");
```


(Notice how to break a long string to make it more readable.) While this is definitely more functional than our previous solution, we can actually do better. For example, what happens when the user presses the No button in the above interaction? One thing you know for sure is that the function will return the string "No". However, the way we are using this function, it really does not matter which key the user presses. `askQuestion` is designed so it can be customized by you so that you can specify how many button choices you want to have in the dialog window as well as what the names of those buttons would be. Here is an illustration of how you would write a better version of the above command:

```
string userInput = askQuestion("Change my pen to a different"  
    " color and press 'OK' when ready", "OK");
```



Now this is certainly better. Notice that the function `askQuestion` can be used with either one parameter or two. If only one parameter is specified, then the default behavior of the function is to offer two button choices: "Yes" and "No". However, using the second parameter you can specify, in a string,² any number of labels that will become the choice buttons. For example,

```
string userInput = askQuestion(  
    "What is your favorite ice cream flavor?",  
    "Vanilla,Chocolate,Mango,Hazelnut,Other");
```



This will be a very handy function to use in many different situations. In the next exercise, try and use this function to become familiar with it.

Do This: Write a Scribbler program of your own that exploits the Scribbler's movements to make random drawings. Make sure you generate drawings with at least three or more colors. Because of random movements, your robot is likely to run into things and get stuck. Help your robot out by picking it up and placing it elsewhere when this happens.

Back to time...

Most programming languages also allow you to access the internal clock to keep track of time, or time elapsed (as in a stop watch), or in any other way you may want to make use of time (as in the case of the `wait`) function. The

² This interface may not represent the final C++/Myro API.

Myro library provides a simple function that can be used to retrieve the current time:

```
cout << currentTime() << endl;
```

```
1169237231.836
```

The value returned by `currentTime` is a number that represents the seconds elapsed since some earlier time, whatever that is. Try issuing the command several times and you will notice that the difference in the values returned by the function represents the real time in seconds. For example:

```
cout << currentTime() << endl;
```

```
1169237351.5580001
```

```
cout << 1169237351.5580001 - 1169237231.836 << endl;
```

```
119.72200012207031
```

That is, 119.722 seconds had elapsed between the two commands above. This provides another way for us to write robot behaviors. So far, we have learned that if you wanted your robot to go forward for 3 seconds, you could either do:

```
robot.forward(1.0, 3.0);
```

OR

```
robot.forward(1.0);  
wait(3.0);  
robot.stop();
```

OR

```
while (timeRemaining(3.0)) {  
    robot.forward(1.0);  
}  
robot.stop();
```

Using the `currentTime` function, there is yet another way to do the same thing:

```
double startTime = currentTime(); // record start time
while (currentTime() - startTime < 3.0) {
    robot.forward(1.0);
}
robot.stop();
```

The above solution uses the internal clock. First, it records the start time. Next it enters the loop which first gets the current time and then checks to see if the difference between the current time and start time is less than 3.0 seconds. If so, the `forward` command is repeated. As soon as the elapsed time gets over 3.0 seconds, the loop terminates. This is another way of using the `while`-loop that you learned in the previous chapter. In the last chapter, you learned that you could write a loop that executed forever as shown below:

```
while (true) {
    <do something>
}
```

The more general form of the `while`-loop is:

```
while (<some condition is true>) {
    <do something>
}
```

That is, you can specify any condition in `<some condition is true>`. The condition is *tested* and if it results in a `true` value, the step(s) specified in `<do something>` is/are performed. The condition is tested again, and so on. In the example above, we use the expression:

```
currentTime() - startTime < 3.0
```

If this condition is true, it implies that the elapsed time since the start is less than 3.0 seconds. If it is false, it implies that more than 3.0 seconds have

elapsed and it results in a `false` value, and the loop stops. Learning about writing such conditions is essential to writing smarter robot programs.

While it may appear that the solution that specified time in the `forward` command itself seemed simple enough (and it is!), you will soon discover that being able to use the internal clock as shown above provides more versatility and functionality in designing robot behaviors. This, for example is how one could program a vacuum cleaning robot to clean a room for 60 minutes:

```
double startTime = currentTime();
while ((currentTime() - startTime)/60.0 < 60.0) {
    cleanRoom();
}
```

You have now seen how to write robot programs that have behaviors or commands that can be repeated a fixed number of times, or forever, or for a certain duration:

```
// do something N times
for (int step = 1; step <= N; step++) {
    do something...
}

// do something forever
while (true) {
    do something...
}

// do something for some duration
while (timeRemaining(duration)) {
    do something...
}

// do something for some duration
double duration = <some time in seconds>;
double startTime = currentTime();
while (currentTime() - startTime < duration) {
    do something...
}
```

All of the above are useful in different situations. Sometimes it becomes a matter of personal preference.

Writing Conditions

Let us spend some time here to learn about conditions you can write in `while`-loops. The first thing to realize is that all conditions result in either of two values: `true` or `false` (interchangeable with 1 and 0, respectively). These are C++ values, just like numbers. You can use them in many ways. Simple conditions can be written using comparison (or relational) operations: `<` (less than), `<=` (less than or equal to), `>` (greater than), `>=` (greater than or equal to), `==` (equal to), and `!=` (not equal to). These operations can be used to compare all kinds of values. Here are some examples:

```
cout << (42 > 23) << endl;
```

```
1
```

```
cout << (42 < 23) << endl;
```

```
0
```

```
cout << (42 == 23) << endl;
```

```
0
```

```
cout << (42 != 23) << endl;
```

```
1
```

```
cout << ((42 + 23) < 100) << endl;
```

```
1
```

```
int a = 10;  
int b = 20;  
int c = 10;  
cout << (a == b) << endl;
```

```
0
```

```
cout << (a == c) << endl;
```

```
1
```



```
cout << (a == a) << endl;
```

```
1
```

```
cout << (true == 1) << endl;
```

```
1
```

```
cout << (false == 1) << endl;
```

```
0
```

The last two examples above also show how the values `true` and `false` are related to 1 and 0. `true` is interchangeable 1 and 0 is interchangeable with `false`. You can form many useful conditions using the comparison operations and all conditions result in either `true` (or 1) or `false` (or 0). You can also compare other values, like strings, using these operations:

```
cout << ("Hello" == "Good Bye") << endl;
```

```
0
```

```
cout << ("Hello" != "Good Bye") << endl;
```

```
1
```

```
cout << ("Elmore" < "Elvis") << endl;
```

```
1
```

```
cout << ("New York" < "Paris") << endl;
```

```
1
```

```
cout << ("A" < "B") << endl;
```

```
1
```

```
cout << ("a" < "A") << endl;
```

```
0
```

Unicode

Text characters have an internal computer coding or representation that enforces lexicographic ordering. This internal encoding is very important in the design of computers and this is what enables all computers and devices like iPhones etc. to exchange information consistently. All language characters in the world have been assigned a standard computer encoding. This is called *Unicode*.

Study the above examples carefully. Two important things to notice are: strings are compared using alphabetical ordering (i.e. lexicographically). Thus "Elmore" is less than "Elvis" since "m" is less than "v" in those strings ("El" being equal in both). That is also why "New York" is less than "Paris" (since "N" is less than "P"). The second important thing to note is that uppercase letters are less than their equivalent lowercase counterparts ("A" is less than "a"). This is by design (see box on the previous page).

Besides relational operations you can build more complex conditional expressions using the *logical operations* (also called *Boolean operations*): `&&` (and), `||` (or), and `!` (not). Here are some examples:

```
cout << ((5 < 7) && (8 > 3)) << endl;
```

```
1
```

```
cout << (! ((5 < 7) && (8 > 3))) << endl;
```

```
0
```

```
cout << ((6 > 7) || (3 > 4)) << endl;
```

```
0
```

```
cout << ((6 > 7) || (3 > 2)) << endl;
```

```
1
```

We can define the meaning of logical operators as follows:

- **<expression-1> && <expression-2>**: Such an expression will result in a value `true` only if both `<expression-1>` and `<expression-2>` are `true`. In all other cases (i.e. if either one or both of `<expression-1>` and `<expression-2>` are `false`) it results in a `false`.
- **<expression-1> || <expression-2>**: Such an expression will result in a value `true` if either `<expression-1>` or `<expression-2>` are `true` or if both are `true`. In all other cases (i.e. if both of

<expression-1> **and** <expression-2> are false) it results in a false.

- **! <expression>**: Such an expression will result in a value `true` if <expression> is false or false if <expression> is true). I.e., it flips or complements the value of expression.

These operators can be combined with relational expressions to form arbitrarily complex conditional expressions. In fact, any decision making in your programs boils down to forming the appropriate conditional expressions. The logical operators were invented by the logician George Boole in the mid 19th century. Boolean algebra, named after Boole, defines some simple, yet important laws that govern the behavior of logical operators. Here are some useful ones:

- `(A || true)` is always true.
- `(! (! A))` is just `A`
- `(A || (B && C))` is the same as `((A || B) && (A || C))`
- `(A && (B || C))` is the same as `((A && B) || (A && C))`
- `(! (A || B))` is the same as `((! A) && (! B))`
- `(! (A && B))` is the same as `((! A) || (! B))`

These identities or properties can help you in simplifying conditional expressions. The conditional expressions can be used to write several useful conditions to control the execution of some program statements. These allow you to write conditional repetitions as:

```
while (<some condition is true>) {
    <do something>
}
```

Now you can see why the following is a way of saying, “do something forever”:

```
while (true) {
    <do something>
}
```

```
}
```

Since the condition is always `true` the statements will be repeated forever. Similarly, in the loop below:

```
while (timeRemaining(duration)) {  
    <do something>  
}
```

As soon as the `duration` is up, the value of the `timeRemaining(duration)` expression will become `false` and the repetition will stop. Controlling the repetitions based on conditions is a powerful idea in computing. We will be using these extensively to control the behaviors of robots.

Sensing Stall

We mentioned in the beginning of this chapter that the Scribbler also has a way of sensing that it is stalled when trying to move. This is done by using the Myro function `getStall`:

`robot.getStall()` Returns `true` if the robot has stalled, `false` otherwise.

You can use this to detect that the robot has stalled and even use it as a condition in a loop to control behavior. For example:

```
while (! robot.getStall()) {  
    <do something>  
}
```

That is, keep doing `<do something>` until the robot has stalled. Thus, you could write a robot behavior that goes forward until it bumps into something, say a wall, and then stops.

```
while (! robot.getStall()) {  
    forward(1.0);  
}  
robot.stop();  
speak("Ouch! I think I bumped into something!");
```

In the above example, as long as the robot is not stalled, `getStall()` will return `false` and hence the robot will keep going forward (since `! false` is `true`). Once it does bump into something, `getStall()` will return `true` and then the robot will stop and speak.

Do This: Write a complete program for the Scribbler to implement the above behavior and then observe the behavior. Show this to some friends who are not in your course. Ask them for their reactions. You will notice that people will tend to ascribe some form of *intelligence* to the robot. That is, your robot is sensing that it is stuck, and when it is, it stops trying to move and even announces that it is stuck by speaking. We will return to this idea of *artificial intelligence* in a later chapter.

Sensing Battery Power Levels

Your Scribbler robot runs on 6 AA batteries. As with any other electronic device, with use, the batteries will ultimately drain and you will need to replace with fresh ones. Myro provides an internal battery-level sensing function, called `getBattery` that returns the current voltage being supplied by the battery. When the battery levels go down, you will get lower and lower voltages causing erratic behavior. The battery voltage levels of your Scribbler will vary between 0.0 and 9.0 volts (0.0 being totally drained). What low means is something you will have to experiment and find out. The best way to do this is to record the battery level when you insert a fresh set of batteries. Then, over time, keep recording the battery levels as you go.

The Scribbler also has some built-in battery-level indicator lights. The red LED on the robot remains lit when the power levels are high (or in the good range). It starts to flash when the battery level runs low. There is also a similar LED on the Fluke dongle. Can you find it? Just wait until the battery levels run low and you will see it flashing.

Disposing Batteries

Make sure that you dispose used batteries properly and responsibly. Batteries may contain hazardous materials

You can use battery-level sensing to define behaviors for robots so that they are carried out only when there is sufficient power available. For example:

```
while ((robot.getBattery() >= 5.0) && timeRemaining(duration))
{
    <do something>
}
```

That is, as long as battery power is above 5.0 and the time limit has not exceeded duration, <do something>.

World Population, revisited

The ability to write conditional expressions also enables us to define more sophisticated computations. Recall the world population projection example from previous chapter. Given the population growth rate and the current population, you can now write a program to predict the year when the world's population will increase to beyond a given number, say 9 billion. All you have to do is write a condition-driven repetition that has the following structure:

```
int year = <current year>;
double population = <current population>;
double growthRate = <rate of growth>;

// repeat as long as the population stays below 9000000000
while (population < 9000000000.0) {
    // compute the population for the next year
    year = year + 1;
    population = population * (1+growthRate);
}

cout << "By the year " << year << " the world's population"
    << endl;
cout << "will have exceeded 9 billion." << endl;
```

That is, add population growth in the next year if the population is below 9 billion. Keep repeating this until it exceeds 9 billion.

Do This: Complete the program above and compute the year when the world's population will exceed 9 billion. To make your program more useful make sure you ask the user to input the values of the year, population, growth rate, etc. In fact, you can even ask the user to enter the population limit so you make use the program for any kinds of predictions (8 billion? 10 billion?). How would you change the program so it prints the population projection for a given year, say 2100?

Summary

In this chapter you have learned about proprioception or internal sensory mechanisms. The Scribbler robot has three internal sensory mechanisms: time, stall, and battery-level. You have learned how to sense these quantities and also how to use them in defining automated robot behaviors. You also learned about random number generation and used it to define unpredictable robot behaviors. Later, we will also learn how to use random numbers to write games and to simulate natural phenomena. Sensing can also be used to define conditional repetitive behaviors using conditional expressions in `while`-loops. You learned how to construct and write different kinds of conditions using *relational* and *logical* operations. These will also become valuable in defining behaviors that use external sensory mechanisms and also enable us to write more explicit decision-making behaviors. We will learn about these in the next chapter.

Myro Review

`randomNumber ()`

Returns a random number in the range 0.0 and 1.0. This is an alternative Myro function that works just like `(double) rand() / RAND_MAX`, which uses the `rand ()` function from the C++ `cstdlib` library (see below).

`askQuestion(MESSAGE-STRING)`

A dialog window with `MESSAGE-STRING` is displayed with choices: "Yes" and "No". Returns "Yes" or "No" depending on what the user selects.

`askQuestion(MESSAGE-STRING, LIST-OF-OPTIONS)`³

A dialog window with `MESSAGE-STRING` is displayed with choices indicated in `LIST-OF-OPTIONS`, which is string of labels separated by commas. Returns option string depending on what the user selects.

`currentTime()`

The current time, in seconds from an arbitrary starting point in time, many years ago.

`robot.getStall()`

Returns `true` if the robot is stalled when trying to move, `false` otherwise.

`robot.getBattery()`

Returns the current battery power level (in volts). It can be a `double` number between 0.0 and 9.0 with 0.0 indicating no power and 9.0 being the highest. There are also LED power indicators present on the robot. The robot behavior becomes erratic when batteries run low. It is then time to replace all batteries.

C++ Review

`true, false`

These are Boolean or logical values in C++. C++ also defines `true` as 1 and `false` as 0 and they can be used interchangeably.

`<, <=, >, >=, ==, !=`

These are relational operations in C++. They can be used to compare values. See text for details on these operations.

³ This description may not reflect the final C++/Myro API.

`&&`, `||`, `!`

These are logical operations. They can be used to combine any expression that yields Boolean values.

`%`

The modulus operator. The expression `<dividend> % <divisor>` returns the remainder after dividing integer `<dividend>` by integer `<divisor>`.

`rand()`

Returns a random integer between 0 and `RAND_MAX` (inclusive). This function is a part of the `cstdlib` library in C++.

`RAND_MAX`

A constant, which is the maximum random integer that can be returned by the `rand()` function in C++.

Exercises

1. Write a robot program to make your Scribbler draw a five point star. [Hint: Each vertex in the star has an interior angle of 36 degrees.]
2. Experiment with Scribbler movement commands and learn how to make it transcribe a path of any given radius. Write a program to draw a circle of any input diameter.
3. Write a program to draw other shapes: the outline of a house, a stadium, or create art by inserting pens of different colors. Write the program so that the robot stops and asks you for a new color.
4. If you had an open rectangular lawn (with no trees or obstructions in it) you could use a Zanboni like strategy to mow the lawn. Start at one end of the lawn, mow the entire length of it along the longest side, turn around and mow the entire length again, next to the previously mowed area, etc. until you are done. Write a program for your Scribbler to implement this strategy (make the Scribbler draw its path as it goes).

5. Enhance the random drawing program from this chapter to make use of speech. Make the robot, as it is carrying out random movements, to speak out what it is doing. As a result you will have a robot artist that you have created!

6. Rewrite your program from the previous exercise so that the random behavior using each different pen is carried out for 30 seconds.

7. The Myro library also provides a function called, `randomNumber()` that returns a random number in the range 0.0 and 1.0. This is similar to the result of `(double) rand() / RAND_MAX`, which uses the `cstdlib` library function `rand()` that was introduced in this chapter. You can use either based on your own preference. You will have to import the appropriate library depending on the function you choose to use. Experiment with both to convince yourself that these two are equivalent.

8. You need only the function `rand()` to generate random integers in any range. For example, you can get a random integer between 1 and 5 (inclusive) as shown below:

```
int randomValue = 1 + rand() % 5;
```

Given that `rand() % 5` returns an integer between 0 and 4 (inclusive), the above expression will assign a random value between 1 and 5 (inclusive) to `randomValue`. Based on this example, write a new function called `randRange(A, B)` that returns a random integer between A and B (inclusive):

```
int randRange (int A, int B) {  
    // generate a random integer between A..B (inclusive)  
}
```

9. You can also use `rand()` to generate floating point numbers in any required range. For example, the following will generate a random double in the range -1.0 to 1.0 (inclusive):

```
double randomValue = -1.0 + 2.0 * rand() / RAND_MAX;
```

Based on this example, write a new function called `randInterval(A, B)` that returns a random `double` number between `A` and `B`:

```
double randInterval (double A, double B) {  
    // generate a random floating point number in (A, B).  
}
```

10.⁴ What kinds of things can your robot talk about? You have already seen how to make the robot/computer speak a given sentence or phrase. But the robot can also “talk” about other things, like the time or the weather.

One way to get the current time and date is to import another C++ library called `ctime`:

```
#include <ctime>
```

The `time` library provides functions called `time` and `localtime` that work as follows. The invocation `time(NULL)` returns a `time_t` object, which represents the current date and time (usually as the number of seconds from some fixed date many years ago). It can be saved in a variable by:

```
time_t now = time(NULL);
```

The `localtime` function takes a pointer to a `time_t` object and returns a pointer to a structure `tm` with the following `int` fields:

- `tm_year`: years since 1900
- `tm_month`: month (0–11)
- `tm_mday`: day (1–31)
- `tm_hour`: hour (0–23)
- `tm_min`: minute (0–59)
- `tm_sec`: seconds (0–59)

⁴ This exercise uses *structures* and *pointers*, C++ features that have not been discussed so far. Do not attempt this exercise unless you know about them from another source or want to learn about them on your own!

- `tm_wday`: weekday (0–6)
- `tm_yday`: day of the year (0–365)
- `tm_isdst`: whether it is using daylight savings time or not (>0 if DST, 0 if not DST, <0 if unknown)

For example, to print out the current year you could write:

```
cout << "The year is "  
  << localtime(&now)->tm_year + 1900  
  << endl;
```

```
2008
```

There are several other useful functions in the `ctime` library. For example `asctime` takes a pointer to a `tm` structure and returns a string containing a human-readable date and time:

```
cout << "Local time is " << asctime(localtime(&now)) << endl;
```

```
Local time is Tue Nov 11 14:14:52 2008
```

Write a C++ program that speaks out the current date and time.

11. Write a function `printTime` that takes an `int` in the range 0 to 23 representing the hour and an `int` in the range 0 to 59 representing the minute, and that prints out the time in 12-hour format with AM or PM appended. For example:

```
printTime (8, 30);
```

```
8:30 AM
```

```
printTime (16, 30);
```

```
4:30 PM
```

```
printTime (12, 15);
```

```
12:15 PM
```

```
printTime (0, 15);
```

```
12:15 AM
```




5

Sensing The World

I see all obstacles in my way.
From the song *I can see clearly now*,
Johnny Nash, 1972.

Opposite page: The Senses
Photo courtesy of Blogosphere (cultura.blogosfere.it)

In the previous chapter you learned how proprioception: sensing time, stall, and battery-level can be used in writing simple yet interesting robot behaviors. All robots also come equipped with a suite of external sensors (or exteroceptors) that can sense various things in the environment. Sensing makes the robot *aware* of its environment and can be used to define more intelligent behaviors. Sensing is also related to another important concept in computing: *input*. Computers act on different kinds of information: numbers, text, sounds, images, etc. to produce useful applications. Acquiring information to be processed is generally referred to as input. In this chapter we will also see how other forms of input can be acquired for a program to process. First, let us focus on Scribbler's sensors.

Scribbler Sensors

The Scribbler robot can sense the amount of ambient light, the presence (or absence) of obstacles around it, and also take pictures from its camera. Several devices (or sensors) are located on the Scribbler (see picture on the previous page). Here is a short description of these:

Camera: The camera can take a still picture of whatever the robot is currently “seeing”.

Light: There are three light sensors present on the robot. These are located in the three holes present on the front of the robot. These sensors can detect the levels of brightness (or darkness). These can be used to detect variations in ambience light in a room. Also, using the information acquired from the camera, the Scribbler makes available an alternative set of brightness sensors (left, center, and right).

Proximity: There are two sets of these on the Scribbler: IR Sensors (left and right) on the front; and Obstacle Sensors (left, center, and right) on the Fluke dongle. They can be used to detect objects on the front and on its sides.

Getting to know the sensors

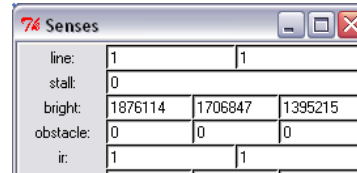
Sensing using the sensors provided in the Scribbler is easy. Once you are familiar with the details of sensor behaviors you will be able to use them in your programs to design interesting creature-like behaviors for your Scribbler. But first, we must spend some time getting to know these sensors; how to access the information reported by them; and what this information looks like. As for the internal sensors, Myro provides several functions that can be used to acquire data from each sensor device. Where multiple sensors are available, you also have the option of obtaining data from all the sensors or selectively from an individual sensor.

Do This: perhaps the best way to get a quick look at the overall behavior of all the sensors is to use the Myro function `senses`:

```
robot.senses();
```

This results in a window (see picture on right) showing all of the sensor values (except the camera) in real time. They are updated every second. You should move the robot around and see how the sensor values change. The window also displays the values of the stall sensor as well as the battery level. The leftmost value in each of the sensor sets (light, IR, obstacle, and bright) is the value of the left sensor, followed by center (if present), and then the right.


Scribbler Sensors



7% Senses			
line:	1		1
stall:	0		
bright:	1876114	1706847	1395215
obstacle:	0	0	0
ir:	1		1

The Camera

The camera can be used to take pictures of the robot's current view. As you can see, the camera is located on the Fluke dongle. The view of the image taken by the camera will depend on the orientation of the robot



(and the dongle). To take pictures from the camera you can use the `takePicture` command:

```
takePicture()  
takePicture("color")  
takePicture("gray")
```

Takes a picture and returns a picture object. By default, when no parameters are specified, the picture is in color. Using the `"gray"` option, you can get a grayscale picture. Example:

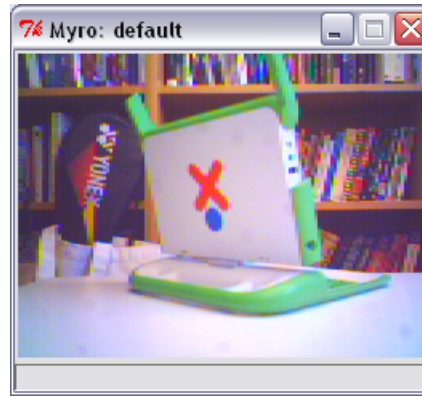
```
picture p = robot.takePicture();  
show(p);
```

Alternately, you can also do:

```
show(robot.takePicture());
```

Once you take a picture from the camera, you can do many things with it. For example, you may want to see if there is a laptop computer present in the picture. Image processing is a vast subfield of computer science and has applications in many areas. Understanding an image is quite complex but something we do quite naturally. For example, in the picture above, we have no problem locating the laptop, the bookcase in the background, and even a case for a badminton racket (or, if you prefer racquet). The camera on the Scribbler is its most complex sensory device that will require a lot of computational effort and energy to use it in designing behaviors. In the simplest case, the camera can serve as your remote "eyes" on the robot. We may not have mentioned this earlier, but the range of the Bluetooth wireless on the robot is 100 meters. In Chapter 9 we will learn several ways of using the pictures. For now, if you take a picture from the camera and would like to save it for later use, you use the Myro command, `savePicture`, as in:

```
savePicture(p, "office-scene.jpg");
```



The file `office-scene.jpg` will be saved in the same folder as your C++ program. You can also use `savePicture` to save a series of pictures from the camera and turn it into an animated “movie” (an animated gif image). This is illustrated in the example below.

Do This: First try out all the commands for taking and saving pictures. Make sure that you are comfortable using them. Try taking some grayscale pictures as well. Suppose your robot has ventured into a place where you cannot see it but it is still in communication range with your computer. You would like to be able to look around to see where it is. Perhaps it is in a new room. You can ask the robot to turn around and take several pictures and show you around. You can do this using a combination of `rotate` and `takePicture` commands as shown below:

```
while (timeRemaining(30)) {
    show(robot.takePicture());
    robot.turnLeft(0.5, 0.2);
}
```

That is, take a picture and then turn for 0.2 seconds, repeating the two steps for 30 seconds. If you watch the picture window that pops up, you will see successive pictures of the robot’s views. Try this a few times and see if you can count how many different images you are able to see. Next, change the `takePicture` command to take grayscale images. Can you count how many images it took this time? There is of course an easier way to do this:

```
int N = 0;
while (timeRemaining(30)) {
    show(robot.takePicture());
    robot.turnLeft(0.5, 0.2);
    N = N + 1;
}
cout << N << endl;
```

Now it will output the number of images it takes. You will notice that it is able to take many more grayscale images than color ones. This is because color images have a lot more information in them than grayscale images (see

text on right). A 256x192 color image requires 256x192x3 (= 147, 456) bytes of data whereas a grayscale image requires only 256x192 (= 49,152) bytes. The more data you have to transfer from the robot to the computer, the longer it takes.

You can also save an animated GIF of the images generated by the robot by using the `savePicture` command by accumulating a series of images in a vector. This is shown below:

```
vector<picture> Pics;
while (timeRemaining(30)) {
    picture pic =
        robot.takePicture();
    show(pic);
    Pics.push_back(pic);
    robot.turnLeft(0.5, 0.2);
}
savePicture(Pics,
    "office-movie.gif");
```

First we create an empty vector called, `Pics`.

Then we append each successive picture taken by the camera onto the vector. This is accomplished by `Pics.push_back(pic)`, which inserts `pic` at the back (end) of `Pics`. Once all the images are accumulated, we use `savePicture` to store the entire set as an animated GIF. You will be able to view the animated GIF inside any web browser. Just load the file into a web browser and it will play all the images as a movie.

There are many more interesting ways that one can use images from the camera. In Chapter 9 we will explore images in more detail. For now, let us take a look at Scribbler's other sensors.

Pixels

Each image is made up of several tiny picture elements or *pixels*. In a color image, each pixel contains color information which is made up of the amount of red, green, and blue (RGB). Each of these values is in the range 0..255 and hence it takes 3 bytes or 24-bits to store the information contained in a single pixel. A pixel that is colored pure red will have the RGB values (255, 0, 0).


A grayscale image, on the other hand only contains the level of gray in a pixel which can be represented in a single byte (or 8-bits) as a number ranging from 0..255 (where 0 is black and 255 is white).

Light Sensing

The following function is available to obtain values of light sensors:

getLight(<POSITION>) Returns the current value in the <POSITION> light sensor. <POSITION> can either be one of "left", "center", "right" or one of the numbers 0, 1, 2. The positions 0, 1, and 2 correspond to the left, center, and right sensors.

Examples:



```
cout << robot.getLight("left") << endl;
```

```
135
```

```
cout << robot.getLight(0) << endl;
```

```
135
```

```
cout << robot.getLight("center") << endl;
```

```
3716
```

```
cout << robot.getLight(1) << endl;
```

```
3716
```

```
cout << robot.getLight("right") << endl;
```

```
75
```

```
cout << robot.getLight(2) << endl;
```

```
75
```

The values being reported by these sensors can be in the range [0..5000] where low values imply bright light and high values imply darkness. The above values were taken in ambient light with one finger completely covering the center sensor. Thus, the darker it is, the higher the value reported. In a way, you could even call it a darkness sensor. Later, we will see how we can easily transform these values in many different ways to affect robot behaviors.

It would be a good idea to use the `senses` function to play around with the light sensors and observe their values. Try to move the robot around to see how the values change. Turn off the lights in the room, or cover the sensors with your fingers, etc.

The camera present on the Fluke dongle can also be used as a kind of brightness sensor. This is done by averaging the brightness values in different zones of the camera image. In a way, you can think of it as a virtual sensor. That is, it doesn't physically exist but is embedded in the functionality of the camera. The function `getBright` is similar to `getLight` in how it can be used to obtain brightness values:

`getBright (<POSITION>)` Returns the current value in the `<POSITION>` light sensor. `<POSITION>` can either be one of "left", "center", "right" or one of the numbers 0, 1, 2. The positions 0, 1, and 2 correspond to the left, center, and right sensors. Examples:

```
cout << robot.getBright("left") << endl;
```

```
2124047
```

```
cout << robot.getBright(0) << endl;
```

```
2124047
```

```
cout << robot.getBright("center") << endl;
```

```
1819625
```

```
cout << robot.getBright(1) << endl;
```


```
1819625
```

```
cout << robot.getBright("right")  
<< endl;
```

```
1471890
```

```
cout << robot.getBright(2) << endl;
```

```
1471890
```




The above values are from the camera image of the Firefox poster (see picture above). The values being reported by these sensors can vary depending on the view of the camera and the resulting brightness levels of the image. But you will notice that higher values imply bright segments and lower values imply darkness. For example, here is another set of values based on the image shown on this page.

```
cout << robot.getBright("left") << ", "  
      << robot.getBright("center") << ", "  
      << robot.getBright("right") << endl;
```

```
1590288, 1736767, 1491282
```

As we can see, a darker image is likely to produce lower brightness values. In the image, the center of the image is brighter than its left or right sections.

It is also important to note the differences in the nature of information being reported by the `getLight` and `getBright` sensors. The first one reports the amount of ambient light being sensed by the robot (including the light above the robot). The second one is an average of the brightness obtained from the image seen from the camera. These can be used in many different ways as we will see later.



Do This: The program shown below uses a normalization function to normalize light sensor values in the range $[0.0..1.0]$ relative to the values of ambient light. Then, the normalized left and right light sensor values are used to drive the left and right motors of the robot.

```
#include "Myro.h"

double Ambient;
// This function normalizes light sensor values to 0.0..1.0
double normalize (int v) {
    if (v > Ambient) {
        v = Ambient;
    }

    return 1.0 - v/Ambient;
}

int main() {
    connect();


    // record average ambient light values
    Ambient = (robot.getLight("left") +
               robot.getLight("center") +
               robot.getLight("right")) / 3.0;

    // Run the robot for 60 seconds
    while (timeRemaining(60)) {
        int L = robot.getLight("left");
        int R = robot.getLight("right");
        // motors run proportional to light
        robot.motors(normalize(L), normalize(R));
    }
    disconnect();
}
```

Run the above program on your Scribbler robot and observe its behavior. You will need a flashlight to affect better reactions. When the program is running, try to shine the flashlight on one of the light sensors (left or right). Observe the behavior. Do you think the robot is behaving like an insect? Which one? Study the program above carefully. We will also return to the idea of making robots behave like insects in the next chapter.

Proximity Sensing

The Scribbler has two sets of proximity detectors. There are two infrared (IR) sensors on the front of the robot and there are three additional IR obstacle sensors on the Fluke dongle. The following function is available to obtain values of the front IR sensors:



getIR(<POSITION>) Returns the current value in the <POSITION> IR sensor. <POSITION> can either be one of "left" or "right" or one of the numbers 0, 1. The positions 0 and 1 correspond to the left, center, and right sensors.

Examples:

```
cout << robot.getIR("left") << endl;
```

```
1
```

```
cout << robot.getIR(0) << endl;
```

```
1
```

```
cout << robot.getIR("right") << endl;
```

```
0
```

```
cout << robot.getIR(1) << endl;
```

```
0
```


IR sensors return either a 1 or a 0. A value of 1 implies that there is nothing in close proximity of the front of that sensor and a 0 implies that there is something right in front of it. These sensors can be used to detect the presence or absence of obstacles in front of the robot. The left and right IR sensors are places far enough apart that they can be used to detect individual obstacles on either side.

Do This: Run the `senses` function and observe the values of the IR sensors. Place various objects in front of the robot and look at the values of the IR

proximity sensors. Take your notebook and place it in front of the robot about two feet away. Slowly move the notebook closer to the robot. Notice how the value of the IR sensor changes from a 1 to a 0 and then move the notebook away again. Can you figure out how far (near) the obstacle should be before it is detected (or cleared)? Try moving the notebook from side to side. Again notice the values of the IR sensors.

The Fluke dongle has an additional set of obstacle sensors on it. These are also IR sensors but behave very differently in terms of the kinds of values they report. The following function is available to obtain values of the obstacle IR sensors:

getObstacle (<POSITION>) Returns the current value in the <POSITION> IR sensor. <POSITION> can either be one of "left", "center", or "right" or one of the numbers 0, 1, or 2. The positions 0, 1, and 2 correspond to the left, center, and right sensors. Examples:



```
cout << robot.getObstacle("left") << endl;
```

```
1703
```

```
cout << robot.getObstacle(0) << endl;
```

```
1703
```

```
cout << robot.getObstacle("center") << endl;
```

```
1128
```

```
cout << robot.getObstacle(1) << endl;
```

```
1128
```

```
cout << robot.getObstacle("right") << endl;
```

```
142
```

```
cout << robot.getObstacle(2) << endl;
```

```
142
```

The values reported by these sensors range from 0 to 7000. A 0 implies there is nothing in front of the sensor where as a high number implies the presence of an object. The sensors on the sides can be used to detect the presence (or absence of walls on the sides).

Do This: Modify your program that uses the game pad controller (`gamepad()`) so that it also issues the `senses` command to get the real time sensor display. Place your Scribbler on the floor, turn it on, and run your new driver program. Our objective here is to really "get into the robot's mind" and drive it around without ever looking at the robot. Also resist the temptation to take a picture. You can use the information displayed by the sensors to navigate the robot. Try driving it to a dark spot, or the brightest spot in the room. Try driving it so it never hits any objects. Can you detect when it hits something? If it does get stuck, try to maneuver it out of the jam! This exercise will give you a pretty good idea of what the robot senses, how it can use its sensors, and to the range of behaviors it may be capable of. You will find this exercise a little hard to carry out, but it will give you a good idea as to what should go into the *brains* of such robots when you actually try to design them. We will try and revisit this scenario as we build various robot programs.

Also do this: Try out the program below. It is very similar to the program above that used the normalized light sensors.

```
#include "Myro.h"
int main() {
    connect();
    // Run the robot for 60 seconds
    while (timeRemaining(60)) {
        int L = robot.getIR("left");
        int R = robot.getIR("right");
        // motors run proportional to IR values
        robot.motors(R, L);
    }
    disconnect();
}
```

Since the IR sensors report 0 or 1 values, you do not need to normalize them. Also notice that we are putting the left sensor value (L) into the right motor and the right sensor value (R) into the left motor. Run the program and observe the robot's behavior. Keep a notebook handy and try to place it in front of the robot. Also place it slightly on the left or on the right. What happens? Can you summarize what the robot is doing? What happens when you switch the R and L values to the motors?

You can see how simple programs like the ones we have seen above can result in interesting automated control strategies for robots. You can also define completely automated behaviors or even a combination of manual and automated behaviors for robots. In the next chapter we will explore several robot behaviors. First, it is time to learn about vectors in C++.

Vectors and Lists in C++

You have seen above that we used vectors to accumulate a series of pictures from the camera to generate an animated GIF. Vectors are a very useful way of collecting a bunch of information and C++ provides a whole host of useful operations and functions that enable manipulation of vectors. In C++, a vector is a sequence of objects of the same type. The objects could be anything: numbers, letters, strings, images, etc. To use vectors, you need to include their header file:

```
#include <vector>
using namespace std;
```

Then an empty vector of integers called N can be defined:

```
vector<int> N;
cout << N << endl;
```

```
{ }
```

(The output of vectors is not built into C++, but is defined in the C++/Myro API.) An empty vector does not contain anything initially, but things can be added to it later. For example,

```
<VECTOR>.push_back(<VALUE>);
```

pushes <VALUE> onto the back (end) of <VECTOR>:

```
N.push_back(7);
N.push_back(14);
cout << N << endl;
```

```
{7, 14}
```

A more convenient way to initialize a vector is:

```
int N_init[] = {7, 14, 17, 20, 27};
vector<int> N (N_init, N_init + 5);
cout << N << endl;
```

```
{7, 14, 17, 20, 27}
```

The first line initializes a C++ “array” (a feature we haven’t discussed) to the five specified numbers. The second line initializes the vector `N` to the values of the array between `N_init` (the location of the first element) and `N_init + 5` (the location after the last). This may seem a bit mysterious (and it is), but it is not too hard to get used to. You can have vectors of any type, for example of `string`s or `doubles`:

```
string Cities_init[] =
    {"New York", "Dar es Salaam", "Moscow"};
vector<string> Cities (Cities_init, Cities_init + 3);

double FN_init[] = {3.14159, 2.718, 42};
vector<double> FamousNumbers (FN_init, FN_init + 3);

cout << Cities << endl;
cout << FamousNumbers << endl;
```

```
{New York, Dar es Salaam, Moscow}
{3.14159, 2.718, 42}
```

C++ provides several useful functions that enable manipulation of vectors. Below, we will show some examples using the variables defined above:

```
cout << N.size() << endl;
```

```
5
```

From the above, you can see that the function `size()` takes a vector and returns the size or the number of objects in the vector. An empty vector has zero objects in it. You can also access individual elements in a vector using the indexing operation (as in `FamousNumbers[0]`). The first element in a list has index 0 and the last element in a list of n elements will have an index $n-1$:

```
cout << FamousNumbers[0] << endl;
```

```
3.14159
```

You can insert the elements from one vector into another. For example, to insert the elements of `N` at the beginning of `FamousNumbers`:

```
FamousNumbers.insert (FamousNumbers.begin(),
    N.begin(), N.end());
cout << FamousNumbers << endl;
```

```
{7, 14, 17, 20, 27, 3.14159, 2.718, 42}
```

The first argument to `insert` specifies where to insert the elements, namely, at the beginning of `FamousNumbers`. The second and third arguments specify the elements to be inserted, namely those from the beginning up to the end of `N`. You can also insert new elements at the end of a vector:

```
N.insert (N.end(), FamousNumbers.begin(),
    FamousNumbers.end());
cout << N << endl;
```

```
{7, 14, 17, 20, 27, 7, 14, 17, 20, 27, 3, 2, 42}
```

The eight values from `FamousNumbers` are inserted before the end of `N`. However, because `N` is an integer vector, the values from `FamousNumbers` are converted from `double` to `int`.

These operations are summarized in more detail at the end of the chapter. The vector types are examples of standard *sequence containers* provided by C++; different types of sequences provide different operations. As we have seen, vectors permit arbitrary elements to be accessed by the indexing operation (denoted by square brackets). Indexing is efficient because the vector elements are stored in contiguous memory locations, but that means that it is inefficient to insert elements anywhere but at the beginning or end of the vector.

C++ provides several other sequence containers, including the *list*. It is similar to the vector, but is stored differently in memory (in doubly linked, rather than contiguous, locations), and so it provides different, efficient operations. For example, it doesn't allow indexing, but it provides other useful list operations, such as `sort` and `reverse`:

```
int SwankeyZips_init[] = {90210, 33139, 60611, 10036};
list<int> SwankeyZips (SwankeyZips_init, SwankeyZips_init+4);
cout << SwankeyZips << endl;
```

```
[90210, 33139, 60611, 10036]
```

(Again, printing of lists is not built into C++, but is defined in the C++ Myro API.)

```
SwankyZips.sort();
cout << SwankyZips << endl;
```

```
[10036, 33139, 60611, 90210]
```

```
SwankyZips.reverse()
cout << SwankyZips << endl;
```

```
[90210, 60611, 33139, 10036]
```

```
SwankyZips.push_back(19010);
cout << SwankyZips << endl;
```

```
[90210, 60611, 33139, 10036, 19010]
```

`sort` rearranges elements in the list in ascending order. `reverse` reverses the order of elements in the list, and `push_back` appends an element to the back (end) of the list. Some other useful list operations are listed at the end of the chapter.

Both vectors and lists are sequences and hence they can be used to perform repetitions. For example, we can use indexing to get the individual strings in `Cities`:

```
for (int i = 0; i < Cities.size(); i++) {
    cout << Cities[i] << endl;
}
```

```
New York
Dar es Salaam
Moscow
```

However, there is a more general way to iterate through any sequence container, including vectors:

```
vector<string>::const_iterator city;
for (city = Cities.begin(); city != Cities.end(); city++) {
    cout << *city << endl;
}
```

```
New York
Dar es Salaam
Moscow
```


The first line declares `city` to be a *constant iterator* for string vectors; making it `const` means that it can be used to read the vector, but not write into it. The iterator `city` points at consecutive values in the vector `Cities`, from `Cities.begin()` up to `Cities.end()`, and the statements inside the loop are executed once for each value of `city`. The “*” in front of `city` means that we will use the value *pointed to* by the iterator, rather than the iterator itself. (Pointers are discussed in more detail elsewhere.) The same approach can be used for iterating through lists; for example:

```
list<int>::const_iterator zip;
for (zip = SwankeyZipsList.begin();
     zip != SwankeyZipsList.end(); zip++) {
    cout << *zip << " ";
}
```

```
90210 60611 33139 10036 19010
```

Strings are sequences and have properties similar to vectors. That is, the string:

```
string ABC = "ABCDEFGHIJKLMNOPQRSTUVWXYZ";
```

is a sequence of 26 letters. You can write a loop that runs through each individual letter in the string and speaks it out as follows:

```
for (string::const_iterator letter = ABC.begin();
     letter < ABC.end(); letter++) {
    speak( *letter );
}
```

In light of the list operations presented above review some of the sensing examples from earlier in the chapter. We will be using vectors and lists in many examples in the remainder of the text. For now, let us return to the topic of sensing.

Extrasensory Perception?

You have seen many ways of acquiring sensory information using the robot's sensors. In addition to the robot itself, you should be aware that your computer also has several "sensors" or devices to acquire all kinds of data. For example, you have already seen how, using the `cin` input stream, you can input some values into your C++ programs:

```
int N;
cout << "Enter a number: ";
cin >> N;
cout << "You entered " << N << endl;
```

```
Enter a number: 42
You entered 42
```

Indeed, there are other ways you can acquire information into your C++ programs. For example, you can input some data from a file in your folder. In Chapter 1 you also saw how you were able to control your robot using the game pad controller. The game pad was actually plugged into your computer and was acting as an input device. Additionally, your computer is most likely connected to the internet using which you can access many web pages. It is also possible to acquire the content of any web page using the internet. Traditionally, in computer science people refer to this as a process of *input*. Using this view, getting sensory information from the robot is just a form of input. Given that we have at our disposal all of the input facilities provided by the computer, we can just as easily acquire input from any of the modalities and combine them with robot behaviors if we wish. Whether you consider this as *extra sensory perception* or not is a matter of opinion. Regardless, being able to get input from a diverse set of sources can make for some very interesting and useful computer and robot applications.

Game Pad Controllers¹

The game pad controller you used in Chapter 1 is a typical device that provides interaction facilities when playing computer games. These devices have been standardized enough that, just like a computer mouse or a keyboard, you can purchase one from a store and plug it into a USB port of your computer.

Myro provides some very useful input functions that can be used to get input from the game pad controller. Game pads come in all kinds of flavors and configurations with varying numbers of buttons, axes, and other devices on them. In the examples below, we will restrict ourselves to a basic game pad shown in the picture on previous page.



The basic game pad has eight buttons (numbered 1 through 8 in the picture) and an axis controller (see picture on right). The buttons can be pressed or released (on/off) which are represented by 1.0 (for on) and 0.0 (for off). The axis can be pressed in many different orientations represented by a pair of values (for the x-axis and y-axis) that range from -1.0 to 1.0 with [0.0, 0.0] representing no activity on the axis. Two Myro functions are provided to access the values of the buttons and the axis:

```
getGamepad (<device>)
getGamepadNow (<device>)
returns a vector of double values indicating
the status of the specified <device>.
<device> can be "axis" or "button".
```



¹ The game pad controller is not implemented in the current C++/Myro API.

The `getGamepad` function returns only after `<device>` has been used by the user. That is, it waits for the user to press or use that device and then returns the values associated with the device at that instant. `getGamepadNow` does not wait and simply returns the device status right away. Here are some examples:

```
cout << robot.getGamepadNow("axis") << endl;
```

```
{0.0, 0.0}
```

```
cout << robot.getGamepad("axis") << endl;
```

```
{0.0, -1.0}
```

```
cout << robot.getGamepadNow("button") << endl;
```

```
{0.0, 0.0, 0.0, 0.0, 1.0, 1.0, 0.0, 0.0}
```

Game Pad's Axes



Both `getGamepad` and `getGamepadNow` return the same vector of values: axis values are returned as a `double` vector `{x-axis, y-axis}` (see picture on right for orientation) and the button values are returned as a `double` vector of 0.0's and 1.0's. The first value in the list is the status of button#1, followed by 2, 3, and so on. See picture above for button numbering.

Do This: Modify your game pad controller to try out the game pad commands above and observe the values. Here is another way to better understand the operation of the game pad and the game pad functions:

```
while (timeRemaining(30)) {
    cout << robot.getGamepad("button") << endl;
}
```

Try out different button combinations. What happens when you press more than one button? Repeat the above for axis control and observe the values returned (keep the axes diagram handy for orientation purposes).

The game pad controller can be used for all kinds of interactive purposes, especially for robot control as well as in writing computer games (see Chapter X). Let us write a simple game pad based robot controller. Enter the program below and run it.

```
#include "Myro.h"
int main() {
    // A simple game pad based robot controller
    connect();
    while (timeRemaining(30)) {
        vector<double> XY = robot.getGamePadNow("axis");
        robot.motors(XY[0], XY[1]);
    }
    disconnect();
}
```

The program above will run for 30 seconds. In that time it will repeatedly sample the values of the axis controller and since those values are in the range -1.0..1.0, it uses them to drive the motors. When you run the above program observe how the robot moves in response to pressing various parts of the axis. Do the motions of the robot correspond to the directions shown in the game pad picture on previous page? Try changing the command from `motors` to `move` (recall that `move` takes two values: translate and rotate). How does it behave with respect to the axes? Try changing the command to `robot.move(-XY[0], -XY[1])`. Observe the behavior.

As you can see from the simple example above, it is easy to combine input from a game pad to control your robot. Can you expand the program above to behave exactly like the `gamepad` controller function you used in Chapter 1? (See Exercise 5).

The World Wide Web and file I/O

If your computer is connected to the internet, you can also save a web page to disk and use it as input to your program. Web pages are written using markup languages like HTML and so when you save the content of a web page you will get the content with the markups included. In this section we will show

you how to access the content of a simple web page and print it out. Later we will see how you could use the information contained in it to do further processing.

Go to a web browser and take a look at the web page:

`http://www.fi.edu/weather/data/jan07.txt`

This web page is hosted by the Franklin Institute of Philadelphia and contains recorded daily weather data for Philadelphia for January 2007. You can navigate from the above address to other pages on the site to look at daily weather data for other dates (the data goes back to 1872!). You will see something like this:

```
January 2007
Day   Max   Min   Liquid   Snow   Depth
1     57   44   1.8    0     0
2     49   40   0      0     0
3     52   35   0      0     0
...   ...   ...   ...     ...   ...
31    31   22   0      0     0
#days 31
Sum   1414 1005 4.18  1.80  1.10
```

Use your browser to save this file to your computer's disk in the same directory where you put your C++ programs. We will use it learn how to do file input/output in C++.

By now you are familiar with using `cin` and `cout` for input/output to the console (keyboard and monitor). They are both examples of *streams* (the *console-input stream* and the *console-output stream*), and input/output to files is also handled by connecting the files to streams. The required definitions are in a standard library called `<fstream>`, which you must `#include` in your program. For example, to read the file `jan07.txt` you would put this declaration in `main`:

```
ifstream in ("jan07.txt");
```

This defines an *input file stream* called `in`, which is connected to the file `jan07.txt`. You can read from this file just as from `cin`; for example, the following will read and print the first word of the file:

```
string word;
in >> word;
cout << word << endl;
```

```
January
```

The operation `in >>`, like `cin >>`, reads “meaningful units” (such as integers, floating point numbers, and words), so it stops when it encounters white space, and skips the white space to get to the next unit. To avoid this, you can use `get()`, which reads a single character from the stream. For example, the following code prints the contents of the file:

```
char ch; // declare a variable to hold a single character
while (!in.eof()) { // continue while not at end of file
    in.get(ch);
    cout << ch;
}
```

Notice that we used the member function `eof()`, which tells us whether we have reached the end of the file. Can you write a program to compute the average of the maximum temperatures recorded in a file such as `jan07.txt`? How about a program to compare two files to see if they are identical?

Writing to an output file is very similar. Consider:

```
ofstream out ("output-file.txt");
...
out << "Here's some output: " << 2+2 << endl;
...
out.close(); // close the file when you are done with it
```

Between the file declaration and the close command, you can use `out <<` just the way you are used to using `cout <<`. C++ provides many more facilities

for doing input and output, including careful control of formatting; you will learn some of them in Chapter 7.

A little more about C++ functions

Before we move on, it would be good to take a little refresher on writing C++ commands/functions. In Chapter 2 we learned that the basic syntax for defining new commands/functions is:

```
void <FUNCTION NAME> (<PARAMETERS>) {  
    <SOMETHING>  
    ...  
    <SOMETHING>  
}
```

The Myro module provides you with several useful functions (`forward`, `turnRight`, etc.) that enable easy control of your robot's basic behaviors. Additionally, using the syntax above, you learned to combine these basic behaviors into more complex behaviors (like `wiggle`, `yoyo`, etc.). By using parameters you can further customize the behavior of functions by providing different values for the parameters (for example, `forward(1.0)` will move the robot faster than `forward(0.5)`). You should also note a crucial difference between the movement commands like `forward`, `turnLeft`, and commands that provide sensory data like `getLight` or `getStall`, etc. The sensory commands always *return* a value whenever they are issued. That is:

```
cout << robot.getLight("left") << endl;
```

```
221
```

```
cout << robot.getStall() << endl;
```

```
0
```

Commands that return a value when they are invoked are called *functions* since they actually behave much like mathematical functions. None of the movement commands return any value, but they are useful in other ways. For instance, they make the robot do something. In any program you typically

need both kinds of functions: those that do something but do not return anything as a result; and those that do something and return a value. You can already see the utility of having these two kinds of functions from the examples you have seen so far. Functions are an integral and critical part of any program and part of learning to be a good programmer is to learn to recognize abstractions that can then be packaged into individual functions (like `drawPolygon`, or `degreeTurn`), which can be used over and over again.

Writing functions that return values

C++ provides a `return`-statement that you can use inside a function to return the results of a function. For example:

```
int triple(int x) {  
    // Returns x*3  
    return x * 3;  
}
```

The word `int` that begins the function definition tells us the type of value returned by `triple`. (As you have seen, definitions of commands, which do not return a value, begin with the word `void`.) The function above can be used just like the ones you have been using:

```
cout << triple(3) << endl;
```

```
9
```

```
cout << triple(5000) << endl;
```

```
15000
```

The general form of a `return`-statement is:

```
return <expression>;
```

That is, the function in which this statement is encountered will return the value of the `<expression>`. Thus, in the example above, the `return`-statement returns the value of the expression $3*x$, as shown in the example

invocations. By giving different values for the parameter `x`, the function simply triples it. This is the idea we used in normalizing light sensor values in the example earlier where we defined the function `normalize` to take in light sensor values and normalize them to the range 0.0..1.0 relative to the observed ambient light values:

```
// This function normalizes light sensor values to 0.0..1.0
double normalize (int v) {
    if (v > Ambient) {
        v = Ambient;
    }

    return 1.0 - v/Ambient;
}
```

In defining the function above, we are also using a new C++ statement: the `if`-statement. This statement enables simple decision making inside computer programs. The simplest form of the `if`-statement has the following structure:

```
if (<CONDITION>) {
    <do something>
    <do something>
    ...
}
```

That is, if the condition specified by `<CONDITION>` is `true` then whatever is specified in the body of the `if`-statement is carried out. In case the `<condition>` is `false`, all the statements under the `if`-command are skipped over.

Functions can have zero or more `return`-statements. Some of the functions you have written, like `wiggle` do not have any. Technically, when a non-void function does not have any `return`-statement that returns a value, the function returns an *undefined value* (that is, an unpredictable and possibly illegal value). This is bad programming and a cause of unpredictable program behavior.

Functions, as you have seen, can be used to package useful computations and can be used over and over again in many situations. Before we conclude this section, let us give you another example of a function. Recall from Chapter 4 the robot behavior that enables the robot to go forward until it hits a wall. One of the program fragments we used to specify this behavior is shown below:

```
while (! getStall()) {
    robot.forward(1);
}
robot.stop();
```

In the above example, we are using the value returned by `getStall` to help us make the decision to continue going forward or stopping. We were fortunate here that the value returned is directly usable in our decision making. Sometimes, you have to do little interpretation of sensor values to figure out what exactly the robot is sensing. You will see that in the case of light sensors. Even though the above statements are easy to read, we can make them even better, by writing a function called `stuck()` as follows:

```
bool stuck() {
    /* Is the robot stalled?
       Returns true if it is and false otherwise. */

    return robot.getStall() == 1;
}
```

The function above is simple enough, since `getStall` already gives us a usable Boolean value (0/false or 1/true). But now if we were to use `stuck` to write the robot behavior, it would read:

```
while (! stuck()) {
    robot.forward(1);
}
robot.stop();
```

As you can see, it reads much better. Programming is a lot like writing in this sense. As in writing, there are several ways of expressing an idea in words.

Some are better and more readable than others. Some are downright poetic. Similarly, in programming, expressing something in a program can be done in many ways, some better and more readable than others. Programming is not all about functionality, there *can* be poetry in the way you write a program.

Summary

In this chapter you have learned all about obtaining sensory data from the robot's perceptual system to do visual sensing (pictures), light sensing, and proximity sensing. The Scribbler provides a rich set of sensors that can be used to design interesting robot behaviors. You also learned that sensing is equivalent to the basic input operation in a computer. You also learned how to get input from a game pad, the World Wide Web, and from data files. Programs can be written to make creative use of the input modalities available to define robot behaviors, computer games, and even processing data. In the rest of the book you will learn how to write programs that make use of these input modalities in many different ways.

Myro Review

```
robot.getBright(<POSITION>)
```

Returns the current value in the <POSITION> light sensor. <POSITION> can either be one of "left", "center", "right" or one of the numbers 0, 1, 2.

```
robot.getGamepad(<device>)  
robot.getGamepadNow(<device>)
```

Returns a double vector indicating the status of the specified <device>. <device> can be "axis" or "button". The `getGamepad` function waits for an event before returning values. `getGamepadNow` immediately returns the current status of the device.

```
robot.getIR(<POSITION>)
```

Returns the current value in the <POSITION> IR sensor. <POSITION> can either be one of "left" or "right" or one of the numbers 0, 1.

```
robot.getLight (<POSITION>)
```

Returns the current value in the <POSITION> light sensor. <POSITION> can either be one of "left", "center", "right" or one of the numbers 0, 1, 2. The positions 0, 1, and 2 correspond to the left, center, and right sensors.

```
robot.getObstacle (<POSITION>)
```

Returns the current value in the <POSITION> IR sensor. <POSITION> can either be one of "left", "center", "right" or one of the numbers 0, 1, or 2.

```
savePicture (<picture>, <file>);  
savePicture (<picture vector>, <file>);
```

Saves the picture in the file specified. The extension of the file should be ".gif" or ".jpg". If the first parameter is a vector of pictures, the file name should have an extension ".gif" and an animated GIF file is created using the pictures provided.

```
robot.senses ();
```

Displays Scribbler's sensor values in a window. The display is updated every second.

```
show (<picture>);
```

Displays the picture in a window. You can click the left mouse anywhere in the window to display the (x, y) and (r, g, b) values of the point in the window's status bar.

```
robot.takePicture ()  
robot.takePicture ("color")  
robot.takePicture ("gray")
```

Takes a picture and returns a `picture` object. When no parameters are specified, the picture is in color.

```
cout << <vector>  
cout << <list>
```

Print vector <vector> or list <list>.

C++ review

```
if (<CONDITION>) {  
    <statement-1>  
    ...  
    <statement-N>  
}
```

If the condition evaluates to `true`, all the statements are performed. Otherwise, all the statements are skipped.

```
return <expression>;
```

Can be used inside any function to return the result of the function.

```
<TYPE> <NAME> [] = { <VALUES> };
```

Declares a C++ array called `<NAME>` with elements of type `<TYPE>`. The array is initialized to the specified `<VALUES>`.

```
char <NAME>;
```

Declares character variable `<NAME>` (able to hold one character).

Vectors and Lists:

```
#include <vector>  
#include <list>
```

Includes definitions for C++ vectors or lists, which are both examples sequence containers.

```
vector< <TYPE> > <NAME>;
```

Defines an initially empty vector called `<NAME>` with elements of type `<TYPE>`.

```
list< <TYPE> > <NAME>;
```

Defines an initially empty list called `<NAME>` with elements of type `<TYPE>`.

```
vector< <TYPE> > <NAME> (<ARRAY START>, <ARRAY END>);
```

```
list< <TYPE> > <NAME> (<ARRAY START>, <ARRAY END>);
```

Defines a vector or list called `<NAME>` with elements of type `<TYPE>` initialized with values from location `<ARRAY START>` up to, but not including, `<ARRAY END>`. These may be of the form `<ARRAY NAME> + <CONSTANT>`.

```
<seq>.size()
```

Returns the current size (length) of the vector, list, or string `<seq>`.

```
<vector>[i]
```

```
<string>[i]
```

Returns the `i`th element in the `<vector>` or `<string>`. Indexing starts from 0. (Strings are a lot like vectors of characters.)

```
<seq>.push_back(<value>);
```

Appends the `<value>` at the back (end) of vector, list, or string `<seq>`.

```
<list>.push_front(<value>);
```

Appends the `<value>` at the front (beginning) of `<list>`.

```
<list>.sort();
```

Sorts the `<list>` in ascending order.

```
<list>.reverse();
```

Reverses the elements in the list.

```
<seq>.begin()
```

Returns an iterator referring to the first element of the vector, list, or string

```
<seq>.
```

```
<seq>.end()
```

Returns an iterator referring the last element of the vector, list, or string

```
<seq>.
```

```
vector< <TYPE> >::const_iterator <NAME>;
```

```
vector< <TYPE> >::iterator <NAME>;
```

```
list< <TYPE> >::const_iterator <NAME>;
```

```
list< <TYPE> >::iterator <NAME>;
```

```
string::const_iterator <NAME>;
```

```
string::iterator <NAME>;
```

Declares an iterator called `<NAME>`, which can refer to locations in a vector/list with elements of type `<TYPE>` or in a string. A `const` (constant) iterator does not allow the elements of the vector or list to be modified, whereas a non-`const` iterator does.

```
<iterator>++
```

Increments `<iterator>` to refer to the next element of a vector, list, or string.

```
* <iterator>
```

Returns the value stored in the vector, list, or string element referred to by `<iterator>`.

```
<seq>.insert( <it1>, <it2-begin>, <it2-end> );
```

Inserts elements from one vector, list, or string into the vector, list, or string `<seq>`. The new elements are inserted at a location specified by iterator `<it1>`. For example, use `<seq>.begin()` to insert at the beginning of `<seq>` or use `<seq>.end()` to insert at its end. The elements to be inserted are specified by iterators `<it2-begin>` and `<it2-end>`. For example, to insert all the elements from vector, list, or string `<seq2>`, use `<seq2>.begin()` and `<seq2>.end()`.

```
for ( <it> = <seq>.begin(); <it> != <seq>.end(); <it>++) {  
    ... STATEMENTS using <it> or *<it> ...  
}
```

Executes the loop body with iterator `<it>` referring to the elements of vector or list `<seq>` from its beginning up to its end.

Streams:

```
#include <fstream>
```

Includes definitions for C++ streams for file input/output.

```
ifstream <name> (<string>);
```

Declares `<name>` to be an input file stream connected to file name `<string>`.


```
ofstream <name> (<string>);
```

Declares <name> to be an output file stream connected to file name <string>.

```
<ifstream> >> <var>;
```

Reads input value from stream <ifstream> and stores it in variable <var>.

```
<ofstream> << <expression>
```

Outputs values of <expression> to output stream <ofstream>.

```
<stream>.close();
```

Closes (input or output) file stream <stream>, thus completing operation on it. You should close files when you are done with them.

```
<ifstream>.get (<chvar>);
```


Puts the next character from input stream <ifstream> into character variable <chvar>.

```
<ifstream>.eof()
```

Returns `true` if input stream <ifstream> is at the end of the file, and `false` if not.

Exercises

1. The numbers assigned to the variable `FamousNumbers` in this chapter all have names. Can you find them? What other famous numbers do you know?
2. Besides text, the `speak` command can also vocalize numbers. Try `speak(42)` and also `speak(3.1419)`. Try some really large whole numbers, like `speak(4537130980)`. How is it vocalized? Can you find out the limits of numerical vocalization?
3. The Fluke camera returns pictures that are 256x192 (= 49,152) pixels. Typical digital cameras are often characterized by the number of pixels they use for images. For example, 8 megapixels. What is a megapixel? Do some web search to find out the answer.
4. All images taken and saved using the Fluke camera can also be displayed in a web page. In your favorite browser, use the Open File option (under the File menu) and navigate to a saved picture from the Scribbler. Select it and view it in the browser window. Try it with an animated GIF image file.
5. Modify the game pad input program from this chapter to make the axis controller behave so that when the top of the axis is pressed, the robot goes forward and when the bottom of the axis is pressed, it goes backward.
6. Write a program that reads in the names of two files, and then compares the two files character-by-character to see if they are identical.
7. Write a program that reads in a monthly weather report, such as `jan07.txt`, and computes the average maximum temperature and average minimum temperature for the month.



6

Insect-Like Behaviors

*So you gotta let me know
Should I stay or should I go?*

From the song, *Should I stay or should I go*, Mick Jones (The Clash), 1982.

Opposite page: Ladybug

Photo courtesy of Jon Sullivan (www.pdphoto.org)

Designing robot behaviors is a challenging, yet fun process. There isn't a formal methodology or a technique that one can follow. It involves creativity, the ability to recognize the strengths and limitations of the physical robot, the kind of environment the robot will be carrying out the behavior, and of course the knowledge of available paradigms for programming robot behaviors. Creativity is essential to the design of robot behaviors. You have already seen how even a simple robot like the Scribbler can be programmed to carry out a diverse range of behaviors. We have also spent a considerable effort so far in exploring the range of possible functions a robot can perform. Where a robot is placed when it is running can play an important role in exhibiting a programmed behavior successfully. In this chapter, we take a different look at robot behaviors.

Braitenberg Vehicles


In 1984, Valentino Braitenberg wrote a book titled: *Vehicles: Experiments in Synthetic Psychology* (MIT Press). In it, he describes several thought experiments that are centered around the creation of simple vehicles with very simple control mechanisms that exhibit seemingly complex behavior. The purpose of the thought experiments was to illustrate some fundamental insights into the internal structure of animal (and human) brains. Each experiment involves a description of a simple vehicle that is endowed with a small suite of sensors (much like our Scribbler robot) and how the sensors can be connected to the motors of these imaginary vehicles in ways that parallel neurological connections in animals. He shows how the resulting vehicles are capable of complex behaviors which can be described as: fear, aggression, love, logic, free will, etc.

One central theme underlying Braitenberg's experiments is the demonstration of what he calls the *Law of uphill analysis and downhill invention*: It is much more difficult to guess the internal structure of an entity just by observing its behavior than it is to actually create the structure that leads to the behavior. That is, trying to postulate the internal structure purely by observing certain behavior is an uphill (harder) task whereas trying to create an entity that exhibits a certain behavior is a downhill (easy) task. While all of Braitenberg's

vehicles were imaginary and not really designed to be actually fabricated people have found it a fun and intellectually interesting exercise to create them. Personal robots like Scribblers make perfect platforms to do this and in what follows we will describe some of Braitenberg's (and Braitenberg-type) vehicles and design robot behaviors based on them.

Vehicle 1: Alive

The first vehicle Braitenberg describes has one sensor and one motor. The value transmitted by the sensor directly feeds into the motor. If the value being reported by the sensor is a varying quantity (say light), the vehicle will move at a speed proportional to the amount of quantity being detected by the sensor.



A schematic of the vehicle is shown above on the left. In order to design this vehicle using the Scribbler, you can use the center light sensor and connect what it reports directly to both motors of the robot. This is shown on the right. That is, the same light reading is directly controlling both the motors by the same amount. As you have already seen, there are a many different ways to specify motor movement commands to the Scribbler. Suppose the value obtained from the center light sensor is `c`, you can control both motors using this value by using the command:

```
robot.motors(C, C);
```

Alternately, you can also use the `forward` command:

```
robot.forward(C);
```

Now that we know how the internal structure of this vehicle looks, we can start to write a program that will implement it. But, before we get there, we need to sort out a small issue of compatibility: light sensors report `int` values

in the range 0..5000 whereas motors and movement commands take `double` values in the range -1.0 to 1.0. In this example, we are only concerned with movements that range from a complete stop to full speed forward, so the values range from 0.0 to 1.0. We have to computationally normalize, or map the light sensor values in this range. A first attempt at this would be to write a function called `normalize` that operates as follows:

```
double normalize(int v) {
    // normalizes v to range 0.0 to 1.0
}
```

Once we have this function, we can write the behavior for the vehicle as follows:

```
int main() {
    // Braitenberg vehicle#1: Alive

    connect();
    while (true) {
        int L = robot.getLight("center");
        robot.forward(normalize(L));
    }

    disconnect();
}
```

Normalizing Sensor Values

It is time now to think about the task of the `normalize` function. Given a value received from a light sensor, it has to transform it to a proportional value between 0.0 and 1.0 so that the brighter the light, the higher the value (i.e. closer to 1.0). Vehicle#1 moves in proportion to the amount of light it receives. This is a good time to revisit the `senses` function of Myro to look at the values reported by the light sensors. Go ahead and do this.

After examining the values returned by the light sensors you may notice that they report small values (less than 50) for bright light and larger values (as large as 3000) for darkness. In a way, you can say that the light sensor is

really a darkness sensor; the darker it is the higher the values reported by it. The light sensors are capable of reporting values between 0 and 5000. Now, we can certainly calibrate or normalize using these values using the following definition of `normalize`:

```
double normalize(int v) {  
    // Normalize v (in the range 0..5000) to 0..1.0, inversely  
  
    return 1.0 - v/5000.0;  
}
```

That is, we divide the value of the light sensor by its maximum value and then subtract that from 1.0 (for inverse proportionality). Thus a brighter light value, say a value of 35, will get normalized as:

$$1.0 - 35.0/5000.0 = 0.9929$$

If 0.9929 is sent to the motors (as in the above program), the robot would move full speed forward. Let us also compute the speed of the robot when it is in total darkness. When you place a finger on the center sensor, you will get values in the 2000-3000 range. For 3000, the normalization will be:

$$1.0 - 3000.0/5000.0 = 0.40$$

The robot will still be moving, although at nearly half the speed. Most likely, you will be operating the robot in a room where there is sufficient ambient light. You will notice that under ambient daylight conditions, the values reported by the light sensors are in the 150-250 range. Using the above normalization you will get:

$$1.0 - 200.0/5000.0 = 0.9599$$

That is almost full speed ahead. In order to experience the true behavior of the above vehicle, we have to use a normalization scheme that takes into account the ambient light conditions (they will vary from room to room). Further, let us assume that in ambient light conditions, we will watch the robot respond to

a light source that we will control. A flashlight will work nicely. So, to make the robot appropriately sensitive to the flashlight under ambient light conditions you can write a better version of `normalize` as follows:

```
double normalize(int v) {
    if (v > Ambient) {
        v = Ambient;
    }
    return 1.0 - v/Ambient;
}
```

That is, the darkest condition is represented by the ambient light value (`Ambient`) and then normalization is done with respect to that value. You can either set the ambient value by hand, or, a better way is to have the robot sense its ambient light at the time the program is initiated. This is the same version of `normalize` that you saw in the previous chapter. Now you know how we arrived at it. The complete program for `Vehicle#1` is shown below:

```
// Braitenberg Vehicle#1: Alive
#include "Myro.h"

double Ambient;

double normalize(int v) {
    if (v > Ambient) {
        v = Ambient;
    }
    return 1.0 - v/Ambient;
}

int main() {
    // Braitenberg vehicle#1: Alive

    connect();
    Ambient = robot.getLight("center");
    while (true) {
        int L = robot.getLight("center");
        robot.forward(normalize(L));
    }
    disconnect();
}
```


Do This: Implement the program above and observe the robot's behavior. Does it respond as described above?

You may have also noticed by now that the three light sensors are not necessarily closely matched. That is, they do not report exactly the same values under the same conditions. When writing robot programs that use multiple light sensors, it is a good idea to average the values returned by all the light sensors to represent the ambient value. Modify the program above to use the average of all three values as the ambient value. There shouldn't be a noticeable difference in the robot's behavior. However, this is something you may want to use in later programs.

Vehicle 2: Coward and Aggressive

Vehicle 2a: Coward

The next set of vehicles use two sensors. Each sensor directly drives one motor. Thus the speed of the individual motor is directly proportional to the quantity being sensed by it sensor. There are two possible ways to connect the sensors.



In the first case, Vehicle2a, the sensor on each side connects to the motor on the same side. In the other case, Vehicle2b, the connections are interchanged. That is, the left sensor connects to the right motor and the right sensor connects to the left motor. Let us design the control program for Vehicle 2a first:

```
// Braitenberg Vehicle#2a
#include "Myro.h"

double Ambient;

double normalize(int v) {
    if (v > Ambient) {
        v = Ambient;
    }
}
```

```

    return 1.0 - v/Ambient;
}

int main() {
    // Braitenberg vehicle#2a: Coward

    connect();
    Ambient = (robot.getLight("left") +
              robot.getLight("center") +
              robot.getLight("right")) / 3.0;
    while (true) {
        int L = robot.getLight("left");
        int R = robot.getLight("right");
        robot.motors(normalize(L), normalize(R));
    }
    disconnect();
}


```

The structure of the above program is very similar to that of Vehicle1. We have modified the setting of the ambient light value to that of an average of the three light values. Also, we use the `motors` command, to drive the left and right motors proportional to the left and right light sensor values (after normalizing).

Do This: Implement the control program for Vehicle 2a as shown above. Observe the robot's behaviors by shining the flashlight directly in front of the robot and in each of the left and right sensors.

Next, write the control program for Vehicle2b as shown here. This requires one simple change from the program of Vehicle2a: switch the parameters of the `motors` command to reflect the interchanged connections. Again observe the behaviors by shining the flashlight directly ahead of the robot and also a little to each side.

Vehicle 2b: Aggressive



You will notice that the robots behave the same way when the light is placed directly ahead of them: they are both attracted to light and hence move towards the light source. However, Vehicle 2a will move away from the light if the light source is on a side. Since the nearer sensor will get excited more, moving the corresponding motor faster, and thereby turning the robot away. In the case of Vehicle 2b, however, it will always turn towards the light source and move towards it. Braitenberg calls these behaviors *coward* (2a) and *aggressive* (2b).

Controlling Robot Responses

It is often necessary, when designing and testing robot behaviors, to properly set up the robot's environment and the orientation of the robot in it. In simple cases this is easily achieved by first placing the robot in the desired orientation and then loading and executing the program. However prior to the robot's actual behavior, the robot may need to perform some preliminary observations (for example, sensing ambient light), it becomes necessary to re-orient the robot properly before starting the execution of the actual behavior. This can be easily accomplished by including some simple interactive commands in the robot's program. The resulting program structure is shown below:

```
// include myro library
// define all functions here (like, normalize, etc.)

int main() {
    // Description of the behavior...

    // Give user the opportunity to set up the robot
    string OK = askQuestion("Press OK to begin...", "OK");

    // establish connection with the robot
    // set values of ambient conditions

    // Write your robot's behavior commands here
    // disconnect from the robot
}
```

Do This: Modify the programs for Vehicles 1, 2a, and 2b to include the `askQuestion` command above.

We have introduced a few basic programming patterns above that can be used in many robot programming situations. The thing to remember is that, at any point in the execution of a robot's program, you can also program appropriate interjections to perform various experimental or control functions. We will see several other examples later on.

Other Normalizations


All the normalizations of light sensor values shown above were used to normalize the values in the range 0.0..1.0 in direct proportion to the amount of light being sensed. That is, the darker it is, the closer the normalized values are to 0.0 and the brighter it gets, the closer the normalized values get to 1.0. This is just one way that one can relate the quantity being sensed to the amount of speed applied to the robot's motors. You can imagine other relationships. The most obvious of course is an inverse relationship: the darker it is the closer to 1.0 and vice versa. Braitenberg calls this *inhibitory* (as opposed to *excitatory*) relationship: the more of a quantity being sensed, the slower the robot's motors turn. As in Vehicles 2a and 2b above, there is choice of two kinds of connections: straight and crossed. These are shown below (a plus (+) sign next to a connector indicates an excitatory connection and a minus sign (-) represents an inhibitory connection):

Writing the `normalize` function for an inhibitory connection is quite straightforward:

```
double normalize(int v) {
    if (v > Ambient) {
        v = Ambient;
    }

    return v/Ambient;
}
```

Vehicles 3a (Love) and 3b (Explorer)



Braitenberg describes the behavior of the resulting vehicles as *love* (Vehicle 3a) and *explorer* (Vehicle 3b). That is, if you were to observe the behavior of the two vehicles, you are likely to notice that Vehicle 3a will come to rest facing the light source (in its vicinity) whereas vehicle 3b will come to rest turned away from the source and may wander away depending on the presence of other light sources.


In other variations on sensor value normalizations, Braitenberg suggests using non-monotonic mathematical functions. That is, if you look at the excitatory and inhibitory normalizations, they can be described as monotonic: more light, faster motor speed; or more light, slower motor speed. But consider other kinds of relationships for normalizations. Observe the function shown on the next page. That is, the relationship is increasing in proportion to sensory input but only up to a certain point and after that it decreases. Incorporating such relationships in vehicles will lead to more complex behavior (Braitenberg describes them as vehicles having *instincts*). The following defines a normalization function, based on the curve shown:

$$f(x) = e^{-(x-100)^2/1800}$$

A Non-Monotonic Function

The above function is based on the simpler function:

$$f(x) = e^{-x^2}$$



which in the first definition is stretched to span the range 0..200 for values of x with 100 being the point where it reports the maximum value (i.e. 1.0). Mathematically this function is also known as the *bell curve* or a *Gaussian Curve* in general. A bell curve is defined in terms of a *mean* (π) and *standard deviation* (σ) as shown below:

$$f(x) = e^{-(x-\pi)^2/2*\sigma^2}$$

Thus, in the normalization function we are using 100 as mean and 30 as standard deviation. You can easily scale the curve for the range of sensor values you desire using the following `normalize` function.

```
double normalize(int v) {
    double mean = Ambient/2.0;
    double stddev = Ambient/6.0;
    if (v >= Ambient) {
        v = Ambient;
    }
    return exp(-(v - mean)**2 / 2 * (stddev**2));
}
```

`exp(x)` is a C++ function that computes the value e^x . It is available in the C++ `cmath` library. We will delve into the math library in more detail in the next chapter. In order to use the `exp` function as shown above you have to include the `cmath` library:

```
#include <cmath>
```

There are of course several other possibilities that one could try: a step function; or a threshold; and any other mathematical combinations. The key idea is that there is a clear mapping of the range of sensor values to motor values in the range 0.0..1.0.

Robots using these normalizations and other variations are likely to exhibit very interesting and sometimes unpredictable behaviors. Observers unaware of the internal mapping mechanisms will have a hard time describing precisely the robot's behavior and will tend to use anthropomorphic terms (like, *love*, *hate*, *instincts*, etc.) to describe the behavior of robots. This is what an *uphill analysis* means.

Multiple Sensors

Adding several sensors enriches the design space for robot behaviors. As a designer, you now have a choice of different types of mappings: excitatory, inhibitory, or more complex; and connections: straight or crossed. Suddenly

the resulting robot behavior will seem complex. On the Scribbler, for instance, in addition to light sensors, you also have the stall sensor, and the IR sensors. With the exception of light sensors, all of these other sensors are digital or threshold sensors that are either ON or OFF (i.e. they report values that are either 0 or 1 indicating the presence or absence of the thing they are sensing). In a way you can think that the digital sensors are already normalized, but it is still possible to invert the relationship if need be. You can design several interesting behaviors by combining two or more sensors and deciding whether to connect them straight or crossed.

Do This: In your design of vehicles 2a and 2b substitute the obstacle sensor in place of the light sensors. Describe the behavior of the resulting vehicles. Try the same for Vehicles 3a and 3b. Next, combine the behavior of the resulting vehicles with the light sensors. Try out all combinations of connections, as well as inhibitory and excitatory mappings. Which vehicles exhibit the most interesting behaviors?

More Vehicles

Here are descriptions of several vehicles that are in the spirit of Braitenberg's designs and also exhibit interesting behaviors. Using the concepts and programming techniques from above, try to implement these on the Scribbler robot. Once completed, you should invite some friends to observe the behaviors of these creatures and record their reactions.

Timid

Timid is capable of moving forward in a straight line. It has one threshold light sensor, pointing up. When the light sensor detects light, the creature moves forward, otherwise, it stays still. The threshold of the light sensor should be set to ambient light. That way, when the creature can "see" the light, it will move. When it enters a shadow (which can be cast by a hand or another object) it stops. If whatever is casting the shadow is moved, the creature will move again. Therefore, timid is a shadow seeker.

Indecisive

Indecisive is similar to Timid, except, it never stops: its motors are always running, either in forward direction, or in reverse direction, controlled by the threshold light sensor. When the light sensor detects light, it moves forward, otherwise, it moves backwards. When you run this creature, you will notice that it tends to oscillate back and forth at shadow edges. Thus, Indecisive is a shadow edge seeker.

Paranoid

Paranoid is capable of turning. This is accomplished by moving the right motor forward and moving the left motor in reverse direction at the same time. It has a single threshold light sensor. When the sensor detects light, it moves forward. When the sensor enters a shadow, it reverses the direction of its left motor, thus turning right. Soon the sensor will swing around, out of the shadow. When that happens, it resumes its forward motion. Paranoid, is a shadow fearing creature.

This, That, or the Other

The `if`-statement introduced earlier in Chapter 5 is a way of making simple decisions (also called one-way decisions). That is, you can conditionally control the execution of a set of commands based on a single condition. The `if`-statement in C++ is quite versatile and can be used to make two-way or even multi-way decisions. Here is how you would use it to choose among two sets of commands:

```
if (<condition>) {  
    <this>  
} else {  
    <that>  
}
```

That is, if the `<condition>` is true it will do the commands specified in `<this>`. If, however, the `<condition>` is false, it will do `<that>`. Similarly, you can extend the `if`-statement to help specify multiple options:

```
if (<condition-1>) {
    <this>
} else if (<condition-2>) {
    <that>
} else if (<condition-3>) {
    <something else>
...
} else {
    <other>
}
```

Thus, depending upon whichever condition is true, the corresponding `<this>`, `<that>`, or `<something else>` will be carried out. If all else fails, the `<other>` will be carried out. This structure is called an `if-else-if` *ladder* or *staircase*.

Simple Reactive Behaviors

Using the three light sensors the robot can detect varying light conditions in its environment. Let us write a robot program that makes it detect and orient towards bright light. Recall from Chapter 5 that light sensors report low values in bright light conditions and high values in low light. To accomplish this task, we only need to look at the values reported by left and right light sensors. The following describes the robot's behavior:

```
do for a given amount of time
    if left light is brighter than right light
        turn left
    else
        turn right
```

Thus, by making use of the `if-else` statement, we can refine the above into the following:

```
while (timeRemaining(30)) {
    if (left light is brighter than right light) {
        robot.turnLeft(1.0);
    } else {
        robot.turnRight(1.0);
    }
}
```

The only thing remaining in the commands above is to write the condition to detect the difference between the two light sensors. This can be done using the expression:

```
robot.getLight("left") < robot.getLight("right")
```

Do This: Write a complete program that implements the above behavior and test it on your robot.

You may have noticed that even in uniform lighting conditions sensors tend to report different values. It is generally a good idea to threshold the difference when making the decision above. Say we set the threshold to a difference of at least 50. That is, if the left and right sensors differ by at least 50 then turn towards the brighter sensor. What happens if the difference is less than the threshold? Let us decide that in that case the robot will stay still. This behavior can be captured by the following:

```
int thresh = 50;

while (timeRemaining(30)) {
    // Get sensor values for left and right light sensors
    int L = robot.getLight("left");
    int R = robot.getLight("right");

    // decide how to act based on sensors values
    if ((L - R) > thresh) {
        // left is seeing less light than right so turn right
        robot.turnRight(1.0);
    } else if ((R - L) > thresh) {
        // right is seeing less light than left, so turn left
        robot.turnLeft(1.0);
    } else {
```

```
        // the difference is less than the threshold, stay put
        robot.stop();
    }
}
```

Notice how we have used the variable `thresh` to represent the threshold value. This is good programming practice. Since the performance of sensors varies under different light conditions, this allows you to adjust the threshold by simply changing that one value. By using the name `thresh` instead of a fixed value, say 50, you only have to make such changes in one place of your program.

In the statements above, there is a pattern that you will find recurring in many programs that define robot behaviors using simple decisions:

```
while (timeRemaining(<seconds>)) {
    <sense>
    <decide and then act>
}
```

Such behaviors are called *reactive* behaviors. That is, a robot is reacting to the change in its environment by deciding how to act based on its sensor values. A wide range of robot behaviors can be written using this program structure. Below, we present descriptions of several interesting, yet simple automated robot behaviors. Feel free to implement some of them on your robot.

Simple Reactive Behaviors


Most of the behaviors described below require selection among alternatives using conditional expressions and `if`-statements.

Refrigerator Detective: As a child did you ever wonder if that refrigerator light was always on? Or did it shut off when you closed the door? Well, here is a way to find out. Build a refrigerator detective robot that sits inside the fridge and tells you if the light is on or off!

Burglar Alarm Robot: Design a robot that watches your dorm door. As soon as the door opens, it sounds an alarm (beeps).

Wall Detector: Write a robot program that goes straight and then stops when it detects a wall in front. You will be using the IR sensors for this task.

Hallway Cruiser: Imagine your robot in an environment that has a walled corridor going around a 2 ft by 2 ft square box (see picture on right). Write a program that will enable the robot to go around this box. One strategy you can use is to have the robot go forward in a straight line until it bumps into a wall. After a bump it will proceed to make a 90 degree turn (you may need to have it go backwards a little to enable turning room) and then continue again in a straight line.



Measuring Device: You have calibrated your robot with regards to how far it travels in a given amount of time. You can use that to design a robot that measures space. Write a program that enables a robot to measure the width of a hallway.

Follower: Write a robot program to exhibit the following behavior: The robot prefers to stay near a wall (in front). If it does not have a wall in front of it, it moves forward until it finds it. Test your program first by placing the robot in a play pen. Ensure that your program behaves as described. Next, place it on a floor and hold a blank piece of paper in front of it (close enough so the robot can detect it). Now, slowly move the paper away from the robot. What happens?

Designing Reactive Behaviors

Most of the robot behaviors that are implemented using the Braitenberg style rely on a few simple things: selecting one or more sensors; choosing the kind of wiring (straight or crossed); and selecting normalization functions for each sensor. While you can guess the behavior that may result from these designs the only way to confirm this is by actually watching the robot carry out the behavior. You also saw how, using if-statements you can design simple, yet interesting robots. In this section we will design additional reactive behaviors.

Light Following

To begin, it will be fairly straightforward to extend the behavior of the light orienting behavior from above into one that results in a light follower robot. That is, with a flashlight you will be able to guide the robot to follow you around. Again, you have to start by observing the range of values reported by the robot under various lighting conditions. If a flashlight is going to be the bright light source, you will observe that the light sensors report very low values when a light is shining directly on them (typically in the 0..50 range). Thus, deciding which way to go (forward, turn left, or turn right) can be decided based on the sensor readings from the three light sensors. The structure of the program appears as follows:

```
// Light follower

#include "Myro.h"

// program settings...

int thresh = 50;
double fwdSpeed = 0.8;
double cruiseSpeed = 0.5;
double turnSpeed = 0.7; // left turn, -0.7 will be right turn

int main() {
    connect();
    while (true) {
```

```
// get light sensor values for left, center, and right
int L = robot.getLight("left");
int C = robot.getLight("center");
int R = robot.getLight("right");

// decide how to act based on sensor values
if (C < thresh) {
    // bright light from straight ahead, go forward
    robot.move(fwdSpeed, 0);
} else if (L < thresh) {
    // bright light at left, turn left
    robot.move(cruiseSpeed, turnSpeed);
} else if (R < thresh) {
    // bright light on right side, turn right
    robot.move(cruiseSpeed, -turnSpeed);
} else {
    // no bright light, move forward slowly (or stop?)
    robot.move(cruiseSpeed/2, 0);
}
}
disconnect();
}
```

Notice that, in the program above, we have decided to set values for light threshold (`thresh`) as well as movements to specific values. Also, in all cases, we are using the `move` command to specify robot movement. This is because the `move` command allows us to blend translation and rotation movement. Additionally, notice that regardless of the sensor values, the robot is always moving forward some amount even while turning. This is essential since the robot has to follow the light and not just orient towards it. In the case where there is no bright light present, the robot is still moving forward (at half the cruise speed).

Do This: Implement the light following program as described above and observe the robot's behavior. Try adjusting the value settings (for threshold as well as motor speeds) and note the changes in the robot's behaviors. Also, do you observe that this behavior is similar to any of the Braitenberg vehicles described above? Which one?

In the design of the light following robot above, we used a threshold value for detecting the presence of bright light. Sometimes it is more interesting to use differential thresholds for sensor values. That is, is the light sensor's value different from the ambient light by a certain threshold amount? You can use the `senses` function again observe the differences from ambient light and modify the program above to use the differential instead of the fixed threshold.

Here is another idea. Get several of your classmates together in a room with their robots, all running the same program. Make sure the room has plenty of floor space and a large window with a curtain. Draw closed the curtains so the outside light is temporarily blocked. Place the robots all over the room and start the program. The robots will scurry around, cruising in the direction of their initial orientation. Now, slowly draw the curtains open to let in more light. What happens?

Avoiding Obstacles

Obstacles in the path of a robot can be detected using the IR sensors in front of the robot. Then, based on the values obtained, the robot can decide to turn away from an approaching obstacle using the following algorithm:

```
if obstacle straight ahead, turn (left or right?)
if obstacle on left, turn right
if obstacle on right, turn left
otherwise cruise
```

This can be implemented using the program below:

```
// Avoiding Obstacles

#include "Myro.h"

// program settings...

double cruiseSpeed = 0.6;
double turnSpeed = 0.5; // this is a left turn,
                        // -0.5 will be right turn
```

```
int main() {
    connect();
    while (true) {
        // get sensor values for left and right IR sensors
        // and negate them
        int L = 1 - robot.getIR("left");
        int R = 1 - robot.getIR("right");

        // decide how to act based on sensors values
        if (L && R) {
            // obstacle straight ahead, turn (randomly)
            robot.move(0, turnSpeed);
        } else if (L) {
            // obstacle on left, turn right
            robot.move(cruiseSpeed, -turnSpeed);
        } else if (R) {
            // obstacle on right, turn left
            robot.move(cruiseSpeed, turnSpeed);
        } else {
            // no obstacles
            robot.move(cruiseSpeed, 0);
        }
    }
    disconnect();
}
```

As in the case of the light follower, observe that we begin by setting values for movements. Additionally, we have flipped the values of the IR sensors so that the conditions in the `if`-statements look more natural. Recall that the IR sensors report a 1 value in the absence of any obstacle and a 0 in the presence of one. By flipping them (using `1 - value`) the value is 1 for an obstacle present and 0 otherwise. These values make it more natural to write the conditions in the program above. Remember in C++, a 0 is equivalent to `false` and a 1 is equivalent to `true`. Read the program above carefully and make sure you understand these subtleties. Other than that, the program structure is very similar to the light follower program.

Another way to write a similar robot behavior is to use the value of the stall sensor. Recall that the stall sensor detects if the robots has bumped against

something. Thus, you can write a behavior that doesn't necessarily avoid obstacles, but navigates itself around by bumping into things. This is very similar to a person entering a dark room and then trying to feel their way by touching or bumping slowly into things. In the case of the robot, there is no way to tell if the bump was on its left or right. Nevertheless, if you use the program (shown below) you will observe fairly robust behavior from the robot.

```
// Avoiding Obstacles by bumping

#include "Myro.h"

// program settings...

double cruiseSpeed = 1.0;
double turnSpeed = 0.5; // this is a left turn,
                        // -0.5 will be right turn

int main() {
    connect();
    while (true) {

        if (robot.getStall()) {
            // I am stalled, turn (randomly?)
            robot.move(0, turnSpeed);
        } else {
            // I am not stalled, cruise on
            robot.move(cruiseSpeed, 0);
        }
    }
    disconnect();
}
```


At times, you may notice that the robot gets stuck even when trying to turn. One remedy for this is to stop the robot, back up a little, and then turn.

Do This: Implement the program above, observe the robot behavior. Next, modify the program as suggested above (when stalled stop, backup, then turn).

Maze Solver: Create a simple maze for your robot. Place the robot at one end of the maze and use the obstacle avoidance programs from above (both versions). Does your robot solve the maze? If not, note if your maze is right handed or left handed (i.e. every turn is a right turn or left turn in the maze), or both. Modify the obstacle avoidance programs to solve the right-handed, left-handed mazes. How would you enable the robot to solve a maze that has both right and left turns?

Corral Exiting

Given that a simple obstacle avoidance program can enable a robot to solve simple mazes, we can also design more interesting behaviors on top of that. Imagine a corral: an enclosed area with maze like partitions and an entrance, with a light source at the entrance (see picture on right). Given the robot's position, can we design a behavior that will enable the robot to exit the corral?



One can design a solution for the specific corral shown here: follow a wall (any wall) until it sees bright light then switch to light seeking. Can the Scribbler be designed to follow a wall? Remember the Fluke dongle has left and right obstacle sensors that are pointing to its sides. Another approach will be to combine the obstacle avoidance behavior from above with the light seeker behavior. That is, in the absence of any bright light, the robot moves around the corral avoiding obstacles and when it sees a bright light, it heads towards it. The hard part here will be to detect that it has exited the corral and needs to stop.

Summary

Braitenberg uses very simple ideas to enable people to think about the way animal and human brains and bodies are wired. For example, in humans, the optic nerves (as do some others) have crossed connections inside the brain. That is, the nerves from the left eye are connected to the right side of the brain and vice versa. Actually they cross over and some information from either side is also represented on the same side (that is there are straight as well as crossed connections). However, it is still a puzzle among scientists as to why this is the case and what, if any, are the advantages or disadvantages of this scheme. Similarly, observing the behaviors of Vehicles 2a and 2b one can easily see in them parallels in the behavior of several animals, like flies orienting towards light/heat sources. Simple robot behaviors can provide deep insights into complex behavior: that the observation and analysis of something is an uphill task if one doesn't know the internal structure. And, by constructing simple internal structures one can arrive at seemingly complex behaviors. These seemingly complex behaviors have also been shown to influence group behavior in insects (see the picture of article on next page). That is, robots that do not look anything like insects, and not too different in size than the Scribbler, can be used to influence insect behavior in many situations.

In this chapter, we have attempted to give you a flavor for the idea of synthetic psychology. At the same time you have also learned how to program internal structures in a robot brain and learned several techniques for robot control.

Background

All the numbered vehicles described here were developed in a set of thought experiments designed by Valentino Braitenberg in his book, *Vehicles: Experiments in Synthetic Psychology*, MIT Press, 1984.

Some of the other vehicles described here were designed by David Hogg, Fred Martin, and Mitchel Resnick of the MIT Media Laboratory. Hogg et al

used specialized electronic LEGO bricks to build these vehicles. For more details, see their paper titled, *Braitenberg Creatures*.

Story from *Science Magazine*, January 10, 2008

BEHAVIOR

Robot Cockroach Tests Insect Decision-Making Behavior

Science-fiction writers have long envisioned societies in which the boundaries between humans and lifelike droids blur and man and machine freely intermingle. José Halloy has taken the first steps toward creating that world, at least for insects. His tiny, autonomous robots lack legs, wings, and antennae, but they nonetheless pass muster with cockroaches. Indeed, these wheeled machines are so well accepted by the household pests that the robots become part of the insects' collective decision-making process, Halloy, a theoretical biologist at the Free University of Brussels, Belgium, and his colleagues report on page 1155. The robots persuaded many of their insect "peers" to hide in an unconventional place.


Halloy's innovative approach puts theories of collective behavior among insects into practice. "We can manipulate these behaviors very easily in a model, but doing so in experiments is often challenging," explains ethologist Jerome Buhl of the University of Sydney, Australia. Others have used remote-controlled robots to study animal behavior but not autonomous ones that later

level of darkness of the shelter and the number and activity of its fellow roaches. Halloy's group then used this model to program robots designed by him and Francesco Mondada and other engineers at the École Polytechnique Fédérale de Lausanne, Switzerland.

The roaches usually ran away from the robots but not if the machines smelled like the insects. For the experiments, Halloy and Sempo covered the robots with a filter paper containing the pheromone equivalent of one cockroach.

Halloy initially programmed the robots to have the same darkness preference as the cockroaches, and they joined the cockroaches at whatever shelter the majority chose to rest in. Next, Halloy programmed the robots to prefer the lighter shelter. About 60% of the time, the robots tipped the group's preference

Can't we be friends? Cockroaches seem to accept this robot as one of their own once it's coated with pheromone.



in favor of the light shelter. "This is a true example of automated leadership," says David Sumpter of Uppsala University in Sweden.

To read more about robots influencing insect behavior see the November 16, 2007 issue of *Science* magazine. The primary article that is discussed in the picture above is by Halloy *et al*, *Social Integration of Robots into Groups of Cockroaches to Control Self-Organized Choices*, *Science*, November 16, 2007. Volume 318, pp 1155-1158.

Myro Review

There were no new Myro features introduced in this chapter.

C++ Review

The `if`-statement in C++ has the following forms:

```
if (<condition>) {  
    <this>  
}
```

```
if (<condition>) {  
    <this>  
} else {  
    <that>  
}
```

```
if (<condition-1>) {  
    <this>  
} else if (<condition-2>) {  
    <that>  
} else if (<condition-3>) {  
    <something else>  
...  
...  
} else {  
    <other>  
}
```

The conditions can be any expression that results in a `true`, `false`, `1`, or `0` value. Review Chapter 4 for details on writing conditional expressions.

Exercises

1. An even better way of averaging the ambient light conditions for purposes of normalization is to have the robot sample ambient light all around it. That is, turn around a full circle and sample the different light sensor values. The ambient value can then be set to the average of all the light values. Write a function called, `setAmbient` that rotates the robot for a full circle (or you could use time), samples light sensor values as it rotates, and then returns the average of all light values. Change the line:


```
Ambient = (robot.getLight("left") +  
           robot.getLight("center") +  
           robot.getLight("right")) / 3.0;
```


to the line:

```
Ambient = setAmbient();
```

Try out all of the earlier behaviors described in this chapter to see how this new mechanism affects the robot's behavior.

2. Design and implement a program that exhibits the corral exiting behavior described in this chapter.
3. Implement the refrigerator detective behavior described in this chapter.
4. Implement the burglar alarm robot described in this chapter.
5. Implement the hallway cruiser behavior described in this chapter.
6. In addition to movements try to integrate music/sound output in your robot behaviors and observe how the addition of sounds amplifies the perception of the robot's personality.





Behavior Control

Oh, Behave!
Austin Powers (played by Mike Myers) in the movie
Austin Powers: International Man of Mystery,
New Line Cinema, 1997.

Opposite page: My Traffic Lights
Illustration by Yingshun Wong (yingshun.co.uk)

Writing programs is all about exercising control. In the case of a robot's brain your program directs the operations of a robot. However, it is also important to realize that the program itself is really controlling the computer. That is, when you write C++ programs you are controlling the computer that is then communicating with the robot. Your program is directing the computer to control the robot. If you take Myro out of the picture you are writing programs to control the computer. This is the sense in which learning with robots also leads to learning computing. Every program you write is doing computation. Contrary to popular misconceptions computing is not just about doing calculations with numbers. Controlling your robot is also computing, as is predicting the world's population, or composing an image, etc. This is one aspect of control.

When writing robot control programs, the structure you use to organize the program itself is a control strategy. Programming a robot is specifying automated control. As a programmer or behavior designer you structure your program to accomplish the goals of the behavior: how the sensors are used to decide what to do next. This is another aspect of control. So far, you have seen how to write control programs using Braitenberg style sensor-motor *wiring*. You have also seen how to specify reactive control. These are examples of two robot control paradigms.

In this chapter we delve further into the world of computation and robot control paradigms. We will learn how to write robot control programs for more complex and more robust robot tasks. We will also see how, using the concepts learned so far, we can write useful and interesting computer applications.

Behavior-based Control

When writing robot control programs, so far, you have used a very basic technique in designing control programs:

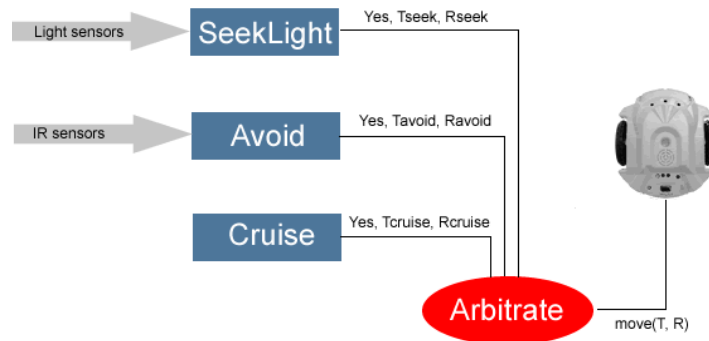
```
int main() {  
    // do forever or for some time  
    // or until a certain condition is satisfied  
  
    // sense and transform sensor values  
    // reason or decide what to do next  
    // do it  
}
```

As you have seen, such a control program works well for many simple tasks. However, you may have already run into many situations where, once the task gets a little complex, it becomes difficult to structure a program in terms of a single stream of control as shown above. For example, the corral exiting behavior from the last chapter requires you to combine two simple behaviors: solve a maze (avoid obstacles) and seek light to get out of the corral. As you have seen before, it is fairly easy to program each of the individual behaviors: obstacle avoidance; light following. But, when you combine these behaviors to accomplish the corral exiting behavior two things happen: you are forced to amalgamate the two control strategies into a single one and it may become difficult to decide which way to combine them; additionally, the resulting program is not very pretty and hard to read. In reality, hardly any robot programs are written that way. In this section, we will look at a different way of structuring robot programs that makes designing behaviors easy, and yet, the resulting structure of the overall program is also clean and straightforward. You can design some very sophisticated behaviors using these ideas.

People in the robotics community call the style of programming shown above as *reactive control* or *direct control*. Also referred to as *sensor fusion*, the resulting programs are purely sensor driven and hence appear to be too bottom-up. That is, the values of the sensors drive the logic of control as opposed to the goals of the robot tasks themselves. In *behavior-based control* you get away from sensors and focus the design of your robot programs based on the number and kinds of behaviors your robot has to carry out.

Let us look at how behavior-based control enables us to design the corral exiting behavior. The robot essentially has to carry out three kinds of behaviors: cruise (in the absence of any obstacles and/or light), avoid

obstacles (if present), and seek light (if present). In a behavior-based style of writing programs, you will define each of these behaviors as an individual decision unit. Thus, each is quite simple and straightforward to write. Next, the control program has to fuse the behaviors recommended by each individual behavior unit. Look at the picture shown below:



In the diagram above, we have shown the three basic behaviors that we are trying to combine: *Cruise*, *Avoid*, *SeekLight*. Each of these behaviors outputs a triple: *Yes/No*, *Translate Speed*, *Rotate Speed*. A *Yes* implies that the behavior module has a recommendation. *No* implies that it doesn't. That is, it allows the possibility of a behavior having no recommendation. For example, in the corral exiting situation, in the absence of a light source being sensed by the robot, the *SeekLight* module will not have any recommendation. It then becomes the task of the arbitrator (or the decision module) to decide which of the available recommendations to use to drive the robot. Notice that in the end to control the robot, all one has to do is decide how much to translate and rotate. Many different arbitration schemes can be incorporated. We will use a simple but effective one: Assign a priority to each behavior module. Then the arbitrator always chooses the highest priority recommendation. This style of control architecture is also called *subsumption architecture*. In the figure above, we have actually drawn the modules in the order of their priority: the higher the module is in the figure, the higher its priority. The lowest behavior,

cruise, does not require any sensors and is always present: it wants the robot to always go forward.

Arranging a control based on combining simple behaviors has several advantages: you can design each individual behavior very easily; you can also test each individual behavior by only adding that behavior and seeing how well it performs; you can incrementally add any number of behaviors on top of each other. In the scheme above, the control regime implies that the robot will always cruise forward. But, if there is an obstacle present, it will override the cruise behavior and hence try to avoid the obstacle. However, if there is a light source detected, it will supersede all behaviors and engage in light seeking behavior. Ideally, you can imagine that all the behaviors will be running simultaneously (asynchronously). In that situation, the arbitrator will always have one or more recommendations to adopt based on priority.

Let us develop the program that implements behavior-based control. First, we define each behavior to possibly recommend an action. To do this, we use a C++ feature that you have not seen before; it is called a *structure*:

```
struct rec {
    bool active;
    double translateSpeed;
    double rotateSpeed;
};
```

In effect `struct` defines a new data type, which is somewhat like a vector, in that it groups several data values together, but unlike a vector in that the values do not have to be of the same type. In this case, the `rec` structure is composed of three *fields*, a Boolean called `active`, and two doubles called `translateSpeed` and `rotateSpeed`. The intent is that if `active` is `true`, then the structure represents an active recommendation that the robot move forward at rate `translateSpeed` while turning at rate `rotateSpeed`.

To make a recommendation, we create a structure of type `rec` and initialize its fields, which can be accomplished with the function:

```
rec make_rec (bool a, double t, double r) {
    rec R;
    R.active = a;
    R.translateSpeed = t;
    R.rotateSpeed = r;
    return R;
}
```

This function has three arguments (a Boolean and two doubles) and it returns a structure of type `rec`. To do this it declares a variable `R` of type `rec`, and initializes its three fields to the supplied parameters. Notice that the fields of `R` are named `R.active`, `R.translateSpeed`, and `R.rotateSpeed`. Given these definitions, each of the desired behaviors can recommend an action:

```
double cruiseSpeed = 0.8;
double turnSpeed = 0.8;
int lightThresh = 80;

rec cruise () {
    // is always active, just move forward
    return make_rec (true, cruiseSpeed, 0);
}

rec avoid () {
    // see if there are any obstacles
    int L = 1 - robot.getIR("left");
    int R = 1 - robot.getIR("right");
    if (L) {
        return make_rec (true, 0, -turnSpeed);
    } else if (R) {
        return make_rec (true, 0, turnSpeed);
    } else {
        return make_rec (false, 0, 0);
    }
}

rec seekLight () {
    int L = robot.getLight("left");
    int R = robot.getLight("right");
}
```



```
if (L < lightThresh) {
    return make_rec (true, cruiseSpeed/2.0, turnSpeed);
} else if (R < lightThresh) {
    return make_rec (true, cruiseSpeed/2.0, -turnSpeed);
} else {
    return make_rec (false, 0, 0);
}
}
```

In the above, you can see that each individual behavior is simple and is easy to read (and write). Also, notice that the behaviors do not always recommend an action. If the first parameter to `make_rec` is `false`, then the behavior is not recommending any action (it is inactive); in this case, the behavior has not been *released* by the stimuli (sensor values). There are several ways to incorporate these behaviors into a behavior-based program. Here is one:

```
int main () {
    connect ();

    while (true) {
        rec decision = arbitrate ();
        robot.move (decision.translateSpeed,
                    decision.rotateSpeed);
    }

    disconnect ();
}
```

The main program calls the `arbitrate` function to decide which recommendation to take, which is saved in the variable `decision`. Then the motor behavior (translation and rotation speeds) from the chosen recommendation is applied to the robot. The function `arbitrate` is simple enough; it returns the first (highest priority) active recommendation:

```
rec arbitrate () {
    rec R;
    R = seekLight(); if (R.active) return R;
    R = avoid();     if (R.active) return R;
    R = cruise();    if (R.active) return R;
}
```

For readability, we have written several things a little differently than previously in this book. First, we have written three statements (an assignment, an `if`, and a `return`) on a single line, to emphasize the close relationship among them. As previously mentioned, C++ really doesn't care how many or how few lines are used to write something. Second, the `return`-statements, which are the bodies of the `if`-statements, are not surrounded by curly braces. As it turns out, all the C++ statements that govern a body (e.g., `for`, `else`, `while`, function definitions) allow the curly braces to be omitted if the body is a single statement; in effect the braces act like parentheses to group a set of statements into a body. However, it is easier to maintain programs if you always use curly braces, even when they are not strictly necessary, since otherwise you may forget to put them in if you change a single-statement body to a multiple-statement body. As in prose, style is important in programming, and coding conventions such as this one make programs easier to read and to maintain.

Do This: Implement the program above and see how well the robot behaves in navigating around and exiting the corral. What happens if you change the priority (ordering in the list) of behaviors?

The C++ Math Library

Most programming languages, C++ included, provide a healthy selection of libraries of useful functions so you do not have to write them. Such libraries are often termed as an *application programming interface* or an *API*. Earlier you were introduced to the `cstdlib` library provided by C++. It contains several useful functions that provide facilities for generating random numbers. Similarly, C++ also provides a `cmath` library that provides frequently used mathematical functions. In Chapter 6, we used the function `exp` from the

`cmath` library to normalize sensor values. Some of the commonly used functions in the `cmath` library are listed below:

Commonly used functions in the `cmath` library

`ceil(x)` Returns the ceiling of x as a float, the smallest integer value greater than or equal to x .

`floor(x)` Returns the floor of x as a float, the largest integer value less than or equal to x .

`exp(x)` Returns e^x .

`log(x[, base])` Returns the logarithm of x to the given base. If the base is not specified, return the natural logarithm of x (i.e., $\log_e x$).

`log10(x)` Returns the base-10 logarithm of x (i.e. $\log_{10} x$).

`pow(x, y)` Returns x^y .

`sqrt(x)` Returns the square root of x (\sqrt{x}).

Some of the other functions available in the `cmath` library are listed at the end of the chapter. In order to use any of these all you have to do is include the `cmath` library:

```
#include <cmath>
```

```
cout << ceil(5.34) << endl;
```

```
6
```

```
cout << floor(5.34) << endl;
```

```
5
```

```
cout << exp(3) << endl;
```

```
20.0855
```

```
cout << log10(1000) << endl;
```

```
3
```

```
cout << log(1024, 2) << endl;
```

```
10
```

```
cout << pow(2, 10) << endl;
```

```
1024
```

```
cout << sqrt(7.0) << endl;
```

```
2.64575
```

Doing Computations

Lets us weave our way back to traditional style computing for now. You will see that the concepts you have learned so far will enable you to write lots of different and more interesting computer applications. It will also give you a clear sense of the structure of typical computer programs. Later, in Chapter 11, we will also return to the larger issue of the design of general computer programs.

A Loan Calculator

Your current car, an adorable 1992 SAAB 93 was bought used and, for past several months, you have had nothing but trouble keeping the car on the road. Last night the ignition key broke off in the key slot when you were trying to start it and now the broken piece would not come out (this used to happen a lot with older SAAB's). The mechanic has to dismantle the entire ignition assembly to get the broken key out and it could cost you upwards of \$500. The car's engine, which has done over 185,000 miles, has left you stranded on the road many times. You have decided that this is it; you are going to go out and get yourself a brand new reliable car. You have been moonlighting at a restaurant to make extra money and have managed to save \$5500.00 for exactly this situation. You are now wondering what kind of new car you can buy. Obviously, you will have to take a loan from a bank to finance the rest of the cost but you are not sure how big a loan, and therefore what kind of car,

you can afford. You can write a small C++ program to help you try out various scenarios.

You can either dream (realistically) of the car you would like to buy, or you can go to any of the helpful car buying sites on the web (www.edmunds.com is a good place to start). Let us say that you have spent hours looking at features and options and have finally narrowed your desires to a couple of choices. Your first choice is going to cost you \$22,000.00 and your second choice is priced at \$18,995.00. Now you have to decide which of these you can actually afford to purchase.

First, you go talk to a couple of banks and also look at some loan offers on the web. For example, go to bankrate.com and look for current rates for new car loans.

Suppose the loan rates quoted to you are: 6.9% for 36 months, 7.25% for 48 months, and 7.15 for 60 months.

You can see that there is a fair bit of variation in the rates. Given all this information, you are now ready to write a program that can help you figure out which of the two choices you may be able to make. In order to secure the loan, you have to ensure that you have enough money to pay the local sales tax (a 6% sales tax on a \$20,000 car will add up to a hefty \$1200!). After paying the sales tax you can use the remainder of the money you have saved up towards the down payment. The remainder of the money is the amount that you would borrow. Depending on the type of loan you choose, your monthly payments and how long you will make those payments will vary. There is a simple formula that you can use to estimate your monthly payment:

$$\text{MonthlyPayment} = \frac{\text{LoanAmount} * \text{MonthlyInterestRate}}{1 - e^{-\text{LoanTerm} * \log(1 + \text{MonthlyInterestRate})}}$$

Whoa! That seems complicated. However, given the formula, you can see that it really requires two mathematical functions: $\log(x)$ and e^x , both of which are

available in the C++ `cmath` library. Suddenly, the problem seems not that hard.

Let us try and outline the steps needed to write the program: First, note the cost of the car, the amount of money you have saved, and the sales tax rate. Also, note the financials: the interest rate, and the term of the loan. The interest rate quoted is generally the annual percentage rate (APR) convert it to monthly rate (by dividing it by 12). Next, compute the sales tax you will pay. Use the money left to make a down payment. Then determine the amount you will borrow. Plug in all of the values in the formula and compute the monthly payment. Also, compute the total cost of the car. Output all the results. Next, we can take each of the above steps and start to encode them into a program. Here is a first order refinement:

```
int main() {
    // First, note the cost of the car (Cost),
    // the amount of money you have saved (Cash),
    // and the sales tax rate (TaxRate) (6% e.g.)

    // Also, note the financials: the interest rate (APR),
    // and the term of the loan (Term)
    // The interest rate quoted is generally the annual
    // percentage rate (APR)
    // Convert it (APR) to monthly rate (divide it by 12) (MR)

    // Next, compute the sales tax you will pay (SalesTax)
    // Use the money left to make a down payment (DownPayment)
    // Then determine the amount you will borrow (LoanAmount)

    // Plug in all of the values in the formula and compute
    // the monthly payment (MP)

    // Also, compute the total cost of the car. (TotalCost)

    // Output all the results
}
```

Above, we have taken the steps and converted them into a skeletal C++ program. All the steps are converted to C++ comments and where needed, we

have decided the names of variables that will hold the values that will be needed for the calculations. This is useful because this also helps determine how the formula will be encoded and also helps determine what values can be programmed in and which ones you will have to supply as input. Making the program require inputs will easily enable you to enter the different parameters and then based on the outputs you get, you can decide which car to buy. Let us encode all the inputs first:

```
int main() {
    // Input variables
    double Cost, Cash, APR;
    int term;

    // First, note the cost of the car (Cost),
    cout << "Enter the cost of the car: $";
    cin >> Cost;

    // the amount of money you have saved (Cash),
    cout << "Enter the amount of money you saved: $";
    cin >> Cash;

    // and the sales tax rate (TaxRate) (6% e.g.)
    double SalesTaxRate = 6.0;

    // Also, note the financials: the interest rate (APR),
    // and the term of the loan (Term)
    // The interest rate quoted is generally the annual
    // percentage rate (APR)
    cout << "Enter the APR for the loan (in %): ";
    cin >> APR;

    // Convert it (APR) to monthly rate (divide it by 12) (MR)

    // Next, compute the sales tax you will pay (SalesTax)
    // Use the money left to make a down payment (DownPayment)
    // Then determine the amount you will borrow (LoanAmount)

    // Plug in all of the values in the formula and compute
    // the monthly payment (MP)

    // Also, compute the total cost of the car. (TotalCost)
```

```
    // Output all the results  
}
```

We have refined the program to include the inputs that will be needed for each run of the program. Notice that we chose not to input the sales tax rate and instead just assigned it to the variable `SalesTaxRate`. If you wanted, you could also have that be entered as input. What you choose to have as input to your program is your design decision. Sometimes the problem may be framed so it explicitly specifies the inputs; sometimes you have to figure that out. In general, whatever you need to make your program more versatile is what you have to base your decisions on. For instance, fixing the sales tax rate to 6.0 will make the program usable only in places where that rate applies. If, for example, you wanted your friend in another part of the country to use the program, you should choose to make that also an input value. Let us go on to the next steps in the program and encode them in C++. These are mainly computations. The first few are simple. Perhaps the most complicated computation to encode is the formula for computing the monthly payment. All of these are shown in the version below.

```
#include <iostream>  
#include <cmath>  
using namespace std;  
  
int main() {  
    // Input variables  
    double Cost, Cash, APR;  
    int term;  
  
    // First, note the cost of the car (Cost),  
    cout << "Enter the cost of the car: $";  
    cin >> Cost;  
  
    // the amount of money you have saved (Cash),  
    cout << "Enter the amount of money you saved: $";  
    cin >> Cash;  
  
    // and the sales tax rate (TaxRate) (6% e.g.)  
    double SalesTaxRate = 6.0;  
  
    // Also, note the financials: the interest rate (APR),
```



```
// and the term of the loan (Term)
// The interest rate quoted is generally the annual
// percentage rate (APR)
cout << "Enter the APR for the loan (in %): ";
cin >> APR;

// Input the term of the loan (Term)
cout << "Enter length of loan term (in months): ";
cin >> term;

// Convert it (APR) to monthly rate (divide it by 12) (MR)
// also divide it by 100 since the value input is in %
double MR = APR/12.0/100.0;

// Next, compute the sales tax you will pay (SalesTax)
double SalesTax = Cost * SalesTaxRate / 100.0;

// Use the money left to make a down payment (DownPayment)
double DownPayment = Cash - SalesTax;

// Then determine the amount you will borrow (LoanAmount)
double LoanAmount = Cost - DownPayment;

// Plug in all of the values in the formula and compute
// the monthly payment (MP)
double MP = (LoanAmount * MR)
            / (1.0 - exp(-term * log(1.0+MR)));

// Also, compute the total cost of the car. (TotalCost)
double TotalCost = SalesTax + DownPayment + MP * term;

// Output all the results
cout << "Here are the details about your new car...\n";
cout << "-----\n";
cout << endl;
cout << "Money you have saved $" << Cash << endl;
cout << "Cost of the car $" << Cost << endl;
cout << "Sales Tax rate is " << SalesTaxRate << "%\n";
cout << "Sales Tax on the car $" << SalesTax << endl;
cout << "Your down payment will be $" << DownPayment
    << endl;
cout << "You will be borrowing $" << LoanAmount << endl;
cout << "A " << term << " month loan at " << APR
```

```
    << "% APR\n";  
    cout << "Your monthly payment will be $" << MP << endl;  
    cout << "Total cost will be $" << TotalCost << endl;  
    cout << endl;  
    return 0;  
}
```

Notice that some of the output strings contain the special “escape sequence” “\n” (“newline”), which means “go to a new line,” and thus has the same effect as `<< endl`; sometimes it’s more convenient, as shown here.

Do This: When you enter the above C++ program and run it, you can enter the data about your car. Here is a sample run:

```
Enter the cost of the car: $20000.00  
Enter the amount of money you saved: $5500.00  
Enter the APR for the loan (in %): 6.9  
Enter the length of loan term (in months): 36  
Here are the details about your new car...  
-----  
  
Money you have saved $5500  
Cost of the car $20000  
Sales Tax rate is 6%  
Sales Tax on the car $1200  
Your down payment will be $4300  
You will be borrowing $15700  
A 36 month loan at 6.9% APR  
Your monthly payment will be $484.053  
Total cost will be $22925.9
```

It appears that at for the \$20000.00 car, for a 36 month 6.9% loan you will end up paying \$484.05 (or \$484.06 depending upon how your loan company rounds pennies!).

When you need to restrict your output values to specific decimal places (two in the case of dollars and cents, for example), you can use the `printf` formatted output function built into C++. For example, using a format string, you can specify how to output a fixed-point value as follows:

```
double pi = acos(-1); // this computes the value of pi

printf ("Value of PI %5.3f in 3 places after the decimal\n",
       pi);
```

The above is a C++ expression that has the syntax:

```
printf (<string>, <expression>, <expression>, ...);
```

The `<expression>`s are printed according to the format `<string>`. Inside `<string>` there is a format specification beginning with a `%`-sign and ending in an `f`. What follows the `%`-sign is a numerical specification of the value to be inserted at that point in the string. The `f` in the specification refers to the fact it is for a floating point value. Between the `%`-sign and the `f` is a number, `5.3`. This specifies that the floating point number to be inserted will take up at least 5 spaces, 3 of which will be after the decimal. One of the spaces is always occupied by the decimal. Thus, the specification `%5.3f` specifies a value to be printed in the following format:

```
Value of PI -.--- expressed in 3 places after the decimal
```

You can see the result of executing this in C++ below:

```
printf ("Value of PI %5.3f in 3 places after the decimal\n",
       pi);
```

```
Value of PI 3.142 in 3 places after the decimal
```

```
printf ("Value of PI %7.4f in 4 places after the decimal\n",
       pi);
```

```
Value of PI 3.1416 in 4 places after the decimal
```

In the second example above, we replaced the specification with `%7.4f`. Notice that the resulting string allocates seven spaces to print that value. If there are more spaces than needed they get padded by blanks on the leading edge (notice the extra space before 3 in the second example above). If the

space specified is less, it will always be expanded to accommodate the number. For example:

```
printf ("Value of PI %1.4f in 4 places after the decimal\n",  
pi);
```

```
Value of PI 3.1416 in 4 places after the decimal
```

We deliberately specified that the value be 1 space wide with 4 spaces after the decimal (i.e. `%1.4f`). As you can see, the space was expanded to accommodate the value. What is assured is that the value is always printed using the exact number of spaces after the decimal. Here is another example:

```
printf ("5 is also %1.3f with 3 places after the decimal.\n",  
5.0);
```

```
5 is also 5.000 with 3 places after the decimal.
```

Thus, the value is printed as `5.000` (i.e. the three places after the decimal are always considered relevant in a specification like `%1.3f`). Similarly, for specifying whole number or integer values you can use the letter-`d`, and for strings you can use the letter-`s`:

```
printf ("Hello %10s, how are you?\n", "Arnold");
```

```
Hello      Arnold, how are you?
```

By default, longer specifications are right-justified. You can use a `%-` specification to left-justify. For example:

```
printf ("Hello %-10s, how are you?\n", "Arnold");
```

```
Hello Arnold , how are you?
```

Having such control over printed values is important when you are trying to output tables of aligned values. Let us modify our program from above to use these formatting features:

```
#include <iostream>
#include <cmath>
using namespace std;

int main() {
    // Input variables
    double Cost, Cash, APR;
    int term;

    // First, note the cost of the car (Cost),
    cout << "Enter the cost of the car: $";
    cin >> Cost;

    // the amount of money you have saved (Cash),
    cout << "Enter the amount of money you saved: $";
    cin >> Cash;

    // and the sales tax rate (TaxRate) (6% e.g.)
    double SalesTaxRate = 6.0;

    // Also, note the financials: the interest rate (APR),
    // and the term of the loan (Term)
    // The interest rate quoted is generally the annual
    // percentage rate (APR)
    cout << "Enter the APR for the loan (in %): ";
    cin >> APR;

    // Input the term of the loan (Term)
    cout << "Enter length of loan term (in months): ";
    cin >> term;

    // Convert it (APR) to monthly rate (divide it by 12) (MR)
    // also divide it by 100 since the value input is in %
    double MR = APR/12.0/100.0;

    // Next, compute the sales tax you will pay (SalesTax)
    double SalesTax = Cost * SalesTaxRate / 100.0;

    // Use the money left to make a down payment (DownPayment)
    double DownPayment = Cash - SalesTax;

    // Then determine the amount you will borrow (LoanAmount)
    double LoanAmount = Cost - DownPayment;
```

```
// Plug in all of the values in the formula and compute
// the monthly payment (MP)
double MP = (LoanAmount * MR)
           / (1.0 - exp(-term * log(1.0+MR)));

// Also, compute the total cost of the car. (TotalCost)
double TotalCost = SalesTax + DownPayment + MP * term;

// Output all the results
printf ("Here are the details about your new car...\n");
printf ("-----\n");
printf ("\n");
printf ("Money you have saved $%1.2f\n", Cash);
printf ("Cost of the car $%1.2f\n", Cost);
printf ("Sales Tax rate is %1.2f%%\n", SalesTaxRate);
printf ("Sales Tax on the car $%1.2f\n", SalesTax);
printf ("Your down payment will be $%1.2f\n",
        DownPayment);
printf ("You will be borrowing $%1.2f\n", LoanAmount);
printf ("A %2d month loan at %1.2f%% APR\n", term, APR);
printf ("Your monthly payment will be $%1.2f\n", MP);
printf ("Total cost will be $%1.2f\n", TotalCost);
printf ("\n");
return 0;
}
```

Notice that we used %% in the format string to print a percent sign after the APR. When you run the program again (say for a slightly different loan term), you get:

```
Enter the cost of the car: $20000.00
Enter the amount of money you saved: $5500.00
Enter the APR for the loan (in %): 7.25
Enter the length of loan term (in months): 48
Here are the details about your new car...
-----
```

```
Money you have saved $5500.00
Cost of the car $20000.00
Sales Tax rate is 6.00%
Sales Tax on the car $ 1200.00
Your down payment will be $4300.00
You will be borrowing $15700.00
A 48 month loan at 7.25% APR
Your monthly payment will be $377.78
Total cost will be $23633.43
```

You can see that for the same amount if you borrow it for a longer period you can reduce your monthly payments by over \$100 (but you pay about \$700 more in the end).

Decision Making in Computer Programs

Decision making is central to all computer programs. In fact there is a famous theorem in computer science that says that if any programmable device is capable of being programmed to do *sequential execution*, *decision making*, and *repetition*, it is capable of expressing any computable algorithm. This is a very powerful idea that is couched in terms of very simple ideas. Given a problem, if you can describe its solution in terms of a combination of the three execution models then you can get any computer to solve that problem. In this section, we will look at decision making in a classic game playing situation.

Computers have been extensively used in game playing: for playing games against people and other computers. The impact of computers on game playing is so tremendous that these days several manufacturers make game playing consoles with computers in them that are completely dedicated to game playing.

Let us design a simple game that you have most likely played while growing up: Paper, Scissors, Rock!

In this game, two players play against each other. At the count of three each player makes a hand gesture indicating that they have selected one of the three items: paper (the hand is held out flat), scissors (two fingers are extended to designate scissors, or rock (indicated by making a fist). The rules of deciding the winner are simple: if both players pick the same object it is a draw; otherwise, paper beats rock, scissors beat paper, and rock beats scissors. Let us write a computer program to play this game against the computer. Here is an outline for playing this game once:


```
// Computer and player make a selection
// Decide outcome of selections (draw or who won)
// Inform player of decision
```

If we represent the three items in a string array, we can have the computer pick one of them at random by using the random number generation facilities provided in C++. If the items are represented as:

```
string items[] = {"Paper", "Scissors", "Rock"};
```

Then we can select any of the items above as the computer's choice using the expression:

```
// Computer makes a selection
string myChoice = items[<0 or 1 or 2 selected randomly>];
```



The exact name of the game can vary, with the three components appearing in a different order, or with “stone” in place of “rock”. Non-English speakers may know the game by their local words for “rock, paper, scissors”, although it is also known as *Jankenpon* in Japan, *Rochambeau* in France, and in South Africa as *Ching-Chong-Cha...*

From: talkingtotan.wordpress.com/2007/08

That is `items[0]` represents the choice "Paper", `items[1]` represents "Scissors", and `items[2]` is "Rock". We saw, in Chapter 4, that we can generate a random number in any range using the `rand` function from the `cstdlib` library in C++. Thus we can model the computer making a random selection using the following statement:

```
#include <cstdlib>
// Computer makes a selection
string myChoice = items[rand() % 3];
```

Recall that `rand() % n` returns a random number in the range $[0..n-1]$. Thus, `rand() % 3` will return either 0, 1, or 2. We can use the `askQuestion` command in Myro to ask the player to indicate their selection (see Chapter 3).

```
string choices = "Paper, Scissors, Rock";
// Player makes a selection
string yourChoice = askQuestion("Pick an item.", choices);
```

Now that we know how to the computer and player make their selection, we need to think about deciding the outcome. Here is an outline:

```
if both picked the same item then it is a draw
if computer wins then inform the player
if player wins then inform the player
```

Rewriting the above using `if`-statements we get the following first draft:

```
if (myChoice == yourChoice) {
    cout << "It is a draw.\n";
} else if (<myChoice beats yourChoice>) {
    cout << "I win.\n";
} else {
    cout << "You win.\n";
}
```

All we need to do is figure out how to write the condition `<myChoice beats yourChoice>`. The condition has to capture the rules of the game mentioned above. We can encode all the rules in a conditional expression as follows:

```
if ((myChoice == "Paper" && yourChoice == "Rock")
    || (myChoice == "Scissors" && yourChoice == "Paper")
    || (myChoice == "Rock" && yourChoice == "Scissors")) {
    cout << "I win.\n";
} else {
    cout << "You win.\n";
}
```

The conditional expression above captures all of the possibilities that should be examined in order to make the decision (recall that `&&` means “and” and `||` means “or”). Another way of writing the above decision would be to use the following:

```
if (myChoice == "Paper" && yourChoice == "Rock") {
    cout << "I win.\n";
} else if (myChoice == "Scissors" && yourChoice == "Paper") {
    cout << "I win.\n";
} else if (myChoice == "Rock" && yourChoice == "Scissors") {
    cout << "I win.\n";
} else {
    cout << "You win.\n";
}
```

That is each condition is examined in turn until one is found that confirms that the computer wins. If none such condition is true, the `else`-part of the `if`-statement will be reached to indicate that the player won.

Another alternative to writing the decision above is to encode the decision in a function. Let us assume that there is a function `beats` that returns `true` or `false` depending on the choices. We could then rewrite the above as follows:

```
if (myChoice == yourChoice) {
    cout << "It is a draw.\n";
} else if (beats(myChoice, yourChoice)) {
    cout << "I win.\n";
} else {
    cout << "You win.\n";
}
```

Let us take a closer look at how we could define the `beats` function. It needs to return `true` if `myChoice` beats `yourChoice`. So all we need to do is encode the rules of the game described above. Here is a draft of the function:

```
bool beats (string me, string you) {
    // Does me beat you? If so, return true, false otherwise.

    if (me == "Paper" && you == "Rock") {
        // Paper beats rock
        return true;
    } else if (me == "Scissors" && you == "Paper") {
        // Scissors beat paper
        return true;
    } else if (me == "Rock" && you == "Scissors") {
        // Rock beats scissors
        return true;
    } else {
        return false;
    }
}
```

Once again, we have used the `if`-statements in C++ to encode the rules of the game. Now that we have completely fleshed out all the critical parts of the program, we can put them all together as shown below:

```
// A program that plays the game of Paper, Scissors, Rock!
#include <iostream>
#include <string>
#include <cstdlib>
#include "Myro.h"
using namespace std;

bool beats (string me, string you) {
    // Does me beat you? If so, return true, false otherwise.

    if (me == "Paper" && you == "Rock") {
        // Paper beats rock
        return true;
    } else if (me == "Scissors" && you == "Paper") {
        // Scissors beat paper
        return true;
    }
}
```

```
    } else if (me == "Rock" && you == "Scissors") {
        // Rock beats scissors
        return true;
    } else {
        return false;
    }
}

int main() {
    string items[] = {"Paper", "Scissors", "Rock"};
    string choices = "Paper,Scissors,Rock";

    // Play a round of Paper, Scissors, Rock!
    cout << "Lets play Paper, Scissors, Rock!\n";
    cout <<
        "In the window that pops up, make your selection>\n";

    // Computer and Player make their selection...
    // Computer makes a selection
    string myChoice = items[rand() % 3];

    // Player makes a selection
    string yourChoice = askQuestion("Pick an item.", choices);

    // inform Player of choices
    cout << endl;
    cout << "I picked " << myChoice << endl;
    cout << "You picked " << yourChoice << endl;

    // Decide if it is a draw or a win for someone
    if (myChoice == yourChoice) {
        cout << "We both picked the same thing.\n";
        cout << "It is a draw.\n";
    } else if (beats(myChoice, yourChoice)) {
        cout << "Since " << myChoice << " beats "
            << yourChoice << "... \n";
        cout << "I win.\n";
    } else {
        cout << "Since " << yourChoice << " beats "
            << myChoice << "... \n";
        cout << "You win.\n";
    }
}
```

```
    cout << "Thank you for playing. Bye!\n";  
    return 0;  
}
```

A few more output commands were added to make the interaction more natural.

Do This: Implement the Paper, Scissors, Rock program from above and play it several times to make sure you understand it completely. Modify the program above to play several rounds. Also, incorporate a scoring system that keeps track of the number of times each player won and also the number of draws.

Summary

In this chapter you have seen a variety of control paradigms: behavior-based control for robots, writing a simple, yet useful computational program, and writing a simple computer game. Behavior-based control is a powerful and effective way of describing sophisticated robot control programs. This is the paradigm used in many commercial robot applications. You should try out some of the exercises suggested below to get comfortable with this control technique. The other two programs illustrate how, using the concepts you have learned, you can design other useful computer applications. In the next several chapters, we will explore several other dimensions of computing: media computation, writing games, etc.

Myro Review

There was nothing new from Myro in this chapter.

C++ Review

```
struct <structure name> { <declaration>; <declaration>; ... };
```

Defines a structure called `<structure name>` with *members* declared by the `<declarations>`s (usually variable and function declarations).

```
<structure name> <variable name>, <variable name>, ...;
```

Defines variables belonging to a structured type. Each variable has the members declared in the preceding `<structure name>` definition.

```
<structure>.<variable name>  
<structure>.<function name> ( <arguments> )
```

Accesses named member-variable (field) or member-function of a `<structure>`, which must have a structured type that declares the referenced member variable or function.

```
printf (<format string>, <expression>, ... );
```

Prints the `<expression>`s according to the `<format string>`. Within it, `%W.Df` prints a floating-point number in a field of minimum width `W` with `D` digits after the decimal point, `%Wd` prints an integer in a field of minimum width `W`, and `%Ws` prints a string in a field of minimum width `W`. The value is right-justified in the field, unless a minus sign follows the `%` (e.g., `%-W.Df`, `%-Wd`, `%-Ws`), in which case it's left-justified. There are many other format specifications besides these. In a `<format string>`, `%%` prints a single percent sign.

The math library module provides several useful mathematics functions. Some of the commonly used functions are listed below:

ceil(x) Returns the ceiling of `x` as a float, the smallest integer value greater than or equal to `x`.

floor(x) Returns the floor of x as a float, the largest integer value less than or equal to x.

exp(x) Returns e^x .

log(x[, base]) Returns the logarithm of x to the given base. If the base is not specified, return the natural logarithm of x (i.e., $\log_e x$).

log10(x) Returns the base-10 logarithm of x (i.e. $\log_{10} x$).

pow(x, y) Returns x^y .

sqrt(x) Returns the square root of x (\sqrt{x}).

Trigonometric functions

acos(x) Returns the arc cosine of x, in radians.

asin(x) Returns the arc sine of x, in radians.

atan(x) Returns the arc tangent of x, in radians.

cos(x) Returns the cosine of x radians.

sin(x) Returns the sine of x radians.

tan(x) Returns the tangent of x radians.

cosh(x) Returns the hyperbolic cosine of x.


sinh(x) Returns the hyperbolic sine of x.

tanh(x) Returns the hyperbolic tangent of x.

Exercises

1. Replace the `avoid` module in the behavior-based control program with a module called: `follow`. The `follow` module tries to detect a wall to the robot's right (or left) and then tries to follow it at all times. Observe the resulting behavior.
2. Do population growth by modeling it as a difference equation.
3. Do population growth by modeling it as a continuous growth function.
4. Can you make the Paper, Scissors, Rock program better?
5. Implement the HiLo game: the computer picks a random number between 0 and 1000 and the user has to try and guess it. With every guess the computer informs the user of their guess was high (Hi), low (Lo), or they got it.
6. Reverse the roles in the HiLo game from above. This time, the user guesses a number between 0 and 1000 and computer program has to guess it. Think about a good strategy for guessing the number based on high and low clues, and then implement that strategy.





Sights & Sounds

*Don't make music for some vast, unseen audience or market or ratings share
or even for something as tangible as money. Though it's crucial to make a
living, that shouldn't be your inspiration. Do it for yourself.*
-Billy Joel

Opposite page: Be My Valentine
Fractal Art by Sven Geier (www.sgeier.net)

We mentioned earlier that the notion of computation these days extends far beyond simple numerical calculations. Writing robot control programs is computation, as is making world population projections. Using devices like iPods you are able to enjoy music, videos, and radio and television shows. Manipulating sounds and images is also the realm of computation and in this chapter we will introduce you to these. You have already seen how, using your robot, you can take pictures of various scenes. You can also take similar images from your digital camera. Using basic computational techniques you have learned so far you will see, in this chapter, how you can do computation on shapes and sound. You will learn to create images using computation. The images you create can be used for visualizing data or even for aesthetic purposes to explore your creative side. We will also present some fundamentals of sound and music and show you how you can also do similar forms of computation using music.

Sights: Drawing¹

If you have used a computer for any amount of time you must have used some kind of a drawing application. Using a typical drawing application you can draw various shapes, color them etc. You can also generate drawings using drawing commands provided in the Myro library module. In order to draw anything you first need a place to draw it: a canvas or a window. You can create such a window using the command:

```
window myCanvas = GraphWin();
```

If you execute the command above in a C++ program, you will immediately see a small gray window pop up (see picture on right). This window will be serving as our canvas for creating drawings. By default, the window created by the `GraphWin` command is 200 pixels high and 200 pixels wide and its name is "Graphics Window". Not a very inspiring way to start, but the

¹ Graphics functions and objects are not implemented in the current C++/Myro API. The following description is subject to change.

`GraphWin` command has some other variations as well. First, in order to make the window go away, you can use the command:

```
myCanvas.close();
```

To create a graphics window of any size and a name that you specify, you can use the command below:

```
window myCanvas = GraphWin
    ("My Masterpiece", 200, 300);
```

The command above creates a window named "My Masterpiece" that will be 200 pixels wide and 300 pixels tall (see picture on right). You can change the background color of a graphics window as shown below:

```
myCanvas.setBackground("white");
```


You can name any of a number of colors in the command above, ranging from mundane ones like "red", "blue", "gray", "yellow", to more exotic colors ranging from "AntiqueWhite" to "LavenderBlush" to "WhiteSmoke". Colors can be created in many ways as we will see below. Several thousand color names have been pre-assigned (Google: color names list) that can be used in the command above.

Now that you know how to create a canvas (and make it go away) and even set a background color, it is time to look at what we can draw in it. You can create and draw all kinds of geometrical objects: points, lines, circles, rectangle, and

A Graphics window



My Masterpiece 200x300



even text and images. Depending on the type of object you wish to draw, you have to first *create* it and then draw it. Before you can do this though, you should also know the coordinate system of the graphics window.

In a graphics window with width, *W* and height *H* (i.e *W*×*H* pixels) the pixel (0, 0) is at the top left corner and the pixel (199, 299) will be at the bottom right corner. That is, x-coordinates increase as you go right and y-coordinates increase as you go down.

The simplest object that you can create is a *point*. This is done as follows:

```
Point p = Point(100, 50);
```

That is, *p* is an object that is a `Point` whose x-coordinate is at 100 and y-coordinate is at 50. This only creates a `Point` object. In order to draw it, you have to issue the command:

```
p.draw(myCanvas);
```

The syntax of the above command may seem a little strange at first. You saw it briefly when we presented lists in Chapter 5 but we didn't dwell on it. Certainly it is different from what you have seen so far. But if you think about the objects you have seen so far: numbers, strings, etc. Most of them have standard operations defined on them (like +, *, /, etc.). But when you think about geometrical objects, there is no standard notation. Programming languages like C++ provide facilities for modeling any kind of object and the syntax we are using here is standard syntax that can be applied to all kinds of objects. The general form of commands issued on objects is:

```
<object>.<function>(<parameters>)
```

Thus, in the example above, `<object>` is the name `p` which was earlier defined to be a `Point` object, `<function>` is `draw`, and `<parameters>` is `myCanvas`. The `draw` function requires the graphics window as the parameter. That is, you are asking the point represented by `p` to be drawn in the window specified as its parameter. The `Point` objects have other functions available:

```
cout << p.getX() << endl;
```

```
100
```

```
cout << p.getY() << endl;
```

```
50
```

That is, given a `Point` object, you can get its x- and y-coordinates. Objects are created using their *constructors* like the `Point(x, y)` constructor above. We will use lots of constructors in this section for creating the graphics objects. A line object can be created similar to point objects. A line requires the two end points to be specified. Thus a line from (0, 0) to (100, 200) can be created as:

```
Line L = Line(Point(0,0),
              Point(100,200));
```

And you can draw the line using the same draw command as above:


```
L.draw(myCanvas);
```

The picture on the right shows the two objects we have created and drawn so far. As for `Point`, you can get the values of a line's end points:

```
Point start = L.getP1();
cout << start.getX() << endl;
```


```
0
```

```
Point end = L.getP2();
cout << end.getY() << endl;
```



Here is a small C++ loop that can be used to create and draw several lines:

```
for (int n = 0; n < 200; n = n+5) {  
    Line L=Line(Point(n,25),  
                Point(100,100));  
    L.draw(myCanvas);  
}
```



In the loop above (the results are shown on the right), the value of n starts at 0 and increases by 5 after each iteration all the way up to but not including 200 (i.e. 195). For each value of n a new `Line` object is created with starting coordinates $(n, 25)$ and end point at $(100, 100)$.

Do This: Try out all the commands introduced so far. Then observe the effects produced by the loop above. Change the increment 5 in the loop above to different values (1, 3, etc.) and observe the effect. Next, try out the following loop:

```
for (int n = 0; n < 200; n = n+5) {  
    Line L = Line(Point(n, 25), Point(100, 100));  
    L.draw(myCanvas);  
    wait(0.3);  
    L.undraw();  
}
```

The `undraw` function does exactly as the name implies. In the loop above, for each value that n takes, a line is created (as above), drawn, and then, after a wait of 0.3 seconds, it is erased. Again, modify the value of the increment and observe the effect. Try also changing the amount of time in the `wait` command.

You can also draw several geometrical shapes: circles, rectangles, ovals, and polygons. To draw a circle, (or any geometrical shape), you first create it and then draw it:


```
Circle C = Circle(centerPoint, radius);  
c.draw(myCanvas);
```

`centerPoint` is a `Point` object and `radius` is specified in pixels. Thus, to draw a circle centered at (100, 150) with a radius of 30, you would do the following commands:

```
Circle C = Circle(Point(100, 150), 30);  
c.draw(myCanvas);
```

Rectangles and ovals are drawn similarly (see details at the end of the chapter). All geometrical objects have many functions in common. For example, you can get the center point of a circle, a rectangle, or an oval by using the command:

```
centerPoint = C.getCenter();
```

By default, all objects are drawn in black. There are several ways to modify or specify colors for objects. For each object you can specify a color for its outline as well as a color to fill it with. For example, to draw a circle centered at (100, 150), radius 30, and outline color red, and fill color yellow:

```
Circle C = Circle(Point(100, 150), 30);  
C.draw(myCanvas);  
C.setOutline("red");  
C.setFill("yellow");
```

By the way, `setFill` and `setOutline` have the same effect on `Point` and `Line` objects (since there is no place to fill any color). Also, the line or the outline drawn is always 1 pixel thick. You can change the thickness by using the command `setWidth(<pixels>):`

```
C.setWidth(5);
```

The command above changes the width of the circle's outline to 5 pixels.

Do This: Try out all the commands introduced here. Also, look at the end of the chapter for details on drawing other shapes.

Earlier, we mentioned that several colors have been assigned names that can be used to select colors. You can also create colors of your own choosing by specifying their red, green, and blue values. In Chapter 5 we mentioned that each color is made up of three values: RGB or red, green and blue color values. Each of these values can be in the range 0..255 and is called a 24-bit color value. In this way of specifying colors, the color with values (255, 255, 255) (that is red = 255, green = 255, and blue = 255) is white; (255, 0, 0) is pure red, (0, 255, 0), is pure blue, (0, 0, 0) is black, (255, 175, 175) is pink, etc. You can have as many as 256x256x256 colors (i.e. over 16 million colors!). Given specific RGB values, you can create a new color by using the command, `color_rgb`:

```
color_rgb myColor = color_rgb(255, 175, 175);
```

Do This: The program below draws several circles of random sizes with random colors. Try it out and see its outcome. A sample output screen is shown on the right. Modify the program to input a number for the number of circles to be drawn. In Chapter 4, Exercise 8, you programmed `randrange(m, n)` to return a random number in range `[m..n-1]`.

```
// Program to draw a bunch of # random colored circles
#include <cstdlib>
#include "Myro.h"
using namespace std;

void makeCircle(double x, double y, double r) {
    // creates a Circle centered at point (x, y) of radius r
    return Circle(Point(x, y), r);
}

color_rgb makeColor() {
    // creates a new color using random RGB values
    int red = randrange(0, 256);
    int green = randrange(0, 256);
    int blue = randrange(0, 256);
}
```

```


    return color_rgb(red, green,blue);
}

int main() {
    // Create and display a
    // graphics window
    int width = 500;
    int height = 500;
    window myCanvas = GraphWin
        ("Circles",width,height);
    myCanvas.
        setBackground("white");

    // draw a bunch of random
    // circles with random
    // colors.
    int N = 500;
    for (int i = 0;
        i < 500; i++) {
        // pick random center
        // point and radius
        // in the window
        int x = randrange(0,width);
        int y = randrange(0,height);
        int r = randrange(5, 25);
        int c = makeCircle(x, y, r);
        // select a random color
        c.setFill(makeColor());

        c.draw(myCanvas);
    }
    return 0;
}

```



Notice our use of functions to organize the program. From a design perspective, the two functions `makeCircle` and `makeColor` are written differently. This is just for illustration purposes. You could, for instance, define `makeCircle` just like `makeColor` so it doesn't take any parameters and generates the values of `x`, `y`, and radius as follows:

```
Circle makeCircle() {  
    // creates a Circle centered at point (x, y) of radius r  
    int x = randrange(0,width);  
    int y = randrange(0,height);  
    int r = randrange(5, 25);  
  
    return Circle(Point(x, y), r);  
}
```

Unfortunately, as simple as this change seems, the function is not going to work. In order to generate the values of `x`, and `y` it needs to know the width and height of the graphics window. But width and height are defined in the function `main` and are not available or accessible in the function above. This is an issue of *scope* of names in a C++ program: what is the scope of accessibility of a name in a program?

C++ defines the scope of a name in a program *textually* or *lexically*. That is, any name is visible in the text of the program/function *after* it has been defined. Note that the notion of *after* is a textual notion. Moreover, C++ restricts the accessibility of a name to the text of the function in which it is defined. That is, the names `width` and `height` are defined inside the function `main` and hence they are not visible anywhere outside of `main`. Similarly, the variables `red`, `green`, and `blue` are considered local to the definition of `makeColor` and are not accessible outside of the function, `makeColor`.

So how can `makeCircle`, if you decided it would generate the `x` and `y` values relative to the window size, get access to the width and height of the window? There are two solutions to this. First, you can pass them as parameters. In that case, the definition of `makeCircle` will be:

```
Circle makeCircle(int w, int h) {  
    // creates a Circle centered at point (x, y) of radius r  
    // such that (x, y) lies within width, w and height, h  
    int x = randrange(0,w);  
    int y = randrange(0,h);  
    int r = randrange(5, 25);  
    return Circle(Point(x, y), r);  
}
```

Then the way you would use the above function in the main program would be using the command:

```
Circle C = makeCircle(width, height);
```

That is, you pass the values of `width` and `height` to `makeCircle` as parameters. The other way to define `makeCircle` would be exactly as shown in the first instance:

```
Circle makeCircle() {
    // creates a Circle centered at point (x, y) of radius r
    int x = randrange(0,width);
    int y = randrange(0,height);
    int r = randrange(5, 25);

    return Circle(Point(x, y), r);
}
```

However, you would move the definitions of `width` and `height` outside and before the definitions of all the functions:

```
#include <cstdlib>
#include "Myro.h"
using namespace std;

int width = 500;
int height = 500;

Circle makeCircle() {
    ...

color_rgb makeColor() {
    ...

int main() {
    ...
```

Since the variables are defined outside of any function and before the definitions of the functions that use them, you can access their values. You

may be wondering at this point, which version is better? Or even, why bother? The first version was just as good. The answer to these questions is similar in a way to writing a paragraph in an essay. You can write a paragraph in many ways. Some versions will be more preferable than others. In programming, the rule of thumb one uses when it comes to the scope of names is: ensure that only the parts of the program that are supposed to have access to a name are allowed access. This is similar to the reason you would not share your password with anyone, or your bank card code, etc. In our second solution, we made the names `width` and `height` *globally* visible to the entire program that follows. This implies that even `makeColor` can have access to them whether it makes it needs it or not.

You may want to argue at this point: what difference does it make if you make those variables visible to `makeColor` as long as you take care not to use them in that function? You are correct, it doesn't. But it puts an extra responsibility on your part to ensure that you will not do so. But what is the guarantee that someone who is modifying your program chooses to?

We used the simple program here to illustrate simple yet potentially hazardous decisions that dot the landscape of programming like land mines. Programs can crash if some of these names are *mishandled* in a program. Worse still, programs do not crash but lead to incorrect results. However, at the level of deciding which variables to make local and which ones to make global, the decisions are very simple and one just needs to exercise safe programming practices. We will conclude this section of graphics with some examples of creating animations.

Any object drawn in the graphics window can be moved using the command `move(dx, dy)`. For example, you move the circle 10 pixels to the right and 5 pixels down you can use the command:

```
C.move(10, 5);
```

Do This: Let us write a program that moves a circle about (randomly) in the graphics window. First, enter this following program and try it out.

```
// Moving circle; Animate a circle...
#include <cstdlib>
#include "Myro.h"
using namespace std;

int main() {
    // create and draw the graphics window
    window w = GraphWin("Moving Circle", 500, 500);
    w.setBackground("white");

    // Create a red circle
    Circle c = Circle(Point(250, 250), 50);
    c.setFill("red");
    c.draw(w);

    // Do a simple animation for 200 steps
    for (int i = 0; i < 200; i++) {
        c.move(randrange(-4, 5), randrange(-4, 5));
        wait(0.2);
    }
    return 0;
}
```

Notice, in the above program, that we are moving the circle around randomly in the x- and y-directions. Try changing the range of movements and observe the behavior. Try changing the values so that the circle moves only in the horizontal direction or only in the vertical direction. Also notice that we had to *slow down* the animation by inserting the wait command after every move. Comment the wait command and see what happens. It may appear that nothing did happen but in fact the 200 moves went so quickly that your eyes couldn't even register a single move! Using this as a basis, we can now write a more interesting program. Look at the program below:

```
// Moving circle; Animate a circle...
#include <cstdlib>
#include "Myro.h"
using namespace std;
```

```
int main() {
    // create and draw the graphics window
    int winWidth = 500, winHeight = 500;
    window w = GraphWin("Bouncing Circle",
        winWidth, winHeight);
    w.setBackground("white");

    // Create a red circle
    int radius = 25;
    Circle c = Circle(Point(53, 250), radius);
    c.setFill("red");
    c.draw(w);

    // Animate it
    int dx = 3, dy = 3;
    while (timeRemaining(15)) {
        // move the circle
        c.move(dx, dy);

        // make sure it is within bounds
        Point center = c.getCenter();
        int cx = center.getX();
        int cy = center.getY();

        if ((cx+radius >= winWidth) || (cx-radius <= 0))
            dx = -dx;

        if ((cy+radius >= winHeight) || (cy-radius <= 0))
            dy = -dy;

        wait(0.01);
    }
    return 0;
}
```

For 15 seconds, you will see a red circle bouncing around the window. Study the program above to see how we keep the circle inside the window at all times and how the direction of the ball bounce is being changed. Each time you change the direction, make the computer beep:

```
computer.beep(0.005, 440);
```


If you are excited about the possibility of animating a circle, imagine what you can do if you have many circles and other shapes animated. Also, plug in the game pad controller and see if you can control the circle (or any other object) using the game pad controls. This is very similar to controlling your robot. Design an interactive computer game that takes advantage of this new input modality. You can also design multi-user games since you can connect multiple game pad controllers to your computer. See the Reference documentation for details on how to get input from several game pad controllers.

Drawing Text & Images

Like shapes, you can also place text in a graphics window. The idea is the same. You first create the text using the command:

```
Text myText = Text(<anchor point>, <string>);
```

and then draw it. You can specify the type face, size, and style of the text. We will not detail these here. They are summarized in the reference at the end of the text. When we get an opportunity, we will use these features below in other examples.

Images in this system are treated just like any other objects. You can create an image using the `Image` command:

```
Image myPhoto = Image(<centerPoint>, <file name>);
```

You have to have an already prepared image in one of the common image formats (like JPEG, GIF, etc.) and stored in a file (`<file name>`). Once the image object is created, it can be drawn, undrawn, or moved just like other shapes.

What you do with all this new found functionality depends on your creativity. You can use graphics to visualize some data: plotting the growth of world population, for example; or create some art, an interactive game, or even an

animated story. You can combine your robot, the game pad controller, or even sound to enrich the multi-media experience. The exercises at the end of the chapter present some ideas to explore further. Below, we delve into sounds.

Sound

As we have seen, you can have your robot make *beeps* by calling the `beep()` function, like so:

```
robot.beep(1, 440);
```

This command instructs the robot to play a tone at 440 Hz for 1 second. Let us first try and analyze what is in the *440 Hz* tone. First, the letters *Hz* are an abbreviation for *Hertz*. The name itself comes from a German physicist, Heinrich Rudolph Hertz who did pioneering work in the production of electromagnetic waves in the late 19th century. Today, we use **Hertz** (or *Hz*) as a unit for specifying frequencies.

$$1\text{Hertz} = 1\text{cycle} / \text{second}$$

The most common use of frequencies these days is in specifying the clock speeds of computer CPU's. For example, a typical PC today runs at clock speeds of a few GigaHertz (or GHz).

$$1\text{GigaHertz} = 10^9\text{cycles} / \text{second}$$

Frequencies are related to periodic (or repetitive) motions or vibrations. The time it takes for a motion or vibration to repeat is called its *time period*. Frequency and time period are inversely related. That is the number of cycles or repetitions in a second is called the frequency. Thus 1 Hertz refers to any motion or vibration that repeats every 1 second. In the case of computer clock frequencies then, a computer running at 4 Gigahertz is repeating 4 billion times a second! Other examples of periodic motions include: the earth's rotation on its axis (1 cycle every $24 * 60 * 60 = 86400$ seconds or at a frequency of 0.00001157 cycles/second), a typical audio CD spins 400 times a

second, a CD drive on your computer rated at 52x spins the CD at $52 * 400 = 20800$ times per second, hummingbirds can flap their wings at frequencies ranging from 20-78 times/second (some can go even as high as 200!).

Sound is a periodic compression and refraction (or return to its original state) of air (for simplicity, let us assume that the medium is air). One Cycle of a sound comprises one compression and one refraction. Thus, producing a beep at 440 Hz represents 440 complete cycles of compression and refraction. Generally, a human ear is capable for hearing frequencies in the 20 Hz to 20000 Hz (or 20 Kilo Hertz) range. However the capability varies from person to person. Also, many electronic devices are not capable for producing frequencies in that entire range. 20-20KHz is considered hi-fidelity for stereo or home theater audio components. Let us first examine the range of audible sounds the Scribbler can produce. To make a sound out of the Scribbler, you have to give a frequency and the duration (in seconds) that the sound should be played. For example, to play a 440 Hz sound for 0.75 seconds:

```
robot.beep(0.75, 440);
```

The human ear is capable of distinguishing sounds that differ only by a few Hertz (as little as 1 Hz) however this ability varies from person to person. Try the commands:

```
robot.beep(1, 440);  
robot.beep(1, 450);
```

Can you distinguish between the two tones? Sometimes it is helpful to place these in a loop so you can repeatedly hear the alternating tones to be able to distinguish between them. The next exercise can help you determine what frequencies you are able to distinguish.

Exercise: Using the example above, try to see how close you can get in distinguishing close frequencies. As suggested, you may want to play the tones alternating for about 5-10 seconds. Start with 440 Hz. Can you hear the difference between 440 and 441? 442? Etc. Once you have established your

range, try another frequency, say 800. Is the distance that you can distinguish the same?

Do This: You can program the Scribbler to create a siren by repeating two different tones (much like in the example above). You will have to experiment with different pairs of frequencies (they may be close together or far apart) to produce a realistic sounding siren. Write your program to play the siren for 15 seconds. The louder the better!

You can also have Myro make a beep directly out of your computer, rather than the robot, with the command:

```
computer.beep(1, 440);
```

Unfortunately, you can't really have the robot and computer play a duet. Why not? Try these commands:

```
robot.beep(1, 440);  
computer.beep(1, 440);  
robot.beep(1, 880);  
computer.beep(1, 880);  
robot.beep(1, 440);  
computer.beep(1, 440);
```

What happens? Try your solutions to the above exercises by making the sounds on the computer instead of the Scribbler.

Musical Scales

In western music, a *scale* is divided into 12 notes (from 7 major notes: ABCDEFG). Further there are *octaves*. An octave in C comprises of the 12 notes shown below:

C C#/Db D D#/Eb E F F#/Gb G G#/Ab A A#/Bb B

C# (pronounced "C sharp") is the same tone as Db (pronounced "D flat").

Frequencies corresponding to a specific note, say C, are multiplied (or divided) by 2 to achieve the same note in a higher (or lower) octave. On a piano there are several octaves available on a spread of keys. What is the relationship between these two tones?

```
robot.beep(1, 440);
robot.beep(1, 880);
```

The second tone is exactly one octave the first. To raise a tone by an octave, you simply multiply the frequency by 2. Likewise, to make a tone an octave lower, you divide by 2. Notes indicating an octave can be denoted as follows:

C0 C1 C2 C3 C4 C5 C6 C7 C8

That is, C0 is the note for C in the lowest (or 0) octave. The fifth octave (numbered 4) is commonly referred to as a middle octave. Thus C4 is the C note in the middle octave. The frequency corresponding to C4 is 261.63 Hz. Try playing it on the Scribbler. Also try C5 (523.25) which is twice the frequency of C4 and C3 (130.815). In common tuning (*equal temperament*) the 12 notes are equidistant. Thus, if the frequency doubles every octave, each successive note is $2^{1/12}$ apart. That is, if C4 is 261.63 Hz, C# (or Db) will be:

$$C\#/Db4 = 261.63 * 2^{1/12} = 277.18$$

Thus, we can compute all successive note frequencies:

$$D4 = 277.18 * 2^{1/12} = 293.66$$

$$D\#/Eb = 293.66 * 2^{1/12} = 311.13$$

etc.

The lowest tone that the Scribbler can play is A0 and the highest tone is C8. A0 has a frequency of 27.5 Hz, and C8 has a frequency of 4186 Hz. That's quite a range! Can you hear the entire range?

```
robot.beep(1, 27.5);  
robot.beep(1, 4186);
```

Exercise: Write a Scribbler program to play all the 12 notes in an octave using the above computation. You may assume in your program that C0 is 16.35 and then use that to compute all frequencies in a given octave (C4 is $16.35 * 2^4$). Your program should input an octave (a number from 0 through 8), produce all the notes in that octave and also printout a frequency chart for each note in that octave.

Making Music

Playing songs by frequency is a bit of a pain. Myro contains a set of functions to make this task a bit more abstract.² A Myro song is a string of characters composed like so:

```
NOTE1 [NOTE2] WHOLEPART
```

where [] means optional. Each of these notes/chords is composed on its own line, or separated by semicolons where:

```
NOTE1 is either a frequency or a NOTENAME  
NOTE2 is the same, and optional. Use for Chords.  
WHOLEPART is a number representing how much of  
a whole note to play.
```

NOTENAMES are case-insensitive strings. Here is an entire scale of NOTENAMES:

```
C C#/Db D D#/Eb E F F#/Gb G G#/Ab A A#/Bb B C
```

This is the default octave. It is also the 5th octave, which can also be written as:

```
C5 C#5/Db5 D5 D#5/Eb5 E5 F5 F#5/Gb5 G5 G#5/Ab5 A5 A#5/Bb5 B5  
C6
```

² These functions are not implemented in the current C++/Myro interface.

The Myro Song Format replicates the keys on the piano, and so goes from A0 to C8. The middle octave on a keyboard is number 4, but we use 5 as the default octave. See http://en.wikipedia.org/wiki/Piano_key_frequencies for additional details. Here is a scale:

```
"C 1; C# 1; D 1; D# 1; E 1; F 1; F# 1; G 1; G# 1; A 1; A# 1; B 1; C 1;"
```

The scale, one octave lower, and played as a polka:

```
"C4 1; C#4 1/2; D4 1/2; D#4 1; E4 1/2; F4 1/2; F#4 1; G4 1/2; G#4 1/2; A4 1; A#4 1/2; B4 1/2; C4 1;"
```

There are also a few other special note names, including PAUSE, REST, you can leave the octave number off of the default octave notes if you wish. Use "#" for sharp, and "b" for flat.

WHOLEPART can either be a decimal notation, or division. For example:

```
Ab2 .125
```

OR

```
Ab2 1/8
```

represents the A flat in the second octave (two below middle).

As an example, try playing the following:

```
c 1
c .5
c .5
c 1
c .5
c .5
e 1
c .5
c .5
```

```
c 2
e 1
e .5
e .5
e 1
e .5
e .5
g 1
e .5
e .5
e 2
```

Do you recognize it??

You may leave blank lines, and comments should begin with a # sign. Lines can also be separated with a semicolon.

Using a song

For the following exercises, you will need to have an object to play the song. Now that you have a song, you probably will want to play it. If your song is in a file, you can read it:³

```
song s = readSong(filename);
```

and play it on the robot:

```
robot.playSong(s);
```

or on the computer:

```
computer.playSong(s);
```

You can also use `makeSong(text)` to make a song. For example:

```
song s = makeSong
```

³ These functions are not implemented in the current C++/Myro interface.


```
("c 1; d 1; e 1; f 1; g 1; a 1; b 1; c7 1;");
```

and then play it as above.

If you want to make it play faster or slower, you could change all of the `WHOLENOTE` numbers. But, if we just want to change the tempo, there is an easier way:

```
robot.playSong(s, .75);
```

The second argument to `playSong` is the duration of a whole note in seconds. Standard tempo plays a whole note in about .5 seconds. Larger numbers will play slower, and smaller numbers play faster.

Summary

You can use the graphics window as a way of visualizing anything. In the graphics window you can draw all kinds of shapes: points, line, circles, rectangles, ovals, polygons, text, and even images. You can also animate these shapes as you please. What you can do with these basic drawing capabilities is limited only by your creativity and your programming ability. You can combine the sights you create with sounds and other interfaces, like the game pad controller, or even your robot. The multimedia functionalities introduced in this chapter can be used in all kinds of situations for creating interesting programs.

Myro Reference

Below, we summarize all of the graphics commands mentioned in this chapter. You are urged to review the reference manual for more graphics functionality.

```
GraphWin()  
GraphWin(<title>, <width>, <height>)
```

Returns a `graphics window` object. It creates a graphics window with title,

`<title>` and dimensions `<width> x <height>`. If no parameters are specified, the window created is 200x200 pixels.

```
<window>.close();
```

Closes the displayed graphics window `<window>`.

```
<window>.setBackground(<color>);
```

Sets the background color of the window to be the specified color. `<color>` can be a named color (Google: color names list), or a new color created using the `color_rgb` command (see below).

```
color_rgb(<red>, <green>, <blue>)
```

Creates a new RGB color object using the specified `<red>`, `<green>`, and `<blue>` values. The values can be in the range 0..255.

```
Point(<x>, <y>)
```

Creates a `Point` object at `(<x>, <y>)` location in the window.

```
<point>.getX()
```

```
<point>.getY()
```

Returns the x and y coordinates of the `Point` object `<point>`.

```
Line(<start point>, <end point>)
```

Creates a `Line` object starting at `<start point>` and ending at `<end point>`.

```
<line>.getP1()
```

```
<line>.getP2()
```

Returns an endpoint (`Point` object) of a `Line` object.

```
Circle(<center point>, <radius>)
```

Creates a `Circle` object centered at `<center point>` with radius `<radius>` pixels.

```
Rectangle(<point1>, <point2>)
```

Creates a `Rectangle` object with opposite corners located at `<point1>` and `<point2>`.

Oval(<point1>, <point2>)

Creates an Oval object in the bounding box defined by the corner points <point1> and <point2>.

Polygon(<point1>, <point2>, <point3>,...)

Polygon([<point1>, <point2>, ...])

Creates a Polygon object with the given points as vertices.

Text(<anchor point>, <string>)

Creates a Text object anchored (bottom-left corner of text) at <anchor point>. The text itself is defined by <string>.

Image(<centerPoint>, <file name>)

Creates an Image centered at <center point> from the image file <file name>. The image can be in GIF, JPEG, or PNG format.

All of the graphics objects respond to the following commands:

<object>.draw(<window>);

Draws the <object> in the specified graphics window <window>.

<object>.undraw();

Undraws <object>.

<object>.getCenter();

Returns the center point of the <object>.

<line-object>.getP1();

<line-object>.getP2();

Returns the end points of the <line-object>.

<object>.setOutline(<color>);

<object>.setFill(<color>);

Sets the outline and the fill color of the <object> to the specified <color>.

<object>.setWidth(<pixels>);

Sets the thickness of the outline of the <object> to <pixels>.

```
<object>.move (<dx>, <dy>);
```

Moves the object `<dx>`, `<dy>` from its current position.

The following sound-related functions were presented in this chapter.

```
<robot/computer object>.beep (<seconds>, <frequency>);
```

```
<robot/computer object>.beep (<seconds>, <f1>, <f2>);
```

Makes the robot or computer beep for `<seconds>` time at frequency specified.

You can either specify a single frequency `<frequency>` or a mix of two: `<f1>` and `<f2>`.

```
robot.playSong (<song>);
```

Plays the `<song>` on the robot.

```
readSong (<filename>)
```

Reads a song object from `<filename>`.

```
song2text (song)
```

Converts a `song` to text format.

```
makeSong (<text>)
```

```
text2song (<text>)
```

Converts `<text>` to a `song` format.


C++ reference

In this chapter we presented informal *scope rules* for names in C++ programs. While these can get fairly complicated, for our purposes you need to know the distinction between a *local name* that is local within the scope of a function versus a *global name* defined outside of the function. The text ordering defines what is accessible.

Exercises


1. Using the graphics commands learned in this chapter, write a program to generate a seasonal drawing: winter, summer, beach, or a festival scene.

2. Write a program that has a function `drawRobot(x, y, w)` such that it draws the figure of a robot shown on the right. The robot is anchored at (x, y) and is drawn in the window, w . Also notice that the robot is entirely made up of rectangles (8 of them).



3. Using the `drawRobot` function from the previous exercise, draw a pyramid of robot as shown on the right.

4. Suppose that each robot from the previous exercise is colored using one of the colors: red, blue, green, and yellow. Each time a robot is drawn, it uses the next color in the sequence.



Once the sequence is exhausted, it recycles. Write a program to draw the above pyramid so robots appear in these colors as they are drawn. You may decide to modify `drawRobot` to have an additional parameter: `drawRobot(x, y, c, w)` or you can write it so that `drawRobot` decides which color it chooses. Complete both versions and compare the resulting programs. Discuss the merits and/or pitfalls of these versions with friends. [Hint: Use a vector of color names.]

5. The `beep` command can play two tones simultaneously. In this exercise we will demonstrate an interesting phenomena. First, Try the following commands:


```
robot.beep(2, 660);  
robot.beep(2, 665);
```

You may or may not be able to tell the two tones apart. Next, try playing the two tones together using the command:

```
robot.beep(2, 660, 665);
```

You will hear pulsating beats. This phenomenon called *beating* is very common when you play pure tone together. Explain why this happens.

6. Do a web search on *fractals* and write programs to draw some simple fractals. For example, *Koch Snowflakes*, *Sierpinski Triangles*, *Mandelbrot Sets*, *Julia Sets*, etc.



9

Image Processing & Perception

Seeing is believing.
Proverb

Opposite page: Rotating Snakes (A Visual Illusion)
Created by Akiyoshi Kitaoka (www.ritsumei.ac.jp/~akitaoka/index-e.html)

Your robot has a small digital camera that can be used to take pictures. A picture taken by a digital camera is represented as an *image*. As you have seen in the previous chapter images can be drawn and moved about in a graphics window just as if it were any other graphics object (like a line, circle, etc.). We also saw in Chapter 5 how an image taken from the Scribbler's camera can be viewed on your monitor in a separate window. In this chapter we will learn how to do computing on images. We will learn how images are represented and how we can create them via computation and also process them in many different ways. The representation of images we will use is same as those used by most digital cameras and cell phones and the same representation can be used to display them in a web page. We will also learn how an image taken by the robot's camera can be used to serve as the camera's eye into its world using some image understanding techniques. Image understanding is the first step in visual perception. A robot equipped with even the most rudimentary visual perception capabilities can be designed to carry out more interesting behaviors.

What is an Image?


In Myro you can issue a command for the robot to take a picture and display it on the computer's monitor using the commands:

```
picture pic = robot.takePicture();  
show(pic);
```

The picture on the next page shows an example image taken from the Scribbler's camera. An image is made up of several tiny picture elements or *pixels*. In a color image, each pixel contains color information which is made up of the amount of red, green, and blue (also called, RGB) values. Each of these values can be in the range [0..255] and hence it takes 3 bytes or 24 bits to store the information contained in a single pixel. A pixel that is colored pure red will have the RGB values (255, 0, 0). A grayscale image, on the other hand, only contains the level of gray in a pixel which can be represented in a single byte (8 bits) as a number ranging from 0..255 where 0 is black and 255 is white. The entire image is just a 2-dimensional array of pixels. For example, the images obtained from the Scribbler have 256x192 (WxH) or a

total of 49,152 pixels. Since each pixel requires 3 bytes of data, this image has a size of 147,456 bytes.

All digital cameras are sold by specifying the number of megapixels. For example, the camera shown below is rated at 6.3 megapixels. This refers to the size of the largest image that it can take. The more pixels in an image the better the image resolution can be when it is printed. With a 6.3 megapixel image you will be able to create good quality prints as large as 13x12 inches (or even larger). By comparison, a conventional 35mm photographic film has roughly 4000x3000 or 12 million pixels. Professional digital cameras easily surpass this resolution and that is why, in the last decade or so, we have seen a rapid decline in film-based photography. For displaying sharp images on a computer you need less than half the resolution offered by the camera shown here. The Scribbler camera, with an image size of 147,456 bytes is only about 0.14 megapixels. While this is low resolution it is sufficient for doing visual perception for the robot.



To make electronic storage and transfer of images (web, e-mail, etc.) fast and convenient the data in an image can be compressed.

Several formats are available to electronically store images: JPEG, GIF, PNG, etc. JPEG is the most common format used by digital cameras, including the Scribbler's. JPEG enables excellent compression coupled with a wider range of colors compared with the GIF format making it useful for most image applications. Myro supports both JPEG and GIF image formats. When we intend to process an image, we will always use the JPEG format. We will primarily use the GIF format for creating animated images.

Myro Image Basics¹

After you take a picture with the Scribbler as above you can get some information about the size of the picture with the `getWidth()` and `getHeight()` functions:

```
int picWidth = getWidth(pic);
int picHeight = getHeight(pic);
cout << "Image WxH is " << picWidth << "x" << picHeight
      << " pixels.\n";
```

If you wish to save the image for later use, you can use the Myro command:

```
savePicture(pic, "OfficeScene.jpg");
```

The file `OfficeScene.jpg` will be saved in the current folder. The `.jpg` extension signals the command to save the image as a JPEG image. If you wanted to save it as a GIF image, you can use the `.gif` extension as shown below:

```
savePicture(pic, "OfficeScene.gif");
```

Later, you can load the picture from disk with the `makePicture()` function:

```
picture mySavedPicture = makePicture("OfficeScene.jpg");
show(mySavedPicture);
```

A nice command combination that allows you to navigate and then select the image to load is:

```
picture mySavedPicture = makePicture(pickAFile());
show(mySavedPicture);
```

¹ The C++/Myro interface is under development and might not agree with the description in this chapter.

The `pickAFile` command gives you a navigational dialog box, just like the one you see when you open and select files in many other computer applications. You can navigate to any folder and select a file to open as an image. In fact, you can use the `makePicture` command to load any JPEG picture file regardless of whether it was created from your own digital camera or one you downloaded from the web. Below, we show you how to load a picture and display it:

```
picture lotusTemple = makePicture(pickAFile());
show(lotusTemple, "Lotus Temple");
```

If you move your mouse and click anywhere in the picture, you can also get the x- and y-coordinates of the pixel you selected along with its RGB values. This is shown in the picture on the right. The mouse was clicked on the pixel at location (65, 196) and its RGB values were (168,174,104). Can you guess where that is in the picture?

By the way, the `show` command takes an optional second parameter which is a string that becomes the title of the image window. One advantage of being able to load and view any image is that we can also learn to process or manipulate such images computationally. We will return to image processing later in this chapter.



A Robot Explorer

If you do not need a full color picture, you can tell Myro to capture a gray-scale image by giving the `takePicture()` function the "gray" parameter.

```
picture grayPic = robot.takePicture("gray");
show(grayPic);
```

You will notice that taking the gray-scale picture took less time than taking the color picture. As we explained earlier, a gray scale picture only uses one byte per pixel, instead of three. Because gray-scale images can be transferred from the robot to your computer faster than full color images, they can be useful if you want the images to update quickly. For example, you can use the `joyStick()` function combined with a loop that takes and displays pictures to turn your robot into a remotely piloted explorer, similar to the Mars rovers.

```
joyStick();
for (int i = 1; i <= 25; i++) {
    picture pic = robot.takePicture("gray");
    show(pic);
}
```

The above code will open a joy stick window so that you can control your robot and then capture and show 25 pictures, one after the other. While the pictures are being captured and displayed like a movie, you can use the joystick to drive your robot around, using the pictures to guide it. Of course, if you removed the "gray" parameter from the `takePicture()` function call, you would get color pictures instead of grayscale pictures, but they would take much longer to transfer from the robot to your computer, and make it more difficult to control the robot.

Animated GIF movies

The `savePicture()` function also allows you to make an animated GIF, which is a special type of picture that in a web browser will show several pictures one after another in an animation. To save an animated GIF, you must give the `savePicture()` function a list of pictures (instead of a single picture) and a filename that ends in ".gif". Here is an example:

```
vector<picture> listOfPictures (4);
listOfPictures[1] = robot.takePicture();
robot.turnLeft(0.5,0.25);
```

```

listOfPictures[2] = robot.takePicture();
robot.turnLeft(0.5,0.25);
listOfPictures[3] = robot.takePicture();
robot.turnLeft(0.5,0.25);
listOfPictures[4] = robot.takePicture();

savePicture(listOfPictures, "turningMovie.gif")

```

The best way to view an animated GIF file is to use a web browser. In your favorite browser use the FILE->Open File menu, and then pick the turningMovie.gif file. The web browser will show all frames in the movie, but then stop on the last frame. To see the movie again, press the "Reload" button. You can also use a loop to make a longer movie with more images:

```

vector<picture> pictureList (0); //Start with an empty vector.
for (int i = 1; i <= 15; i++) {
    picture pic = robot.takePicture();
    pictureList.push_back(pic);    //Append the new picture
    robot.turnLeft(0.5,0.1);
}
savePicture(pictureList, "rotatingMovie.gif");

```

The above commands create an animated GIF movie from 15 pictures taken while the robot was rotating in place. This is a nice way to capture a complete scene around the robot


Making Pictures

Since an image is just an array of pixels it is also possible to create or compose your own images by coloring in each individual pixel. Myro provides simple commands to fill in or examine a color or grayscale value in individual pixels. You can use computations to determine what to fill in each pixel. To start with, you can create a blank image as follows:

```

int W = 100, H = 100;
picture newPic = makePicture(W, H, black);
show(newPic);

```



The 100x100 pixel image starts out with all its pixels colored pure black (i.e. RGB = (0,0,0)). If you'd rather like all pixels to be a different color, you can specify its RGB values:

```
picture newPic = makePicture(W, H, makeColor(R,G,B));
```

Alternately, if you ever need to, you can also set all the pixels to any color, say white, using the loop:

```
for (int x = 0; x < W; x++)
    for (int y = 0; y < H; y++)
        Pixel pixel = getPixel(newPic, x, y);
        setColor(pixel, white);
    }
repaint(newPic);
```

The `getPixel` command returns the pixel at specified `x`- and `y`- locations in the picture. `setColor` sets the given pixel's color to any specified color. Above we're using the predefined color white. You can create a new color by specifying its RGB values in the command:

```
color_rgb myRed = makeColor(255, 0, 0);
```

To visually select a color and its corresponding RGB values, you can use the command:

```
color_rgb myColor = pickAColor();
```

A color palette will be displayed from which you can select a color of your choosing. The palette also shows you the chosen color's RGB values. After you select a color and press OK, the value of `myColor` will be the selected color.

The `repaint` command refreshes the displayed image so you can view the changes you made. In the example loop above, we are using it after all the pixels are modified. If you want to see the changes as they are being made, you can include the `repaint` command inside the loop:


```
for (int x = 0; x < W; x++)
  for (int y = 0; y < H; y++)
    Pixel pixel = getPixel(newPic, x, y);
    setColor(pixel, white);
    repaint(newPic);
}
```

You will be able to view each pixel being changed. However, you will also notice that repainting this way takes a considerable amount of time even on the small image we are creating. Thus, it is a good idea to refresh once all the pixels are modified.

In addition to the `setColor` command, Myro also has the `setRGB` command that can be used to set a pixel's color. This command uses RGB values themselves, instead of a color.

```
setRGB(pixel, 255, 0, 0);
```

There is also a corresponding command to get the RGB values of a given pixel:

```
color_rgb rgb = getRGB(pixel);
int R = rgb.red, G = rgb.green, B = rgb.blue;
```

The `getRGB` command returns a `color_rgb` object with instance variables (components) `red`, `blue` and `green`, which can be assigned to individual variables as shown above. Additionally, given a pixel, you get the individual RGB values using the commands `getRed`, `getGreen`, and `getBlue`. These are described in more detail at the end of this chapter and are illustrated in examples below.

Many image creation and processing situations require the processing of every pixel in the image. In the loops above, we are using the `x`- and `y`- variables to reach every pixel in the image. An alternate way of accomplishing the same result is to use the following form of the loop:

```
Pixels PicPixels = getPixels(newPic);
for (Pixels pixel = PicPixels.begin();
     pixel != PicPixels.end(); pixel++)
    setColor(*pixel, gray);
repaint(newPic);
```

The function `getPixels` returns an *iterator* of type `Pixels`, which returns the next pixel to be processed each time around the loop thus guaranteeing that all pixels of the image are processed. `PicPixels.begin()` is a pointer to the first pixel in the vector of pixels, and `PicPixels.end()` is a pointer to the end of the vector. Just to watch how the loop above works you may want to try and put `repaint` inside the loop. The difference between the two methods is that the first loop gets a pixel at the given *x*- and *y*- coordinates, while the latter gives you a pixel at a time without worrying about its *x*- and *y*- values.

Shades of Gray

Using the basic image pixel accessing and modifying commands one can write computations to create interesting and creative images. To introduce you to basic image creation and processing techniques, we will create some images entirely in the grayscale spectrum. In a JPEG image, all shades of gray have equal RGB values. The darkest shade of gray is (0,0,0) or black and the brightest is (255,255,255) or white. Rather than worrying about the triple of RGB values, we can think of just a single value in the range 0..255 to represent the grayscale spectrum. We can write a simple function that will return the RGB shade of gray as follows:

```
color_rgb gray(int v) {
    // returns an rgb gray color for v
    return makeColor(v, v, v);
}
```

Let's create an image that displays the entire range of shades of gray using the following rule:


pixel at x, y is colored using the grayscale x+y

That is, a pixel in the image at 0, 0 will get the shade $0+0=0$ or black, which the pixel at 50,50 will get the shade $50+50=100$ which will be a midrange gray. We can accomplish this using the following loop:

```
int main() {
    int MAX = 255;
    int W = 100;
    int H = 100;

    // Create the blank image
    picture myPic = makePicture(W, H);
    show(myPic, "Shades of Gray");

    // Fill in each pixel with x+y shade of gray
    for (int x = 0; x < W; x++)
        for (int y = 0; y < H; y++)
            setPixel(myPic, x, y, gray((x+y)%(MAX+1)));
}
```



We have used the variable `MAX` above to represent the maximum value of the grayscale range. In the last line we take the remainder $(x+y) \% (MAX+1)$ to ensure that the gray value will be in the range 0..255. The image generated by the above commands is shown on the right. It may seem strange at first that we have used the `x`- and `y`- locations of a pixel to compute its grayscale or brightness value. Since `x`- and `y`- are numbers and since the grayscale value itself is a number it is ok to do this. Recognizing that `x`- and `y`- values (and their sum) might be larger than the range of the grayscale spectrum, we take the remainder to wrap the values around.

Do This: Write a complete program and try this out. Once you have the program, try it out on an image of size 123x123 (why?). Then again, try it out on a much larger image, say 500x500.

The grayscale range basically represents levels of brightness. We can have some more fun by changing the range, say to 0..50. Then we can modify the gray function as shown below to proportionately map any value in the range 0..50 to the range 0..255:

```
int MAX = 50;
color_rgb gray(int v){
    // returns an rgb gray color for v
    int brightness = v*255/MAX;
    return makeColor(brightness, brightness, brightness);
}
```

Do This: Rewrite the program to use the above definition of `gray` and create a 500x500 image. Modify the value of `MAX` to 25 and create a 100x100 image (shown to the right). Note that if you set `MAX` to 255, it will revert to the earlier version of the function.



Do This: Next use the rule:

*pixel at x, y is colored using the grayscale $x*y$*

Remember to take remainder as above. Change `MAX` to 250 and create a 500x500 image. Try other sizes and values of `MAX`.

You can experiment with other rules, geometric transformations, and trigonometric functions. See exercises at the end of this chapter for more suggestions.

Image Processing

One doesn't always have to start with a blank image to create a new one. You can take existing images and transform them in interesting ways. The basic principles remain the same: you access an individual pixel and its color value and transform it in any way you like. In this section we will do some simple image transformations to illustrate this. All the image transformations we will learn will require the same transformation on all pixel values of a given image. You can of course choose to select specific regions if you like.

We will use the sample images shown here to learn image transformations. You can feel free to select one or more images from your own digital camera. In most of the examples we have used transformations on grayscale images. You should feel free to use color images and experiment with them. You will also notice that the bigger the size of the image, the longer it will take to transform it. This is a direct consequence of the amount of computation you have to do. For example, to transform a 500x500 image you will be operating on 250,000 pixels. If each pixel transformation involves 10 operations, you will be specifying 25 million operations! How long this will take depends on the computer you are using. We will return to this in more detail later in the chapter. For now, try and restrict yourself to small size images, say no bigger than 300x300 pixels.

As we mentioned earlier, even the most rudimentary digital cameras these days give you several megapixels of image resolution. So the first thing we should learn to do is how to shrink a given image into a smaller size.



Shrinking & Enlarging



Original Image (300x213)

After Shrinking
(150x106, F=2)

Lets us write a program that will take an input image and shrink it by a given factor, say F . For example, if the original image is 3000x3000 pixels and we shrink it by a factor of 10, we would end up with an image of size 300x300 pixels. This transformation boils down to selecting the value of each of the pixel in the new image based on the old image. We can use the following rule as our transformation:

*New pixel at x, y is a copy of the old pixel at $x * F, y * F$*

The program below prompts the user to pick an image, display it, and then ask the user to specify the shrinking factor (which has to be larger than 1).

```
int main() {  
  
    // read an image and display it  
    picture oldPic = makePicture(pickAFile());  
    show(myPic, "Before");  
  
    int X = getWidth(oldPic);  
    int Y = getHeight(oldPic);
```

```
// Input the shrink factor and computer size of new image
int F;
cout << "Enter the shrink factor.";
cin >> F;
int newx = X/F;
int newy = Y/F;

// create the new image
picture newPic = makePicture(newx, newy);

for (int x=0; x<newx; x++)
    for (int y=0; y<newy; y++)
        setPixel(newPic, x, y, getPixel(myPic, x*F, y*F));

show(newPic, "After");
return 0;
}
```

Do This: Select one or more images and make sure they are in JPEG format and are high resolution images (for our purposes, 400x400 pixels or larger). Run the program on each image and observe the output for different values of F . Modify the program to use the `savePicture` command to save the images into a file for use in examples below.

Enlarging is done in a complementary manner using the rule:

New pixel at x, y is a copy of the old pixel at $x/F, y/F$

The size of the new image will be F times the size of the original image. When you enlarge an image from a smaller image you are really generating data where there wasn't any. Therefore, the resulting image is likely to be more blurred than the original. This is similar to so called *upconverting* DVD players that claim to upconvert a movie to high definition from a standard signal. If you look closely, you will be able to see flaws in the unconverted images.

Blurring & Sharpening

You can blur an image by taking the average color values of neighboring pixels. The size of the neighborhood determines the amount of blurring. Below, we show how to blur each pixel using a neighborhood of 4 (see picture on right):

```
Pixel n1 = getPixel (oldPic, x,y);
Pixel n2 = getPixel (oldPic, x,y-1);
Pixel n3 = getPixel (oldPic, x-1,y);
Pixel n4 = getPixel (oldPic, x+1,y);
Pixel n5 = getPixel (oldPic, x,y+1);
```


```
int v =
    (getRed(n1)+...+getRed(n5))/5;
setPixel(newPic, x, y, gray(v));
```

Notice in the above that `n1` is the original pixel. You can use different sizes and shapes of neighborhoods to create more elaborate blurring effects. The next page shows the effect of using a larger, still square, neighborhood.

Instead of averaging the pixel values, if you subtracted the pixel values you end up sharpening the image. That is,

```
int v = 5*getRed(n1) -
    (getRed(n2)+...+getRed(n5));
setPixel(newPic, x, y, gray(v));
```

Do This: Write a complete program that uses the above neighborhood averaging formula to blur and sharpen images.



4-pixel neighborhood



Blurred Image



Sharpened Image

Depending on the size and shape of the neighborhood you choose, the loop index ranges will vary. For example, the loop for the 4-pixel neighborhood will appear as follows:

```
for (int x=1; x<X-1; x++)
  for (int y=1,y<Y-1; y++)
    // Average pixel values...
```

As the size of the neighborhood grows, you will have to adjust the start and stop values of the x and y variables in the loop above.


Negative & Embossing

Creating a negative of an image is obtained by inverting the pixel values. For grayscale values this amounts to subtracting from 255. For color pixels you have to subtract from 255 for each of the RGB values. The rule for grayscale is:

New pixel at x,y is 255-value of old pixel at x,y

You can specify the entire transformation using the loop:

```
for (x = 0; x < X; x++)
  for (y = 0; y < Y; y++) {
    Pixel pixel = getPixel(oldPic, x, y);
    int v = MAX - getRed(pixel);
    setPixel(newPic, x, y, gray(v));
  }
```



15-pixel neighborhood



Blurred Image

Similarly, you can create an embossing or relief effect by subtracting from the current pixels value, the value (or a fraction thereof) of a pixel two pixels away:


```
Pixel pixel1 = getPixel(oldPic, x, y);  
Pixel pixel2 = getPixel(oldPic, x+2, y+2);  
int v = getRed(pixel1) - getRed(pixel2);  
setPixel(newPic, x, y, gray(v));
```

or

```
int v = getRed(pixel1) + (MAX/2 - getRed(pixel2));
```



Negative



Emboss

Do This: Implement the two rules shown above. Try creating a negative of a color image, as well as a grayscale image. Try different fractions for embossing to see their effect.

Most likely, as you were reading and experimenting with the above techniques, you were wondering about the kind of image processing done by some of the popular photo-editing software programs. You may already be familiar with or might have heard of Adobe Photoshop, Fireworks, GIMP, etc. Most of these programs use techniques similar to the ones we have described above and provide many more features and options for image manipulation.

Image Understanding & Robot Vision

The picture shown on right was taken from a newspaper article in a local newspaper in western New York. Read the accompanying caption and see if you can identify each individual in the picture. After reading the caption, you will no doubt confirm that the person in the front row center is Coach Kevin Starr. It would be really hard to do identify the people if the caption weren't present. Still you would be able to answer questions like, how many people are in the picture. This is an example of the *image understanding* problem: Given an image, how can you determine, computationally, that there are seven people in the image, that three of them are kneeling and four are standing, etc. You would also not have any trouble drawing a box on the face of each person in the image. In this section we introduce you some basic techniques in image understanding and use them for creating a visual perception system for the Scribbler robot. You are already familiar with the basic image processing commands. The same commands will be used to do image understanding.



"THE BEST - The All-WNY boys volleyball team surrounds the coach of the year, Frontier's Kevin Starr. Top row, from left, are Eden's Brian Heffernan, Lake Shore's Tom Kait, Sweet Home's Greg Gegenfurtner, Kenmore West's George Beiter. Bottom row: from left, are Eden's David Nagowski and Orchard Park's Rich Bland."

For a robot to perceive its environment visually it has to start by taking a picture. Once the image is obtained perception of the contents of the image are done computationally. Instead of taking on the image understanding task shown above, we will start simple: how does a Scribbler recognize a ball? Once it recognizes it, can it follow the ball wherever it goes?

The second problem becomes easy once we know how to recognize a ball. We can write a simple control program like this:


```

while (timeRemaining(T)) {
    // take picture and locate the ball
    ...
    if (ball is on the left of visual field)
        robot.turnLeft(cruiseSpeed);
    else if (ball is in the middle of visual field)
        robot.forward(cruiseSpeed);
    else
        robot.turnRight(cruiseSpeed);
}

```

The main problem is locating the ball in the image. Given the pixel manipulation commands you have learned so far, think about how we might accomplish this?

The basic principles in image understanding rely on starting at the pixel level. Let us make things a little more concrete. Take a look at the image shown here on the right. The picture was taken by a Scribbler robot. For our purposes we have exaggerated the scene by choosing a bright colored (pink) ball. When you click your mouse anywhere on the image it will give you the RGB values of that pixel. The picture shows the values of a pixel in the center of the ball. Notice that the red pixel is 253 and the green pixel's value is 66. By the way, true pink has RGB values (255,175,175). We can use this information as a way of identifying the ball. Here is a simple image transformation loop that turns all pixels in the image into white or black pixels. If a pixel has a high red value and a lower green value it will be turned white else it will be turned black.




```

Pixels pixels = getPixels(p);
for (Pixels pixel = pixels.begin();
     pixel != pixels.end(); pixel++) {
    color_rgb rgb = getRGB(*pixel);
    if (rgb.red > 200 && rgb.green < 100)
        setRGB(pixel, makeColor(255, 255, 255));
}

```

```
else
    setRGB(pixel, makeColor(0, 0, 0));
```

The resulting image is shown here. As you can see, by just selecting the appropriate values of individual colors you can filter out the unwanted regions and focus on just the colors you are interested in. Even though the entire ball has not been filtered, it is sufficient for use in identifying its location. By tweaking the RGB threshold values you can further refine this filter. Further processing will be required to identify the ball's location in the visual field (which is the image).



We are interested in locating the ball in the left, right, or center of the image. One way to accomplish this is to locate the x -coordinate of the center point of the ball. One way to do this is to take the average of the x -locations of all white pixels in the image:

```
int locateBall(picture p) {
    int tot_x = 0;
    int count = 0;
    Pixels pixels = getPixels(p);
    for (Pixels pixel = pixels.begin();
         pixel != pixels.end(); pixel++) {
        color_rgb rgb = getRGB(*pixel);
        if (rb.red > 200 && rgb.green < 100) {
            tot_x = tot_x + getX(pixel);
            count = count + 1;
        }
    }
    return (double) tot_x/count;
}
```

Alternately, you could also average the x-locations of all pink pixels without using the filter. On the image shown, the average x-value returned was 123. Given that the image is 256 pixels wide, 123 is just about in the center. Below, we show the locations obtained for each of the images shown:



location = 41



location = 123



location = 215


Given that the image's width is 256 pixels, you can divide the visual field into three regions: left (0..85), center (86..170), and right (171, 255). The robot control to follow the ball can then be defined as shown below:

```
while (timeRemaining(T)) {
  // take picture and locate the ball
  picture pic = robot.takePicture();
  int ballLocation = locateBall(pic);


  if (ballLocation <= 85)
    robot.turnLeft(cruiseSpeed);
  else if (ballLocation <= 170)
    robot.forward(cruiseSpeed);
  else
    robot.turnRight(cruiseSpeed);
}
```

Do This: Implement the above program to recognize a given object in the picture. The object could be a box, a ball, a trash can, etc. Experiment with different colored objects. Can you use these techniques to recognize any object? Say an animal? A person? Etc.

Blobs



Setting a blob on the ball.



Result of `takePicture("blob")`

The kind of threshold filtering used above is commonly called *blob filtering*. All pixels satisfying the selected color properties are to be identified and the rest eliminated. This is such a common operation in image processing that many cameras and camera hardware integrate blob filtering as a primitive. The Scribbler also has this capability. Using this, you do not have to worry about identifying specific color values for the filter threshold. All you have to do is take a picture and then display it. Then using your mouse, define the rectangular region that contains your desired blob. To do this, take a picture, display it, and then click your mouse on the top left corner of the ball in the image and then drag it to the bottom right corner before releasing it. You have to make sure that you are connected to the robot because the selected blob region is transferred back to the robot. Once that is defined, you will be able to make use of the predefined blob filter by taking pictures using the command:

```
picture pic = takePicture("blob");
```

As soon as the picture is taken, the camera's blob detector goes to work and returns the resulting image (shown above). Now you can use this image as

before to control your robot. You will have to average the locations of all black pixels.


Alternately, you could use the `getBlob()` function, which returns a three-element vector of integers:

```
vector<int> blobData = getBlob();  
int onPixels = blobData[0];  
int avgX = blobData[1];  
int avgY = blobData[2];
```

The function returns three values: the number of ‘on’ pixels (i.e. the pixels in the image that were in the detected blob), and the average x- and y- location of the blob pixels. This is a better and much faster way of doing blob recognition and you will notice that your robot tracks the ball much better.

Do This: Redo the ball tracking example from above to use the blob features introduced here.

Further computations on an image can result in identifying object boundaries. This is called *edge detection*. These can be further used to do object recognition: Is there a car in the picture? A face? Etc. Object recognition requires the ability to identify features that are typical of an object. The domain of image recognition and computational vision is rich with techniques and applications like these. On the right, we show the results of recognizing faces and identifying the



individuals in the picture. As researchers advance these techniques and as computational capabilities of digital cameras increase, we are starting to see many of these incorporated into even the most basic digital cameras. Such techniques can be used to automatically focus on the main subject and also adjust exposure to get a perfect picture even for amateur photographers.

Summary

In this chapter, you have seen various Myro facilities that can be used for image creation, processing and understanding. Most image creation, processing, and understanding techniques involve pixel-level manipulations of an image. However, many interesting effects as well as applications can be developed using these simple ideas. Many of these features are increasingly found in everyday devices like digital cameras, biometric identification devices, and even automated mail sorting systems. Combined with robot behaviors these facilities can be used for modeling smarter behaviors.

Myro review²

`getBlob()`

Return a three-element vector of integers of the number of pixels that matched the blob, the average x and y locations of blob pixels.

`getHeight(<picture>)`
`getWidth(<picture>)`

Returns the height and width of the `<picture>` object (in pixels).

`getPixel(<picture>, x, y)`

Returns the pixel object at `x, y` in the `<picture>`.

² The C++/Myro interface is under development and might not agree with the description in this chapter.

`getPixels (<picture>)`

Returns an iterator of type `Pixels`, which returns one `Pixel` at a time from `<picture>`.

`getRGB (pixel)`
`getRed (<pixel>)`
`getGreen (<pixel>)`
`getBlue (<pixel>)`

Returns the RGB values of the `<pixel>`. `getRGB` returns a `color_rgb` object.

`getBlob ()`

Returns a three-element `vector<int>` containing the number of on-pixels and the average x- and y-locations of blob pixels. See chapter text for example.

`makeColor (<red>, <green>, <blue>)`

Creates a `color_rgb` object with the given `<red>`, `<green>`, and `<blue>` values (all of which are in the range `[0..255]`).

`makePicture (<file>)`
`makePicture (<width>, <height>)`
`makePicture (<width>, <height>, <color>)`

Creates a `picture` object either by reading a picture from a `<file>`, or of the given `<width>` and `<height>`. If `<color>` is not specified, the picture created has a white background.

`pickAColor ()`

Creates an interactive dialog window to select a color visually. Returns the color object corresponding to the selected color.

`pickAFile ()`

Creates an interactive dialog window that allows user to navigate to a folder and select a file to open. Note: it cannot be used to create new files.

`repaint ()`;
`repaint (<picture>)`;

Refreshes the displayed `<picture>`.

```
savePicture(<picture>, <file>);  
savePicture(<picture vector>, <gif file>);
```

Saves the <picture> in the specified file (a GIF or JPEG as determined by the extension of the <file>: .gif or .jpg). <picture vector> is saved as an animated GIF file.

```
setColor(<pixel>, <color_rgb>);  
setRed(<pixel>, <value>);  
setGreen(<pixel>, <value>);  
setBlue(<Pixel>, <value>);
```

Sets the color of <pixel> to specified <color_rgb> or <value>.

```
setPixel(<picture>, x, y, <color_rgb>);  
setPixel(<picture>, x, y, <pixel>);
```

Sets the pixel at x, y in <picture> to specified <color_rgb> or color of <pixel>.

```
show(<picture>);  
show(<picture>, <name>);
```

Displays the <picture> on the screen in a window named <name> (string).

```
takePicture()  
takePicture("gray")  
takePicture("blob")
```


Takes a picture from the Scribbler camera. It is a color picture by default, or grayscale ("gray"), or a filtered image based on the defined blob ("blob"). See chapter text for examples.

C++ Review


There were no new C++ features introduced in this chapter.

Exercises

1. Instead of viewing each pixel in an image as x - and y - values in the Cartesian coordinate system, imagine that it is in a polar coordinate system with the origin in the center of the picture (See picture below).



Cartesian Coordinates



Polar Coordinates

Write a function called `polar(x, y)` that returns the polar coordinates a two-element `double`-vector (r, θ) of the given point. Using this function write a transformation as follows:

```
for (int x = 1; x < W; x++)
  for (int y = 1; y < H; y++) {
    vector<double> r_theta = polar(x, y);
    double r = r_theta[0], theta = r_theta[1];
    setPixel(myPic, x, y,
             gray(((r+theta)%16)-8)*MAX/16+MAX/2));
  }
```

Create a 400x400 picture using the above transformation. Start with a black picture.

2. Based on the previous exercise, try transforming an image (like a portrait). For example use the rule: new x, y is a copy of original picture's $\sqrt{[r \times W/2]}$, *heta*.
3. Write a transformation that flips an image clockwise (or counter clockwise) by 90 degrees. Try it on a square image of your choice.
4. Write an image transformation that averages two images (same size) into a single image.
5. Write a program that sews two or more images together to create a panorama.
6. Use blob detection techniques described in this chapter to recognize Coke and Pepsi cans. Write a program to hunt for a Coke or Pepsi can.
7. Can you expand on Exercise above to recognize small toy figurines of Disney cartoon characters like Mickey, Minnie, Daisy, Goofy, Pluto, and Donald?



10

Artificial Intelligence

David: Martin is Mommy and Henry's real son. After I find the Blue Fairy then I can go home. Mommy will love a real boy. The Blue Fairy will make me into one.

Gigolo Joe: Is Blue Fairy Mecha, Orga, man or woman?

David: Woman.

Gigolo Joe: Woman? I know women! They sometimes ask for me by name. I know all about women. About as much as there is to know. No two are ever alike, And after they've met me, no two are ever the same. And I know where most of them can be found.

David: Where?

Gigolo Joe: Rouge City. Across the Delaware.

Dialog between two Artificial Intelligence entities: Gigolo Joe (played by Jude Law) and David (played by Haley Joel Osment) in the movie, *Artificial Intelligence*, Directed by Steven Spielberg, Warner Bros., 2001.

Opposite page: A.I. Artificial Intelligence
From the movie poster. Warner Bros., 2001.

The Question of Intelligence

The quest for the understanding of intelligence probably forms the oldest and yet to be fully understood human inquiry. With the advent of computers and robots the question of whether robots and computers can be as intelligent as humans has driven the scientific pursuits in the field of *Artificial Intelligence (AI)*. Whether a computer can be intelligent was lucidly discussed by Professor Alan Turing in 1950. To illustrate the issues underlying machine intelligence, Turing devised a thought experiment in the form of an *imitation game*. It is played with three people, a man, a woman, and an interrogator. They are all in separate rooms and interact with each other by typing text into a computer (much like the way people interact with each other over IM or other instant messaging services). The interrogator's task is to identify which person is a man (or woman). To make the game interesting, either player can try and be deceptive in giving their answers. Turing argues that a computer should be considered intelligent if it could be made to play the role of either player in the game without giving itself away. This *test* of intelligence has come to be called the *Turing Test* and has generated much activity in the community of AI researchers (see Exercises). The dialog shown above, from the movie *Artificial Intelligence*, depicts an aspect of the test of intelligence designed by Alan Turing. Based on the exchange between Gigolo Joe and David, can you conclude that they are both intelligent? Human?

After over five decades of AI research the field has matured and evolved in many ways. For one, the focus on intelligence is no longer limited to humans: insects and other forms of animals with varying degrees and kinds of intelligence have been the subject of study within AI. There has also been a fruitful exchange of ideas and models between AI scientists, biologists, psychologists, cognitive scientists, neuroscientists, linguists and philosophers. You saw examples of such an influence in the models of Braitenberg vehicles introduced earlier. Given the diversity of researchers involved in AI there has also been an evolution of what AI itself is really about. We will return to this later in the chapter. First, we will give you a few examples of models that could be considered *intelligent* that are commonly used by many AI scientists.

Language Understanding

One aspect of intelligence acknowledged by many people is the use of language. People communicate with each other using a language. There are many (several thousand) languages in use on this planet. Such languages are called *natural languages*. Many interesting theories have been put forward about the origins of language itself. An interesting question to consider is: Can people communicate with computers using human (natural) languages? In other words, can a computer be made to understand language? Think about that for a few moments.

To make the question of language understanding more concrete, think of your Scribbler robot. So far, you have controlled the behavior of the robot by writing C++ programs for it. Is it possible to make the Scribbler understand English so that you could interact with it? What would an interaction with Scribbler look like? Obviously, you would not expect to have a conversation with the Scribbler about the dinner you ate last night. However, it would probably make sense to ask it to move in a certain way. Or to ask whether it is seeing an obstacle ahead.

Do this: Write down a series of short 1-word commands like: *forward*, *right*, *left*, *stop*, etc. Create a vocabulary of commands and then write a program that inputs a command at a time interprets it and makes the Scribbler carry it out. For example:

```
You: forward
Scribbler: starts moving forward...
You: right
Scribbler starts turning right...
You: stop
...
```

Experiment with the behavior of the robot based on these commands and think about the proper interpretation that may make its behavior more natural.

You will find yourself making several assumptions about interpretation of even the simplest commands in the exercise above. For example, what happens when after you command the Scribbler to move forward, you ask it to turn right? Should the Scribbler stop going forward or should it stop and then start turning?

Decisions like these also give deep insights into our own abilities of understanding language. You can also see that, as in the case of visual perception, processing of language (or text) begins at a very primitive level: words. If the input is speech, the basic units are electrical signals, perhaps coming from a microphone. Just like processing individual pixels to try and understand the contents of an image, one has to start at a low level of representation for beginning to understand language.

Researchers working in the field of *computational linguistics* (or *natural language understanding*) have proposed many theories of language processing that can form the basis of a computational model for a Scribbler to understand a small subset of the English language. In this section, we will examine one such model which is based on the processing of syntax and semantics of language interaction. Imagine, interacting with the Scribbler using the following set of sentences:

```
You: do you see a wall?  
Scribbler: No
```

```
You: Beep whenever you see a wall.  
You: Turn right whenever you see a wall to your left.  
You: Turn left whenever you see a wall to your right.  
You: Move for 60 seconds.
```

```
[The Scribbler robot moves around for 60 seconds turning  
whenever it sees a wall. It also beeps whenever it sees a  
wall.]
```

Earlier, you have written C++ programs that perform similar behaviors. However, now imagine interacting with the robot in the fashion described. From a physical perspective, imagine that you are sitting in front of a

computer, and you have a Bluetooth connection to the robot. The first question then becomes: Are you actually speaking or typing the above commands? From an AI perspective, both modalities are possible: You could be sitting in front of the computer and speaking into a microphone; or you could be typing those commands on the keyboard. In the first instance, you would need a speech understanding capability. Today, you can obtain software (commercial as well as freeware) that will enable you to do this. Some of these systems are capable of distinguishing accents, intonations, male or female voices etc. Indeed, speech and spoken language understanding is a fascinating field of study that combines knowledge from linguistics, signal processing, phonology, etc.

You can imagine that the end result of speaking into a computer is a piece of text that transcribes what you said. So, the question posed to the Scribbler above: *Do you see a wall?* will have to be processed and then transcribed into text. Once you have the text, that is, a string "Do you see a wall?" it can be further processed or analyzed to understand the *meaning* or the content of the text. The field of *computational linguistics* provides many ways of syntactic parsing, analyzing, and extracting meaning from texts. Researchers in AI itself have developed ways of representing knowledge in a computer using symbolic notations (e.g. *formal logic*). In the end, the analysis of the text will result in a `getIR()` or `getObstacle()` command to the Scribbler robot and will produce the response shown above.

Our goal of bringing up the above scenario here is to illustrate to you various dimensions of AI research that can involve people from many different disciplines. These days, it is entirely possible even for you to design and build computer programs or systems that are capable of interacting with robots using language.

Game Playing

In the early history of AI, scientists posed several challenging tasks which if performed by computers could be used as a way of demonstrating the feasibility of machine intelligence. It was common practice to think of games

in this realm. For example, if a computer could play a game, like chess, or checkers, at the same level or better than humans we would be convinced into thinking that it was indeed feasible to think of a computer as a possible candidate for machine intelligence. Some of the earliest demonstrations of AI research included attempts at computer models for playing various games. Checkers and chess seemed to be the most popular choices, but researchers have indulged themselves into examining computer models of many popular games: Poker, Bridge, Scrabble, Backgammon, etc.



Two-person zero-sum games: Chess, Tic Tac Toe, Konane

In many games, it is now possible for computer models to play at the highest levels of human performance. In Chess, for example, even though the earliest programs handily beat novices in the 1960's, it wasn't until 1996 when an IBM computer Chess program, named Deep Blue, beat the world champion Gary Kasparov at a tournament-level game, though Kasparov did manage to win the match 4-2. A year later, in New York, Deep Blue beat Kasparov in a 6 game match representing the very first time a computer defeated the best human player in a classical style game of Chess. While these accomplishments are worthy of praise it also now clear that the quest for machine intelligence is not necessarily answered by computer game playing. This has resulted in much progress in game playing systems and game playing technology which is now a multi-billion dollar industry.

It turns out that in many Chess-like games the general strategy for a computer to play the game is very similar. Such games are classified as two-person zero-sum games: two people/computers play against each other and the result of the game is either a win for one player and loss for the other, or it is a draw.

In many such games, the basic strategy for making the next move is simple: look at all the possible moves I have and for each of them all the possible moves the other player might have and so on until the very end. Then, trace back from wins (or draws) and make the next move based on those desirable outcomes. You can see this even in simple games like Tic Tac Toe where it is easy to mentally look ahead future moves and then make more informed decisions about what to do next. The best way to understand this is to actually write a program that plays the game.

Tic Tac Toe

Also known as *Noughts and Crosses* or *Hugs and Kisses*, Tic Tac Toe is a popular children's game (see description on Wikipedia under Tic-tac-toe). We will develop a program that can be used to play this game against a person. Almost any board game can be programmed using the basic loop shown below:

```
void play() {
    // Initialize board
    Board board = makeBoard();

    // set who moves first/next: X always moves first
    char player = 'X';

    // Display the initial board
    display(board);

    // The game loop
    while (winner(board) == " " &&
           !gameOver(board, player)) {
        move(board, player);
        display(board);
        player = opponent(player);
    }

    // game over, show outcome
    string winningPiece = winner(board);
```

```
if (winningPiece != "Tie")
    cout << winningPiece << " won.\n";
else
    cout << "It is a tie.\n";
}
```

The function above can be used to play a round of any two-person board game. The variable `player` is the player (or piece) whose move is next. We are already using the Tic Tac Toe piece 'X' in the function above. Six basic functions (shown highlighted above) make up the basic building blocks of the game. For Tic Tac Toe, they can be defined as:

1. **makeBoard()** : Returns a fresh new board representing the start of the game. For Tic Tac Toe, this function will return an empty board representing the nine squares.
2. **displayBoard(board)** : Displays the board on the screen for the user to see. The display can be as simple or elaborate as you wish. It is good to start with the easiest one you can write. Later you can make it fancier.
3. **opponent(player)** : Returns the opponent of the current player/piece. In Tic Tac Toe, if the player is X, it will return an O, and vice versa.
4. **move(board, player)** : Updates the board by making one move for the player. If the player is the user, it will input the move from the user. If the player is the computer, it will decide how to make the best move. This is where the smarts will come in.
5. **gameOver(board)** : Returns `true` if there are no more moves left to be made, `false` otherwise.
6. **winner(board)** : Examines the board and returns the winning piece or that the game is not yet over, or that it is a tie. In Tic Tac Toe, it will return either an X, O, a blank (representing game is not over yet), or a TIE.

We will need a few more functions to complete the game but these six form the basic core. We will write them first.

The board itself consists of nine squares that start out being empty. Players choose a game piece: an 'X' or an 'O'. We will assume that 'X' always goes first. The first thing to do then is to design a representation of the game board. We will use the following simple representation:

```
typedef vector<char> Board;
Board board (9, ' ');
```

The first line allows `Board` to be used as an abbreviation for the type `vector<char>`. The second line defines `board` to be a vector of 9 single characters initialized to blanks. Note that we are using this linear representation of the board instead of a 2-dimensional one. However, as you will see, this representation makes it easier to do many manipulations for the game. During play, the board can be displayed in its natural format. Below, we show two functions: one creates a fresh new board each time it is called; and one displays it:

```
Board makeBoard() {
    /* A 3x3 board is represented as a list of 9 elements.
     * We will use the following numbering to locate a square
     *  0 | 1 | 2
     * ---|---|---
     *  3 | 4 | 5
     * ---|---|---
     *  6 | 7 | 8
     */
    return Board (9, ' ');
}

void display (Board board) {
    for (int i = 0; i < 9; i += 3) {
        if (i > 0)
            printf ("---|---|---\n");
        printf (" %c | %c | %c \n",
            board[i], board[i+1], board[i+2]);
    }
    cout << endl;
}
```

One advantage of writing the display function as shown is that it gives us a quick way of creating and displaying the game. Later, when you are done, you can write a fancier version that displays the game graphically (see Exercises). With the above functions, we can easily create a fresh new board and display it as follows:

```
Board board = makeBoard();
display(board);
```

To determine the opponent of a given piece is simple enough:

```
char opponent(char player) {
    if (player == 'X')
        return 'O';
    else
        return 'X';
}
```

Next, we have to write the function `move`. It first determines whose move it is. If it is the user's move, it will input the move from the user. Otherwise, it will determine the best move for the computer. Then it will actually make the move. As a first design, we will let `move` make a random choice out of the possible moves for the computer. Later, we will make a more informed decision.

```
#include <cstdlib>
using namespace std;

char You = 'X';
char Me = 'O';

void move(Board& board, char player) {
    int square;

    if (player == You) { // user's move?
        cout << "Enter your move: ";
        cin >> square;
        square--;
    } else // my turn
```



```

    // player is the computer, make a random choice
    square = choice(possibleMoves(board, player));
    // place player's piece on the chosen square
    applyMove(board, player, square);
}

```

Notice the ampersand & after type `Board` in the first line of `move`. This means that when we call `move`, the computer will not make a fresh copy of the board for `move` to use (which is what usually happens), but will pass a *reference* to (the address of) the board in the caller. This is because the purpose of `move` is to make a move, which changes the board, and so we want to change the original game board and not a local copy of it. This way of passing a parameter to a function is called *pass by reference* (as opposed to the usual *pass by value*).

We will leave it as an exercise for you to define an integer function `choice(vec)` that returns a random element of the integer vector `vec`.

We have set the global variables `You` and `Me` to specific pieces. This is a simplification for now. Later you can come back and rewrite them so that for each game, the user gets to select their piece. In Tic Tac Toe, X always moves first. So by making the user's piece X, we have made an assumption that the user always goes first. Again, later we can come back and modify this (see exercises). Also notice that we are not doing any error checking in user input to ensure that a legal move was input (see exercises).

The user inputs their move by entering a number from 1..9 where the square numbering is as shown below. This is slightly different from our internal numbering of squares and more natural for people. In the `move` function above, we subtract 1 from the input number so it maps to the proper square in our internal scheme.

```

 1 | 2 | 3
---|---|---
 2 | 5 | 6
---|---|---
 7 | 8 | 9

```

Again, we have simplified the interface for now. The exercises suggest how to improve on this design.

The `move` function defined above requires two additional functions (shown highlighted). These are also core functions in any board-based game and are described below:

7. **`possibleMoves(board, player)`** : Returns a list of possible moves for the given player.
8. **`applyMove(board, player, square)`** : Given a specific square and a player/piece, this function actually applies the move on the board. In Tic Tac Toe all one has to do is actually place the piece in the given square. In other games, like Chess or Checkers, there may be pieces that get removed.

In Tic Tac Toe all empty squares are possible places where a player can move. Below, we write a function that, given a board, returns a list of all the possible locations where a piece can be placed:

```
vector<int> possibleMoves(Board board, char player) {
    vector<int> moves (0);
    for (int i = 0; i < 9; i++)
        if (board[i] == ' ')
            moves.push_back(i);
    return moves;
}
```

To complete the game playing program, we need to write two more functions defined above. The `winner` function examines the board and determines who won. It returns the winning piece (an "X", "O") or the string "Tie". In case the game is not yet over, it returns " ". Below, we first define all the winning positions in Tic Tac Toe, based on our board representation. Next, we define the function itself to examine the board.

```

// These square triples represent wins (three in a row).

int Wins[][3] = {{0, 1, 2},{3, 4, 5},{6, 7, 8}, // the rows
                 {0, 3, 6},{1, 4, 7},{2, 5, 8}, // the columns
                 {0, 4, 8},{2, 4, 6}};        // diagonals

string winner(Board board) {
    for (int win = 0; win < 8; win++) {
        char posWinner = board[Wins[win][0]];
        if (posWinner != ' ' &&
            posWinner == board[Wins[win][1]] &&
            posWinner == board[Wins[win][2]])
            return string(1, posWinner);
    }

    // No winner yet, are there empty squares left?
    for (int i = 0; i < 9; i++)
        if (board[i] == ' ')
            return " ";
    // The board is full and no one has three in a row
    return "Tie";
}

```

The declaration of `Wins` illustrates a C++ feature that you have seen before (ch. 7), but uses it in a new way. It declares `Wins` to be a two-dimensional array; `[][3]` means that it has three columns but the number of rows is determined by the following initialization, which lists the contents of the array row by row. In the body of `winner`, expressions such as `Wins[win][1]` refer to row `win` and column 1 of `Wins`.

Last, the `gameOver` function can be written either by relying on the fact that `winner` returns " " when the game is not yet over. Alternately, we can write it using `possibleMoves` as follows:

```

bool gameOver(Board board, char player) {
    return possibleMoves(board).size() == 0;
}

```

Implement `applyMOVE` and you have all the ingredients for a program to play a game of Tic Tac Toe. The version above is simplified in many ways, and yet it captures the essential elements of playing a round of Tic Tac Toe.

Do This: Implement the program above and play a few rounds of Tic Tac Toe. When you play against the computer, you are anticipating possible moves down the road and then playing your own moves with those in mind.

You probably had no problems beating the computer at the game. That is because the computer is just picking its moves by random from a set of possible moves:

```
// player is the computer, make a random choice
square = choice(possibleMoves(board, player));
```

This is the line in `move` where we can put some intelligence into the program. However, the basic question that needs to be asked is: which of the possible moves is the best one for me? Think about this for yourself a little. In Tic Tac Toe, you play in two modes: defensively so as not to lose after the next move, or offensively to try and win. If a win is imminent for you, you will of course make that move, but if it isn't (and neither is a loss) how do you select your move? We will try and generalize this idea next in a way that will also be applicable to other board games.

Let us delegate the responsibility of finding the best move to a function: `bestMove` that will return the best move from among a set of possible moves. That is, we can change the line in `move` as follows:

```
// player is the computer, make the best choice
square = bestMove(board, player,
                  possibleMoves(board, player));
```

If `bestMove` were to make a random choice (as above) we could write it as:

```
int bestMove(Board board, char player, vector<int> moves) {
    return choice(moves);
}
```


Imagine that you decided to go watch two players playing a game of Chess (or Tic Tac Toe). However, you get there in the middle of the game. You walk into the room, look at the board, and are able to evaluate how each player is doing without knowing anything about how the board came to be. For example, take a look at the Tic Tac Toe boards:

<pre> X O X --- --- --- O --- --- --- O X X </pre>	<pre> O X X --- --- --- --- --- --- X O </pre>	<pre> X O X --- --- --- O X O --- --- --- X </pre>	<pre> X --- --- --- --- --- --- </pre>
1	2	3	4

In all the cases above, the computer is playing O and it is O's turn next. Examine each board and think about what the best move for O would be.

In case 1, the computer has to play defensively and place an O in square 6 to avoid losing the game. In case 2, it can be on the offensive and recognize that it will win the game by placing an O in square 5. In case 3, it has to play to avoid losing. In the last case, the board is wide open. Of the eight possible choices, is there a best move? From experience, you are probably going to place the O in the center square (square 5). Why? Let us elaborate this by looking ahead (see Figure on next page).

Think about how we could quantify each of the above board positions with respect to O's chances of winning. For each board, count the number of possible winning positions still remaining for O and compare those for X. For example, in the first board, O has 5 possible winning positions, and X has 6. We could say that O is $5 - 6 = -1$ has one less possibility than X. If you did this for the second and the third boards, you will see that both X and O are even in the number of possible remaining winning positions (5 each). However, in the fourth case, O has 5 possible winning positions remaining, and X has only 4. Try and work out the remaining board positions to confirm the scores shown above. Clearly, you can see that placing O in square 5 is the best of all the possibilities. Try and elaborate the 5 possible moves for O in



case 2 above. You will find that one of the five leads to a win for O. Similarly, when you elaborate the two options for O in case 1, you will notice that one leads to a tie and the other to a loss for O. We can capture all these situations in a function that takes a look at the board and returns a number representing a score in favor of O: the higher the value returned, the more desirable the board.

$$\text{evaluate}(\text{board}) = \begin{cases} \infty, & \text{if board is a win for computer} \\ -\infty, & \text{if board is a loss for computer} \\ \text{openWins}(\text{computer}) - \text{openWins}(\text{user}), & \text{otherwise} \end{cases}$$

That is, a win for the computer receives a high score, a loss receives a low score. Otherwise, it is the difference in the number of open wins remaining for each. We can capture this in C++ as follows:

```
int INFINITY = 10;
int evaluate(Board board, char player) {
    // if board is a win for player, return INFINITY
    char piece = winner(board, player);
    if (piece == Me)
        return INFINITY;
    else if (piece == You)
        return -INFINITY;
    else
        return openWins(board, Me) - openWins(board, You);
}
```

We define `INFINITY` as 10, a large enough number relative to the other values that might be returned by `evaluate`. `openWins` looks at each winning triple on the board and counts the number of openings for the given player.

```
int openWins(Board board, char player) {
    int possWins = 0;
    for (int pos = 0; pos < 8; pos++) {
        int n = 0;
        for (int i = 0; i < 3; i++)
            if ((board[Wins[pos][i]] == player) ||
                (board[Wins[pos][i]] == ' '))
                n++;
        if (n == 3) possWins++;
    }
    return possWins;
}
```

Do This: Implement the two functions above and then test them. Create several board positions (use the ones from examples above) and confirm that the board is being evaluated to the correct values.

Next, we can rewrite `bestMove` to take advantage of `evaluate`:

```
int bestMove(Board board, char player, vector<int> moves) {
    vector<int> scores (0);
    for (m = 0; m < moves.size(); m++) {
        Board b (board);
        applyMove(b, player, moves[m]);
        scores.push_back(evaluate(b));
    }
    return moves[maxIndex(scores)];
}
```

The declaration `Board b (board)` initializes `b` to be a fresh copy of `board`, so that we won't affect the original board while we are trying out possible moves. Notice how `bestMove` takes each possible move, creates a new board with that move, and then evaluates the score of the resulting board. It pushes each evaluation on the back (end) of the `scores` vector, which is initially

empty. Finally, it returns the move with the highest score as the best move. Modify your program from above to use this version of `bestMove` (you will need to program `maxIndex(v)`, which returns the index (position) of the maximum element of `v`). Play the game several times. You will notice that there is a significant improvement in the computer's playing ability. Can you measure it in some way? (See Exercises).

The above rewrite of `bestMove` will make the program play significantly better but there is still more room for improvement. In most board games good players are able to mentally picture the game several moves ahead. In many games, like Chess, certain recognizable situations lead to well determined outcomes and so a great part of playing a successful game also relies on the ability to recognize those situations.

Looking ahead several moves in a systematic manner is something computers are quite capable of doing and hence anyone (even you!) can turn them into fairly good players. The challenge lies in the number of moves you can look ahead with limited memory capacity and, if time to make the next move is limited, how to choose among the best available options. These decisions lend interesting character to computer game programs and continue to be a constant source of fascination for many people. Let's us look at how your Tic Tac Toe program can easily look at all the possible moves all the way to the end of game in determining its next move (which, in most situations leads to a draw, given the simplicity of the game).

When you look ahead a few moves, you take into account that your opponent is going to try and beat you at every move. Your program, in trying to select the best move, can look ahead at the opponent's moves and take that into consideration when choosing its best move. In fact, it can go further, all the way to the end. The `evaluate` function we wrote above can be used effectively to evaluate future board situations by assuming that when it is the computer's move, it will always try to pick the move that promises the highest score in the future. However, when it examines the opponent's moves, it has to assume that the opponent is going to make the move that is worst for the computer. In other words, when looking ahead, the computer is going to

maximize its possible score while the opponent is going to minimize the computer's chances to win. This can be captured in the following:

```
int lookahead(Board board, char player, bool MAX, int level) {

    vector<int> moves = possibleMoves(board, player);
    if (level <= 0 || moves.size()==0) // limit of look ahead
        return evaluate(board, player);

    int V;
    if (MAX) { // computer's move
        V = -INFINITY;
        for (int m = 0; m < moves.size(); m++) {
            Board b (board);
            b[moves[m]] = player;
            V = max(V,
                lookahead(b, opponent(player), !MAX, level-1));
        }
    } else { // opponent's move
        V = INFINITY;
        for (int m = 0; m < moves.size(); m++) {
            Board b (board);
            b[moves[m]] = player;
            V = min(V,
                lookahead(b, opponent(player), !MAX, level-1));
        }
    }
    return V;
}
```

The `lookahead` function defined above takes the current board, the player whose turn it is, whether the player is the computer (one trying to maximize its outcome) or the opponent (one trying to minimize the computer's outcomes), and the levels still to look ahead, and computes a score based on examining all the moves going forward to the limit of look ahead. In the above function when `MAX` is `true` it represents the computer and `false` represents its opponent. Thus, depending on the value of `MAX`, the evaluation is minimized or maximized accordingly. Each time it looks ahead further, the level is reduced by 1. The final value returned by `lookahead` can be used by `bestMove` as follows:

```
int LEVEL = 9;
int bestMove(Board board, char player, vector<int> moves) {

    vector<int> scores (0);
    for (int m = 0; m < moves.size(); m++) {
        Board b (board);
        b[moves[m]] = player;
        scores.push_back(
            lookahead(b, opponent(player), false, LEVEL-1));
    }

    return moves[maxIndex(scores)];
}
```

As before, the move with the highest value is considered the best move. We have set the value of `LEVEL` above at 9 (i.e. look 9 moves ahead!) implying that each time it will look as far as the end of the game before making the decision. There can only be a maximum of 9 moves in a Tic Tac Toe game. The quality of the computer's decision maker can in fact be adjusted by lowering the value of `LEVEL`. At `LEVEL = 1`, it will be equivalent to the version we wrote earlier that only used the `evaluate` function.

How many levels ahead one looks in a game like this can depend on the game itself. Can you guess how many board situations the computer will have to look at in doing a look ahead at `LEVEL = 9` after the user's first move? It will be 40,320 different board situations! Why? Additionally, by the time it is the computer's second move, it will only need to look at 720 board positions. This is because, in Tic Tac Toe, as the board gets filled, there are fewer possible moves remaining. In fact, by the time it is the computer's third move, it only needs to look at a total of 24 boards. And, if the computer makes it to its fourth move, it will only have to look at two possible moves. Thus, an exhaustive search for all the possible board positions until the end of the game each time the computer has to make a move it will be examining a total of 41,066 board positions. However, if you consider a typical game of Chess, in which each player makes an average of 32 moves and the number of feasible moves available at any time averages around 10, you would soon realize that the computer would have to examine something of the order of 10^{65} board

positions before making a move! This, even for the fastest computers available today, will take several gazillion years! More on that later. But, to play an interesting two-person zero-sum game, it is not essential to look so far ahead. You can vary the amount of look ahead by adjusting the value of `LEVEL` in such programs.

Do This: Implement `lookahead` as described above and compare how well the computer plays against you. Try and vary the levels from 1, 2, ..., to see if there is any improvement in the computer's play. Would you consider this program intelligent?

The exercises at the end of the chapter will guide you in transforming the above program into a more robust and even efficient game playing program for Tic Tac Toe. However, study the program structure carefully and you will be able to use the same strategy, including much of the core of the program, to play many other two-person board games.

Smarter Paper Scissors Rock

In Chapter 7, you saw an example of a program that played the game of Paper Scissors Rock against a human user. In that version, the program's choice strategy for picking an object was completely random. We reproduce that section of the program here:

```
...
string items[] = {"Paper", "Scissors", "Rock"};

...
// Computer makes a selection
string myChoice = items[rand() % 3];
...
```

In the above program segment, `myChoice` is the program's choice. As you can see, the program uses a random number to select its object. That is, the likelihood of picking any of the three objects is 0.33 or 33%. The game and winning strategies for this game have been extensively studied. Some strategies rely on detecting patterns in human choice behavior. Even though

we may not realize it there are patterns in our seemingly random behavior. Computer programs can easily track such behavior patterns by keeping long histories of player's choices, detect them, and then design strategies to beat those patterns. This has been shown to work quite effectively. It involves recording player's choices and searching through them. Another strategy is to study human choice statistics in this game. Before we present you with some data, do the exercise suggested below:

Do This: Play the game against a few people, Play several dozen rounds. Record the choices made by each player (just write a P/S/R in two columns). Once done, compute the percentages of each object picked. Now read on.

It turns out that most casual human players are more prone towards picking Rock than Paper or Scissors. In fact, various analyses suggest that 36% of the time people tend to pick Rock, 30% Paper, and 34% Scissors. This suggests that RPS is not merely a game of chance there is room for some strategies at winning. Believe it or not, there are world championships of RPS held each year. Even a simple game like this has numerous possibilities. We can use some of this information, for instance, to make our program smarter or better adept at playing the game. All we have to do is instead of using a fair 33% chance of selecting each object we can skew the chances of selection based on people's preferences. Thus, if 36% of the time people tend to pick Rock, it would be better for our program to pick Paper 36% of the time since Paper beats Rock. Similarly, our program should pick Scissors 30% of the time to match the chance of beating Paper, and pick Rock 34% of the time to match the chances of beating Scissors. We can bias the random number generator using these percentages as follows:

First generate a random number in the range 0..99
If the number generated is in the range 0..29, select Scissors (30%)
If the number generated is in the range 30..63, select Rock (34%)
If the number generated is in the range 64..99, select Paper (36%)

The above strategy of biasing the random selection can be implemented as follows:

```
string mySelection() {  
  
    // First generate a random number in the range 0..99  
    int n = rand() % 100;  
  
    // If the n is in range 0..29, select Scissors  
    if (n <= 29)  
        return "Scissors";  
    else if (n <= 63)  
        // if n in range 30..63, select Rock  
        return "Rock";  
    else  
        return "Paper";  
}
```

Do This: Modify your RPS program from Chapter 7 to use this strategy. Play the game several times. Does it perform much better than the previous version? You will have to test this by collecting data from both versions against several people (make sure they are novices!).

Another strategy that people use is based upon the following observation:

After many rounds, people tend to make the move that would have beaten their own previous move.

Say a player picks Paper. Their next pick will be Scissors. A computer program or a player playing against this player should then pick Rock to beat Scissors. Note that the relationship between the choices is cyclical. Paper beats Rock, Rock beats Scissors, and Scissors beat Paper. Therefore, since the player's previous move was Paper, your program can pick Rock in anticipation of the player's pick of Scissors. Try to think over this carefully and make sure your head is not spinning by the end of it. If a player can spot this they can use this as a winning strategy. We will leave the implementation of this strategy as an exercise. The exercises also suggest another strategy.

The point of the above examples is that using strategies in your programs you can make your programs smarter or more intelligent. Deliberately, we have started to use the term intelligence a little more loosely than what Alan Turing

implied in his famous essay. Many people would argue that these programs are not intelligent in the ultimate sense of the word. We agree. However, writing smarter programs is a natural activity. If the programs incorporate strategies or heuristics that people would use when they are doing the same activity, then the programs have some form of artificial intelligence in them. Even if the strategy used by the program is nothing like what people would use, but it would make the program smarter or better, we would call it artificial intelligence. Many people would disagree with this latter claim. To some, the quest for figuring out intelligence is limited to the understanding of intelligence in humans (and other animals). In AI both points of view are quite prevalent and make for some passionate debates among scholars.

Discussion

The very idea of considering a computer as an intelligent device has its foundations in the general purpose nature of computers. By changing the program the same computer can be made to behave in many different ways. At the core of it a computer is just a symbol manipulator: manipulating encodings for numbers, or letters, or images, etc. It is postulated that the human brain is also a symbol manipulator. The foundations of AI lie in the fact that most intelligent systems are physical symbol systems and since a computer is a general purpose symbol manipulator, it can be used for studying or simulating intelligence.

Myro Review

No new Myro functions were introduced in the chapter.

C++ Review

`TYPE & NAME`

Indicates that the specified parameter is to be passed by reference (rather than passed by value), which allows it to be modified from within the function.

`typedef TYPE NAME`

Allows `NAME` to be used as an abbreviation for `TYPE`.

Exercises

1. Read Alan Turing's paper Computing Machinery and Intelligence. You can easily find a copy of it by searching on the web.
2. Do a web search for "Searle Chinese Room argument" to locate Philosopher John Searle's arguments that no matter how intelligent a computer or a program gets, it will never have a "mind".
3. Rewrite display for Tic Tac Toe game to display the board graphically.
4. Design a language of one-word English commands for the Scribbler. Write a program to input one command at a time, interpret it, and then execute the command on the Scribbler.
5. Extend the language from Exercise 4 to include queries (e.g. wall?) and then modify your program to incorporate such queries.

6. Do a survey of speech understanding systems.
7. Do a survey of computational linguistics.
8. In the Tic Tac Toe program designed in this Chapter, we assumed that the user always plays an X. Modify your program so that it gives the user a choice at the beginning of the game. Further, at the end of each game, the pieces are swapped.
9. In the function `move` defined for Tic Tac Toe, the program accepts whatever the user inputs for their move. Try the program and instead of entering a valid move, enter your name instead. What happens? Such an error might be easily detected since it will halt the program's execution. However, next try entering a number from 1-9 using a spare position that is already occupied in the board. What happens? Modify the function to accept only correct moves. If the user enters an incorrect move, the program should point that out and give the user another chance.
10. The function `gameOver` can make use of the `winner` function to make its decision in the Tic Tac Toe program. Rewrite `gameOver` to do this.
11. One way to measure how one strategy compares against another is to play it against another strategy over and over again recording the number of wins, losses, and draws. Modify your Tic Tac Toe or RPS program to substitute the second strategy for the user (instead of taking input from the user, it uses a function that implements the second strategy). Add statements to play the game many times (say 1000) and record the wins, losses, and draws. You may also want to suppress all board output since the game is being played completely inside the program. Do a comparison of all the strategies discussed in this Chapter and how they compare against each other.



11

Computers & Computation

Home computers are being called upon to perform many new functions, including the consumption of homework formerly eaten by the dog.

Doug Larson

Computer Science is no more about computers than astronomy is about telescopes.

Edsger W. Dijkstra

What is the central core of computer science? My answer is simple—it is the art of programming a computer. It is the art of designing efficient and elegant methods of getting a computer to solve problems, theoretical or practical, small or large, simple or complex. It is the art of translating this design into an effective and accurate computer program.

C.A.R. Hoare

Opposite page: The XO Laptop

Photo courtesy of OLPC Project (www.olpc.org)

Today, there are more computers than people on most college campuses in the United States. The laptop computer shown on the previous page is the XO laptop developed by the One Laptop Per Child Project (OLPC). It is a low cost computer designed for kids. The OLPC project aims to reach out to over 2 billion children all over the world. They aim to get the XO computers in their hands to help in education. Their mission is to bring about a radical change in the global education system by way of computers. Of course a project with such a grandiose mission is not without controversy. It has had a bumpy ride since its inception. The technological issues were the least of their problems. They have had to convince governments (i.e. politicians) to buy-in into the mission. Several countries have agreed to buy-in and then reneged for all kinds of reasons. The laptops are not available for open purchase within the United States which has led to other socio-political controversies. Regardless, in the first year of its release the project aims to have over 1 million XO's in the hands of children in many developing countries. The thing about technology that has always been true is that a good idea is infective. Other players have emerged who are now developing ultra-low cost computers for kids. It will not be long before other competitive products will be available to all. The bigger questions are: What kind of change will this bring about in the world? How will these computers be used in teaching elementary and middle school children? Etc. You may also be wondering if your Scribbler robot can be controlled by the XO. It can.

You have been writing programs to control your robot through a computer. Along the way you have seen many different ways to control a robot and also ways of organizing your C++ programs. Essentially you have been engaging in the activity commonly understood as *computer programming* or *computer-based problem solving*. We have deliberately refrained from using those terms because the general perception of these conjures up images of geeks with pocket-protectors and complex mathematical equations that seem to be avoidable or out of reach for most people. Hopefully you have discovered by now that solving problems using a computer can be quite exciting and engaging. Your robot may have been the real reason that motivated you into this activity and that was by design. We are confessing to you that robots were used to attract you to pay some attention to computing. You have to agree, if

you are reading this, that the ploy worked! But remember that the reason you are holding a robot of your own in your hand is also because of computers. Your robot itself is a computer. Whether it was a ploy or not you have assimilated many key ideas in computing and computer science. In this chapter, we will make some of these ideas more explicit and also give you a flavor for what *computer science* is really all about. As Dijkstra puts it, *computer science is no more about computers than astronomy is about telescopes*.

Computers are dumb

Say you are hosting an international exchange student in your home. Soon after her arrival you teach her the virtues of PB&J (Peanut Butter & Jelly) sandwiches. After keenly listening to you, her mouth starts to water and she politely asks you if you can share the recipe with her. You write it down on a piece of paper and hand it to her.

Do This: Go ahead; write down the recipe to make a PB&J sandwich.

Seriously, do try to write it down. We insist!

OK, now you have a recipe for making PB&J sandwiches.

Do This: Go ahead, use your recipe from above to make yourself a PB&J sandwich. Try to follow the instructions *exactly* as *literally* as you can. If you successfully managed to make a PB&J sandwich, congratulations! Go ahead and enjoy it. Do you think your friend will be able to follow your recipe and enjoy PB&J sandwiches?

You have to agree, writing a recipe for PB&J sandwiches seemed trivial at first, but when you sit down to write it you have no choice but to question several assumptions: does she know what peanut butter is? Should I recommend a specific brand? Ditto for jelly? By the way, did you forget to mention it was *grape jelly*? What kind of bread to use? Will it be pre-sliced? If not, you need a knife for slicing the loaf? Did you specify how thick the

slices should be? Should she use the same knife for spreading the peanut butter and the jelly? Etc. The thing is, at such a level of detail you can go on and on...does your friend know how to use a knife to spread butter or jelly on a slice? Suddenly, a seemingly trivial task becomes a daunting exercise. In reality, in writing down your recipe, you make several assumptions: she knows what a sandwich is, it involves slices of bread, spreading things on the slices, and slapping them together. There, you have a sandwich!

Think of the number of recipes that have been published in cookbooks all over the world. Most good cookbook authors start with a dish, write down a recipe, try it a number of times and refine it in different ways. Prior to publication, the recipe is tested by others. After all, the recipe is for others to follow and recreate a dish. The recipe is revised and adjusted based on feedback by recipe testers. The assumption that cookbook authors make is that you will have the competence to follow the recipe and recreate the dish to your satisfaction. It may never result in the same dish that the author prepared. But it will give you a base to improvise upon. Wouldn't it be nice if there was an exact way of following a recipe so that you end up with the exact same result as the cookbook author every time you made that dish? Would it? That may depend on your own tastes and preferences. Just for fun, here is a recipe for some yummy Saffron Chicken Kabobs.

Saffron Chicken Kabobs

Ingredients

- 1 lb boneless chicken breast, cubed into 1-2 inch pieces
- 1 medium onion, sliced
- 1 tsp saffron threads
- 1 lime
- 1 tbsp olive oil
- Salt and black pepper to taste

Preparation

1. Mix the chicken and onions in a non-reactive bowl.
2. With your fingers crush and add saffron threads.
3. Add the juice of the lime, olive oil, and salt and pepper.
4. Marinade in the refrigerator for at least 30 min (or overnight).
5. Preheat a grill (or oven to 400 degrees).
6. Skewer kabobs, discarding the onion slices. Or place everything in a lined baking sheet if using oven.
7. Grill/bake for 12-15 min until done.

In cooking recipes, like the ones above, you can assume many things: they will be used by people (like you and me); they will be able to follow them; perhaps even improvise. For instance, in the recipe above, we do not specify that one will need a knife to cut the chicken, onions, or the lime; or that you will need a grill or an oven; etc. Most recipes assume that you will be able to *interpret* and follow the recipe as written.

Computer programs are also like recipes, to some extent. Think of the program you wrote for choreographing a robot dance, for instance. We have reproduced the version from Chapter 3 here:

```
/*
 * File: dance.cpp
 * Purpose: A simple dance routine
 */

// First include standard declarations

#include "Myro.h"

// Define the new functions...

void yoyo(double speed, double waitTime) {
    robot.forward(speed, waitTime);
    robot.backward(speed, waitTime);
}
```

```
void wiggle(double speed, double waitTime) {
    robot.motors(-speed, speed);
    wait(waitTime);
    robot.motors(speed, -speed);
    wait(waitTime);
    robot.stop();
}

// The main dance program
int main() {
    connect();
    cout << "Running the dance routine..." << endl;
    yoyo(0.5, 0.5);
    wiggle(0.5, 0.5);
    yoyo(1, 1);
    wiggle(1, 1);
    cout << "...Done" << endl;
    disconnect();
}
```

In many ways, this program above is like a recipe:

To do a robot dance

Ingredients

- 1 function `yoyo` for the robot to go back and forth at a given speed
- 1 function `wiggle` that enables the robot to wiggle at a given speed

Preparation

1. `yoyo` at speed 0.5, wait 0.5
2. `wiggle` at speed 0.5, wait 0.5
3. `yoyo` at speed 1, wait 1
4. `wiggle` at speed 1, wait 1

Further, you could similarly specify the steps involved in doing the `yoyo` and `wiggle` motions as a recipe. This may seem like a trivial example, but it makes two very important points: a computer program is like a recipe in that it

lays out the list of ingredients and a method or steps for accomplishing the given task; and, like a recipe, its ingredients and the steps require careful pre-planning and thought. Importantly, computer programs are different from recipes in one aspect: they are designed to be followed by a computer!

A computer is a dumb device designed to follow instructions/recipes. We will save the technical details of how a computer does what it does for a later course. But it is almost common knowledge that everything *inside* is represented as 0's and 1's. Starting from 0's and 1's one can design encoding schemes to represent numbers, letters of the alphabet, documents, images, movies, music, etc. and whatever other abstract entities you would like to manipulate using a computer. A computer program is ultimately also represented as a sequence of 0's and 1's and it is in this form that most computers like to follow recipes. However limiting or degenerate this might sound it is the key to the power of computers. Especially when you realize that it is this simplification that enables a computer to manipulate hundreds of millions of pieces of information every second. The price we have to pay for all this power is that we have to specify our recipes as computer programs in a rather formal and precise manner. So much so that there is no room for improvisation: no pinch of salt vagaries, as in cooking recipes, is acceptable. This is where *programming languages* come in. Computer scientists specify their computational recipes using *programming languages*. You have been using the programming language C++ to write your robot programs. Other examples of programming languages are Java, Python, and C# (pron.: sea sharp). There are well over 2000 programming languages in existence!

Do This: Can you find out how many programming languages there are? What are the ten most commonly used programming languages?

Prior to the existence of programming languages computers were programmed using long sequences of 0's and 1's. Needless to say it drove several people crazy! Programming languages, like C++, enable a friendlier way for programmers to write programs. Programming languages provide easy access to encodings that represent the kinds of things we, humans, relate to. For example, the C++ statement:

```
meaningOfLife = 42;
```

is a command for the computer to assign the value, 42 to the variable named `meaningOfLife`. This way, we can ask the computer to check that it is indeed 42:

```
if (meaningOfLife == 42)
    speak("Eureka!");
else
    speak("What do we do now?");
```

Once again, it would be good to remind you that the choice of the name, `meaningOfLife`, doesn't really mean that we are talking about *the meaning of life*. We could as well have called it `timbuktoo`, as in:

```
timbuktoo = 42;
```

You see, computers are truly dumb!

It is really up to us, the programmer, to ensure that we use our names consistently and choose them, in the first place, carefully. But, by creating a language like C++, we have created a formal notation so that when translated into 0's and 1's each statement will mean only one thing, no other interpretations. This makes them different from a cooking recipe.

Robot goes to buy fresh eggs

Recipes, however, form a good conceptual basis for starting to think about a program to solve a problem. Say, you have in mind to make your favorite *Apple Strudel*. You know you will need apples. Perhaps it is the apple season that prompted the thought in the first place. You will also need pastry. But when you get down to it, you will need that recipe you got from your grandma.

Whenever we are asked to solve a problem using a computer, we begin by laying out a rough plan for solving the problem. That is, sketch out a strategy.

This is further refined into specific steps, perhaps even some variables are identified and named, etc. Once you convince yourself that you have a way of solving the problem, what you have is an *algorithm*

The idea of an algorithm is central to computer science so we will spend some time here developing this notion. Perhaps the best way to relate to it is by an example. Assume that a robot goes into a grocery store to buy a dozen fresh eggs. Assuming it is capable of doing this, how will it ensure that it has selected the freshest eggs available?



Your personal robot is probably not up to this kind of task but imagine that it was. Better yet, leave the mechanics aside, let us figure out how *you* would go and buy the freshest eggs. Well, you would somehow need to know what today's date is. Assume it is September 15, 2007 (why this date? it'll become clear soon!). Now you also know that egg cartons typically carry a freshness date on them. In fact, USDA (the United States Department of Agriculture) offers voluntary, no cost, certification programs for egg farms. An egg farmer can volunteer to participate in USDA's egg certification program whereby the USDA does regular inspections and also provides help in categorizing eggs by various sizes. For example, eggs are generally classified as Grade AA, Grade A, or Grade B. Most grocery stores carry Grade A eggs. They can also come in various sizes: Extra Large, Large, Small, etc. What is more interesting is that the carton labeling system has some very useful information encoded on it.

Egg Carton Labeling



Every USDA certified egg carton has at least three pieces of information (see picture on the

right): a "sell by" date (or a "use by date" or a "best by" date), a code identifying the specific farm the eggs came from, and a date on which the eggs were packed in that carton. Most people buy eggs by looking at the "sell by" date or the "best by" date. However the freshness information is really encoded in the packed on date. To make things more confusing, this date is encoded as the day of the year.

For example, take a look at the top carton shown on the previous page. Its "sell by" date is October 4. "P1107" is the farm code. This carton was packed on the 248th day of the year. Further, USDA requires that all certified eggs be packed within 7 days of being laid. Thus, the eggs in the top carton were laid somewhere between day 241 and day 248 of 2007. What dates correspond to those dates?

Next, look at the bottom carton. Those eggs have a later "sell by" date (October 18) but an earlier packed date: 233. That is those eggs were laid somewhere between day 226 and day 233 of 2007.

Which eggs are fresher?

Even though the "sell by" date on the second carton is two weeks later, the first carton contains fresher eggs. In fact, the eggs in the upper carton were laid at least two weeks later!

The packed on date is encoded as a 3-digit number. Thus eggs packed on January 1 will be labeled: 001; eggs packed on December 31, 2007 will be labeled: 365.

Do This: Go to the USDA web site (www.usda.gov) and see if you can find out which farm the two eggs cartons came from.

For a robot, the problem of buying the freshest eggs becomes that of figuring out, given a packed on date, what the date was when the eggs were packed?

Fasten your seatbelts, we are about to embark on a unique computational voyage...

Designing an algorithm

So far, we have narrowed the problem down to the following specifications:

Input

3-digit packed on date encoding

Output

Date the eggs were packed

For example, if the packed on date was encoded as 248, what will be the actual date?

Well, that depends. It could be September 4 or September 5 depending on whether the year was a leap year or not. Thus, it turns out, that the problem above also requires that we know which year we were talking about. Working out one or two sample problems is always a good idea because it helps identify missing information that may be critical to solving the problem. Given that we do need to know the year, we can ask the user to enter that at the same time the 3-digit code is entered. The problem specification then becomes:

The Etymology of Algorithm

The word algorithm, an anagram of logarithm, is believed to have been derived from Al-Khowarizmi, a mathematician who lived from 780-850 AD. His full name was Abu Ja'far Muhammad ibn Mūsā al-Khwārizmī, (Mohammad, father of Jafar, son of Moses, a Khwarizmian). Much of the mathematical knowledge of medieval Europe was derived from Latin translations of his works.



In 1983, The Soviet Union issued the stamp shown above in honor of his 1200th anniversary.

Input

3-digit packed on date encoding
Current year

Output

Date the eggs were packed

Example:

Input: 248, 2007

Output: The eggs were packed on September 5, 2007

Any ideas as to how you would solve this problem? It always helps to try and do it yourself, with pencil and paper. Take the example above, and see how you would arrive at the output date. While you are working it out, try to write down your problem solving process. Your algorithm or recipe will be very similar.

Suppose we are trying to decode the input 248, 2007. If you were to do this by hand, using a pen and paper, the process might go something like this:

The date is not in January because it has 31 days and 248 is much larger than 31.

Lets us subtract 31 out of 248: $248 - 31 = 217$

217 is also larger than 28, the number of days in February, 2007.

So, let us subtract 28 from 217: $217 - 28 = 189$

189 is larger than 31, the number of days in March.

Subtract 31 from 189: $189 - 31 = 158$

158 is larger than 30, the number of days in April.

So: $158 - 30 = 128$

128 is larger than 31, the number of days in May.

Hence: $128 - 31 = 97$

97 is larger than 30, the number of days in June.

97 - 30 = 67

67 is larger than 31, the number of days in July.

67 - 31 = 36

36 is larger than the number of days in August (31).

36 - 31 = 5

5 is smaller than the number of days in September.

Therefore it must be the 5th day of September.

The answer is: 248th day of 2007 is September 5, 2007.

That was obviously too repetitious and tedious. But that is where computers come in. Take a look at the process above and see if there is a pattern to the steps performed. Sometimes, it is helpful to try another example.

Do This: Suppose the input day and year are: 56, 2007. What is the date?

When you look at the sample computations you have performed, you will see many patterns. Identifying these is the key to designing an algorithm. Sometimes, in order to make this easier, it is helpful to identify or name the key pieces of information being manipulated. Generally, this begins with the inputs and outputs identified in the problem specification. For example, in this problem, the inputs are: day of the year, current year. Begin by assigning these values to specific variable names. That is, let us assign the name `day` to represent the day of the year (248 in this example), and `year` as the name to store the current year (2007). Notice that we didn't choose to name any of these variables `timbuktu` or `meaningOfLife`!

Also, notice that you have to repeatedly subtract the number of days in a month, starting from January. Let us assign a variable named, `month` to keep track of the month under consideration.

Next, you can substitute the names `day` and `year` in the sample computation:

Input:

```
    day = 248;  
    year = 2007;
```

```
// Start by considering January  
month = 1;
```

The date is not in month = 1 because it has 31 days and 248 is much larger than 31.

```
day = day - 31;
```

```
// next month
```

```
month = 2;
```

day (= 217) is also larger than 28, the # of days in month = 2

```
day = day - 28;
```

```
// next month
```

```
month = 3;
```

day (= 189) is larger than 31, the # of days in month = 3.

```
day = day - 31;
```

```
// next month
```

```
month = 4;
```

day (= 158) is larger than 30, the # of days in month = 4.

```
day = day - 30;
```

```
// next month
```

```
month = 5;
```

day (= 128) is larger than 31, the # of days in month = 5.

```
day = day - 31;
```

```
// next month
```

```
month = 6;
```

day (= 97) is larger than 30, the # of days in month = 6.

```
day = day - 30;
```

```
// next month
```

```
month = 7;
```

day (= 67) is larger than 31, the # of days in month = 7.

```
day = day - 31;
```

```
// next month
```

```
month = 8;
```



```
day (= 36) is larger than the # of days in month = 8.  
day = day - 31;
```

```
// next month  
month = 9;  
day (= 5) is smaller than the # of days in month = 9.  
Therefore it must be the 5th day of September.
```

The answer is: 9/5/2007

Notice now how repetitious the above process is. The repetition can be expressed more concisely as shown below:

```
Input:  
  day  
  year  
  
// start with month = 1, for January  
month = 1;  
repeat  
  if day is less than number of days in month  
    day = day - number of days in month;  
    // next month  
    month = month + 1;  
  else  
    done;
```

Output: day/month/year

It is now starting to look like a recipe or an algorithm. Go ahead and try it with the sample inputs from above and ensure that you get correct results. Additionally, make sure that this algorithm will work for boundary cases: 001, 365.

Thirty days hath September

We can refine the algorithm above further: one thing we have left unspecified above is the computation of the number of days in a month. This information has to be made explicit for a computer to be able to follow the recipe. So, how

do we compute the number of days in a month? The answer may seem simple. Many of you may remember the following poem:

*Thirty days hath September
April, June, and November
All the rest have thirty-one
Except for February alone
Which hath twenty-eight days clear
And twenty-nine in each leap year*

From a design perspective, we can assume that we have an ingredient, a function in this case, called `daysInMonth` that, given a month and a year will compute and return the number of days in the month. That is, we can refine our algorithm above to the following:

Ingredients:

```
1 integer function daysInMonth(int m, int y):  
    returns the number of days in month m and in year y.
```

Input:

```
    an integer, day  
    an integer, year
```

```
// start with month = 1, for January  
month = 1;  
repeat  
    if day is less than number of days in month  
        day = day - daysInMonth(month, year);  
        // next month  
        month = month + 1;  
    else  
        done;
```

Output: day/month/year (integers)

Now, we do have to solve the secondary problem:

Input

month, M (an integer)

year, Y (an integer)

Output

Number of days in month M and in year Y

On the surface this seems easy, the poem above specifies that April, June, September, and November have 30 days, and the rest, with the exception of February have 31. February has 28 or 29 days depending upon whether it falls in a leap year or not. Thus, we easily elaborate a recipe or an algorithm for this as follows:

Input:

```
int m, y;
```

```
if m is April(4), June(6), September(9), or November(11)
```

```
    days = 30;
```

```
else if m is February (2)
```

```
    if y is a leap year
```

```
        days = 29;
```

```
    else
```

```
        days = 28;
```

```
else
```

```
    /* m is January, March, May, July, August, October,
       December */
```

```
    days = 31;
```

Output:

```
days (an integer)
```

This still leaves out one more detail: how do we tell if y is a leap year?

First, try and answer the question, *what is a leap year?*

Again, we can refine the algorithm above by assuming that we have another ingredient, a Boolean function: `leapYear`, that determines if a given year is a leap year or not. Then we can write the algorithm above as:

Ingredients:

```
1 function bool leapYear(int y)
   returns true if y is a leap year, false otherwise
```

Input:

```
int m, y;
```

```
if m is April(4), June(6), September(9), or November(11)
    days = 30;
else if m is February (2)
    if (leapYear(y))
        days = 29;
    else
        days = 28;
else
    /* m is January, March, May, July, August, October,
       December */
    days = 31;
```

Output:

```
days (an integer)
```

Most of us have been taught that a leap year is a year that is divisible by 4. That is the year 2007 is not a leap year, since 2007 is not divisible by 4, but 2008 is a leap year, since it is divisible by 4.

Do This: How do you determine if something is divisible by 4? Try your solution on the years 1996, 2000, 1900, 2006.

Leap Years: Papal Bull

To design a recipe or an algorithm that determines if a number corresponding to a year is a leap year or not is straightforward if you accept the definition from the last section. Thus, we can write:

Input:

```
int y, a year
```

Output:

```
true if y is a leap year, false otherwise
```

Method

```
if y is divisible by 4
    it is a leap year, or true
else
    it is not a leap year, or false
```

However, this is not the complete story. The western calendar that we follow is called the *Gregorian Calendar* which was adopted in 1582 by a Papal Bull issued by Pope Gregory XIII. The Gregorian Calendar defines a leap year, by adding an extra day, every fourth year. However, there is a 100-year correction applied to it that makes the situation a little more complicated: Century years are not leap years except when they are divisible by 400. That is the years 1700, 1800, 1900, 2100 are not leap years even though they are divisible by 4. However, the years 1600, 2000, 2400 are leap years. For more information on this, see the exercises at the end of the chapter. Our algorithm for determining if a year is a leap year can be refined as shown below:

```
Input:
    int y, a year

    if y is divisible by 400
        it is a leap year, or true
    else if y is divisible by 100
        it is not a leap year, or false
    else if y is divisible by 4
        it is a leap year, or true
    else
        it is not a leap year, or false
```

Finally, we have managed to design all the algorithms or recipes required to solve the problem. You may have noticed that we used some familiar constructs to elaborate our recipes or algorithms. Next, let us take a quick look at the essential constructs that are used in expressing algorithms.

Essential components of an algorithm

Computer scientists express solutions to problems in terms of algorithms, which are basically more detailed recipes. Algorithms can be used to express any solution and yet are comprised of some very basic elements:

1. Algorithms are step-by-step recipes that clearly identify the inputs and outputs
2. Algorithms name the entities that are manipulated or used: variables, functions, etc.
3. Steps in the algorithm are followed in the order they are written (from top to bottom)
4. Some steps can specify decisions (if-then) over the choice of some steps
5. Some steps can specify repetitions (loops) of steps
6. All of the above can be combined in any fashion.

Computer scientists claim that solutions/algorithms to *any* problem can be expressed using the above constructs. You do not need any more! This is a powerful idea and it is what makes computers so versatile. From a larger perspective, if this is true, then these can be used as tools for thinking about any problem in the universe. We will return to this later in the chapter.

Programming Languages

Additionally, as you have seen earlier, in writing C++ programs, programming languages (C++, for example) provide formal ways of specifying the essential components of algorithms. For example, the C++ language provides a way for you to assign values to variables that you name, it provides a sequential way of encoding the steps, it provides the if-then conditional statements, and also provides the while-loop and for-loop constructs for expressing repetitions. C++ also provides means for defining functions and also ways of organizing groups of related functions into libraries or classes which you can include and use as needed. As an example,

we provide below, the C++ (or Java) program that encodes the `leapYear` algorithm shown above:

```
bool leapYear(int y) {
    // Returns true if y is a leap year, false otherwise.
    if (y % 400 == 0)
        return true;
    else if (y % 100 == 0)
        return false;
    else if (y % 4 == 0)
        return true;
    else
        return false;
}
```

The same algorithm, when expressed in Python will look like this:

```
def leapYear(y):
    '''Returns true if y is a leap year, false otherwise.'''

    if y % 400 == 0:
        return True
    elif y % 100 == 0:
        return False
    elif y % 4 == 0:
        return True
    else:
        return False
```

As you can see, there are definite syntactic variations among programming languages. But, at least in the above examples, the coding of the same algorithm looks very similar. Just to give a different flavor, here is the same function expressed in the programming language CommonLisp.

```
(defun leapYear (y)
  (cond
    ((zerop (mod y 400)) t)
    ((zerop (mod y 100)) nil)
    ((zerop (mod y 4)) t)
    (t nil)))
```

Again, this may look weird, but it is still expressing the same algorithm.

What is more interesting is that given an algorithm, there can be many ways to encode it, even in the same programming language. For example, here is another way to write the `leapYear` function in C++:

```
bool leapYear(int y) {
    // Returns true if y is a leap year, false otherwise.

    if (((y % 4 == 0) && (y % 100 != 0)) || (y % 400 == 0))
        return true;
    else
        return false;
}
```

Again, this is the same exact algorithm. However, it combines all the tests into a single condition: `y` is divisible by 4 or by 400 but not by 100. The same condition can be used to write an even more succinct version:

```
bool leapYear(int y) {
    // Returns true if y is a leap year, false otherwise.

    return ((y % 4 == 0) && (y % 100 != 0)) || (y % 400 == 0);
}
```

That is, return whatever the result is (`true/false`) of the test for `y` being a leap year. In a way, expressing algorithms into a program is much like expressing a thought or a set of ideas in a paragraph or a narrative. There can be many ways of encoding an algorithm in a programming language. Some seem more natural, and some more poetic, or both, and, like in writing, some can be downright obfuscated. As in good writing, good programming ability comes from practice and, more importantly, learning from reading well written programs.

From algorithms to a working program

To be able to solve the fresh eggs problem, you have to encode all the algorithms into C++ functions and then put them together as a working program. Below, we present one version:

```
// File: fresheggs.cpp

#include <iostream>
using namespace std;

bool leapYear(int y) {
    // Returns true if y is a leap year, false otherwise.

    return ((y % 4 == 0) && (y % 100 != 0)) || (y % 400 == 0);
}

int daysInMonth(int m, int y) {
    /* Returns the number of days in month, m (1-12)
       in year, y. */

    if ((m == 4) || (m == 6) || (m == 9) || (m == 11))
        return 30;
    else if (m == 2) {
        if (leapYear(y))
            return 29;
        else
            return 28;
    } else
        return 31;
}

int main() {
    /* Given a day of the year (e.g. 248, 2007),
       convert it to the date (i.e. 9/5/2007) */

    //Input: day, year
    int day, year;
    cout << "Enter day year: ";
    cin >> day >> year; // read both in one statement

    // start with month = 1, for January
```

```
int month = 1;

while (day > daysInMonth(month, year)) {
    day = day - daysInMonth(month, year);

    // next month
    month++;
}

// done, Output: month/day/year
printf ("The date is: %1d/%1d/%4d\n", month, day, year);

return 0;
}
```

If you save this program in a file, `fresheggs.cpp`, and compile it, you will be able to run it and test it for various dates. Go ahead and do this. Here are some sample outputs:

fresheggs

```
Enter day year: 248 2007
The date is: 9/5/2007
```

fresheggs

```
Enter day year: 12 2007
The date is: 1/12/2007
```

fresheggs

```
Enter day year: 248 2008
The date is: 9/4/2008
```

fresheggs

```
Enter day year: 365 2007
The date is: 12/31/2007
```

fresheggs

```
Enter day year: 31 2007
```

```
The date is: 1/31/2007
```

All seems to be good. Notice how we tested the program for different input values to confirm that our program is producing correct results. It is very important to test your program for a varied set of input, taking care to include all the *boundary* conditions: first and last day of the year, month, etc. Testing programs is a fine art in itself and several books have been written about the topic. One has to ensure that all possible inputs are tested to ensure that the behavior of the program is acceptable and correct. You did this with your robot programs by repeatedly running the program and observing the robot's behavior. Same applies to computation.

Testing and Error Checking

What happens, if the above program receives inputs that are outside the range? What if the user enters the values backwards (e.g. 2007 248 instead of 248 2007)? What if the user enters her name instead (e.g. Paris Hilton)? Now is the time to try all this out. Go ahead and run the program and observe its behavior on some of these inputs.

Ensuring that a program provides acceptable results for all inputs is critical in most applications. While there is no way to avoid what happens when a user enters his name instead of entering a day and a year, you should still be able to safeguard your programs from such situations. For example:

fresheggs

```
Enter the day year: 400 2007
```

```
That corresponds to the date: 14/4/2007
```

Obviously, we do not have a month numbered 14!

The thing that comes to rescue here is the realization that, it is your program and the computer will only carry out what you have expressed in the program. That is, you can include *error checking* facilities in your program to account for such conditions. In this case, any input value for day that is outside the range 1..365 (or 1..366 for leap years) will not be acceptable. Additionally, you can also ensure that the program only accepts years greater than 1582 for the second input value. Here is the modified program (we'll only show the main function):

```
int main() {
    /* Given a day of the year (e.g. 248, 2007),
       convert it to the date (i.e. 9/5/2007) */

    //Input: day, year
    int day, year;
    cout << "Enter day year: ";
    cin >> day >> year; // read both in one statement

    // Validate input values...
    if (year <= 1582) {
        cout << "I'm sorry. You must enter a valid year "
             << "(one after 1582). Please try again.\n";
        exit(EXIT_FAILURE);
    } else if (day < 1) {
        cout << "I'm sorry. You must enter a valid day "
             << "(1..365/366). Please try again.\n";
        exit(EXIT_FAILURE);
    } else if (leapYear(year)) {
        if (day > 366) {
            cout << "I'm sorry. You must enter a valid day "
                 << "(1..365/366). Please try again.\n";
            exit(EXIT_FAILURE);
        }
    } else if (day > 365) {
        cout << "I'm sorry. You must enter a valid day "
             << "(1..365/366). Please try again.\n";
        exit(EXIT_FAILURE);
    }
}
```

```
// input values are safe, proceed...
// start with month = 1, for January
int month = 1;

while (day > daysInMonth(month, year)) {
    day = day - daysInMonth(month, year);

    // next month
    month++;
}

// done, Output: month/day/year
printf ("The date is: %ld/%ld/%4d\n", month, day, year);

return 0;
}
```

Notice that if an illegal date is entered, the program invokes `exit(EXIT_FAILURE)`. This stops the program and signals the operating system that the `fresheggs` program terminated abnormally. (As you know, a program can terminate normally by executing `return 0` in `main`; it can also terminate normally by invoking `exit(EXIT_SUCCESS)` anywhere.)

Here are the results of some more tests on the above program:

fresheggs

```
Enter day year: 248 2007
The date is: 9/5/2007
```

fresheggs

```
Enter day year: 0 2007
I'm sorry. You must enter a valid day (1..365/366). Please try
again.
```

fresheggs

```
Enter day year: 366 2007
I'm sorry. You must enter a valid day (1..365/366). Please try
again.
```

fresheggs

```
Enter day year: 400 2007
```

```
I'm sorry. You must enter a valid day (1..365/366). Please try again.
```

fresheggs

```
Enter day year: 248 1492
```

```
I'm sorry. You must enter a valid year (one after 1582). Please try again.
```

fresheggs

```
Enter day year: 366 2008
```

```
The date is: 12/31/2008
```

Starting from a problem description it is a long and carefully planned journey that involves the development of the algorithm, the encoding of the algorithm in a program, and finally testing and improving the program. In the end you are rewarded not just by a useful program, you have also honed your general problem solving skills. Programming forces you to anticipate unexpected situations and to account for them prior to encountering them which itself can be a wonderful life lesson.

Modules to organize components

Often, in the course of designing a program, you end up designing components or functions that can be used in many other situations. For example, in the problem above, we wrote functions `leapYear` and `daysInMonth` to assist in solving the problem. You will no doubt agree that there are many situations where these two functions could come in handy (see Exercises below). C++ provides the *header file* facility to help organize related useful functions into a single library that you can then use over and over whenever they are needed. For example, you can take the definitions of the two functions and put them separately in a file called, `calendar.h`. Then, you can *include* these functions whenever you need them. You have used the C++ `#include` statement to include functionality from several different modules: `Myro`, `iostream`, etc. Well, now you know how to create your own.

Once you create the `calendar.h` file, you can import it in the `fresheggs.cpp` program as shown below:

```
// File: fresheggs.cpp

#include <iostream>
using namespace std;
#include "calendar.h"

int main() {
    /* Given a day of the year (e.g. 248, 2007),
       convert it to the date (i.e. 9/5/2007) */

    //Input: day, year
    int day, year;
    cout << "Enter day year: ";
    cin >> day >> year; // read both in one statement

    // Validate input values...
    if (year <= 1582) {
        cout << "I'm sorry. You must enter a valid year "
             << "(one after 1582). Please try again.\n";
        exit(EXIT_FAILURE);
    } else if (day < 1) {
        cout << "I'm sorry. You must enter a valid day "
             << "(1..365/366). Please try again.\n";
        exit(EXIT_FAILURE);
    } else if (leapYear(year)) {
        if (day > 366) {
            cout << "I'm sorry. You must enter a valid day "
                 << "(1..365/366). Please try again.\n";
            exit(EXIT_FAILURE);
        }
    } else if (day > 365) {
        cout << "I'm sorry. You must enter a valid day "
             << "(1..365/366). Please try again.\n";
        exit(EXIT_FAILURE);
    }

    // input values are safe, proceed...
    // start with month = 1, for January
    int month = 1;
```

```
while (day > daysInMonth(month, year)) {
    day = day - daysInMonth(month, year);

    // next month
    month++;
}

// done, Output: month/day/year
printf ("The date is: %1d/%1d/%4d\n", month, day, year);

return 0;
}
```

It may have also occurred to you by now that for any given problem there may be many different solutions or algorithms. In the presence of several alternative algorithms how do you decide which one to choose? Computer Scientists have made it their primary business to develop, analyze, and classify algorithms to help make these decisions. The decision could be based on ease of programming, efficiency, or the number of resources it takes for a given algorithm. This has also led computer scientists to create a classification of problems: from easy to hard, but in a more formal sense. Some of the hardest open questions in the realm of problems and computing lie in this domain of research. We will not go into details here, but these questions have even shown up in several popular TV shows. We will attempt to give you a flavor of this next.

Homer contemplates $P = NP$?



The second equation on the right is Euler's Equation.

Space & Time Complexity

Let us start with another problem: You have to travel from Washington State to Washington DC in the United States of America. To make things interesting, let's add a restriction that you can only travel through states whose names begin with the letters in the word "woman". That is, it is OK to go from Washington to Oregon since both "W" and "O" are in the word "woman" but it is not OK to go from Washington to California. Is this feasible? If it is, how many ways are possible? Which one goes through the least/most number of states? Etc.

If you are thinking about how to solve this, you have to rely on your geographic knowledge of the United States. Alternately, you can Google a state map and then figure out a solution. But, in doing so, you have stumbled upon the two key ingredients of computing: data & algorithm.

Data gives you a representation of the relevant information which, in this case is a way of knowing which states adjoin which other states. The algorithm gives you a way of finding a solution: Start in Washington, look at its neighbors. Pick a neighbor whose name satisfies the constraint, then look at that state's neighbors, and so on. Additionally, an algorithm also forces you to make certain choices: if there is more than one eligible neighbor, which one do you pick? What if you end up in a dead end, how do you go back to the choices you left behind to explore alternative paths? Etc. Depending on these decisions, you are likely to end up with many different algorithms: one that may suggest exploring all alternatives simultaneously; or one that forces you to choose, only to return to other choices in case of failure. Each of these will further impact on the amount of data you will need to store to keep track of your progress.

Developing computer programs to solve any problem requires one to design data representations and to choose among a set of alternative algorithms. Computer scientists characterize choices of data representations and algorithms abstractly in terms of the computer resources of *space* and *time* needed to implement those choices. Solutions vary in terms of the amount of

space (or computer memory) and time (seconds, minutes, hours, days, years) required on a computer. Here are some examples:

- To compute the product of two numbers requires constant time: to a computer there is no difference between multiplying 5 by 2 or 5,564,198 by 9,342,100. These are called *constant time* algorithms.
- To find a number in a list of N unordered numbers takes time proportional to N , especially if the number you are looking for is not in there. These are called *linear time* algorithms.
- To find a number in a list of N ordered numbers (say, in ascending order) requires at most $\log_2 N$ time. How? Such algorithms are called *logarithmic time* algorithms.
- To transform a $N \times N$ pixel camera image by manipulating all its pixels takes time proportional to N^2 time. These are called *quadratic algorithms*.
- To find a path from a state to another, in a map of N states, given certain constraints, can take time proportional to b^d where b is the average number of neighbors of each state and d is the number of states that make up the solution. In general, many problems fall into the category N^k where N is the size of the problem. These are called *polynomial time* algorithms.
- There are also several unsolvable problems in the world.

In Chapter 9, when doing image transformations, we restricted ourselves to fairly small sized images. You may have noticed that the larger the image, the longer it takes for the transformation. This has to do with the speed of the computer you are using as well as the speed of the programming language implementation. For example, a computer running at 4GHz speed is capable of doing approximately 1 billion arithmetic operations in one second (or 10^9 operations/second). A program written in the C programming language might be able to give you a speed of $\frac{1}{2}$ billion operations per second on the same computer. This is due to extra operations required to implement the C language and additional tasks the operating system is carrying out in the background. Because they are *interpreted* directly rather than *compiled* into machine instructions, Python programs run approximately 20 times slower

that C programs. That is, a Python program running on the same computer might give you 25 million operations per second at best. A typical transformation, say computing the negative, of an image of size $w \times h$ pixels would require the following loop:

```
for (int x = 0; x < W; x++)
    for (int y = 0; y < H; y++) {
        Pixel pixel = getPixel(myPic, x, y);
        color_rgb rgb = getRGB(pixel);
        setRGB(pixel,
            color_rgb(255-rgb.red, 255-rgb.green, 255-rgb.blue));
    }
```

If the image is 1000x1000 pixels (i.e. $w=1000$ and $h=1000$), each of the three statements in the loop is executed 1 million times. Each of those statements in turn requires an average of 8-16 low-level computing operations: the actual calculations, plus calculations to locate pixels in memory, etc. Thus, the transformation above would require over 24-48 million operations. As you can imagine, it will take a few seconds to complete that task.

In the computing industry, computing speeds are classified based on official benchmark computations that calculate speeds in terms of the number of floating-point operations per second, or *flops*. A typical laptop or a desktop these days is capable of delivering speeds between $\frac{1}{2}$ to 1 Gflops (Giga flops). The world's fastest computer can deliver computing speeds as fast as 500 Tflops (Teraflops). That is, it is about a million times faster (and costs many millions to make as well). However, if you stop and think about the Chess playing program we mentioned in Chapter 10 that would require approximately 10^{65} operations before making a single move, even the world's fastest computer is going to take gazillion years to finish that computation! Such a problem would be considered *uncomputable*.

Computer scientists have developed an elaborate vocabulary for discussing problems and classifying them as *solvable* or *unsolvable*, *computable* or *uncomputable*, based on whether there are known models to solve a given problem and whether the models are solvable, computable, etc. There are also

hierarchies of problem solutions, from simple to hard; constant time to polynomial time and longer; and equivalence classes implying the same algorithm can solve all problems in an equivalence class, etc. It is indeed amazing to conceptualize an algorithm that is capable of solving many unrelated problems! For example, the algorithms that optimize shipping and delivery routes can also be used in determining protein folding structures for DNA molecules. This is what makes computer science intellectually interesting and exciting.

Summary

In this chapter we have tied together many fundamental ideas in computing and computer science. While our journey started in Chapter 1 with playing with personal robots, we have, in the process, acquired a wealth of fundamental concepts in computing. As you can see, computing is a rich, diverse, and deeply intellectual discipline of study that has implications for all aspects of our lives. We started this chapter by pointing out that there are now more computers on a typical college campus in the United States than the number of people. It probably will not be long before there are more computers than people on this entire planet. Yet, the idea of an *algorithm* that is central to computing is barely understood by most users despite its simple and intuitive constituents. Many computer scientists believe that we are still sitting at the dawn of *the age of algorithm* and that there are much bigger intellectual and societal benefits yet to be realized, especially if more people were aware of these ideas.

Myro review

No new Myro features were introduced in this chapter.

C++ Review

```
exit(EXIT_MODE)
```

May be invoked anywhere to exit the program immediately. If `EXIT_MODE` is `EXIT_SUCCESS` then the program exits normally (equivalent to invoking `return(0)` from `main()`); if it is `EXIT_FAILURE` then the program signals abnormal termination to the operating system.

```
#include "FILENAME.h"
```

Inserts the contents of `FILENAME.h` at the above point in your code, which is a simple way of including function and other definitions in your program.

Exercises

1. To compute the number of days in a month, we used the following:

```
int daysInMonth(int m, int y) {
    /* Returns the number of days in month, m (1-12)
       in year, y. */

    if ((m == 4) || (m == 6) || (m == 9) || (m == 11))
        return 30;
    else if (m == 2) {
        if (leapYear(y))
            return 29;
        else
            return 28;
    } else
        return 31;
}
```

You can further simplify the writing of the condition in the first if-statement by using the `binary_search` function in the C++ algorithm library:

```
#include <algorithm>

...
int thirtyDayMonths[] = {4, 6, 9, 11};
if (binary_search (thirtyDayMonths, thirtyDayMonths+4, m))
    return 30;
...
```

Rewrite the function to use the above condition.

2. Define a Boolean function called `valid(day, year)` so that it returns true if the `day` and `year` conform to a valid date as defined in this chapter. Use the function to rewrite the program as follows:

```
int main() {
    /* Given a day of the year (e.g. 248, 2007),
       convert it to the date (i.e. 9/5/2007) */

    //Input: day, year
    int day, year;
    cout << "Enter day year: ";
    cin >> day >> year; // read both in one statement

    // Validate input values...
    if (!valid(day, year)) {
        cout << "Please enter a valid date.\n";
        exit (EXIT_FAILURE);
    }

    // input values are safe, proceed...
    // start with month = 1, for January
    int month = 1;

    while (day > daysInMonth(month, year)) {
        day = day - daysInMonth(month, year);

        // next month
        month++;
    }

    // done, Output: month/day/year
```

```
    printf ("The date is: %1d/%1d/%4d\n", month, day, year);  
    return 0;  
}
```

Rewrite the program as shown above to use the new function. Besides developing a correct algorithm it is also important to write programs in a way that makes them more readable for you.

3. Find out the how fast your computer is by noting the clock speed of the CPU. Based on that estimate how many arithmetic operations it will be able to perform in 1 second. Write a program to see how many operations it actually performs in one second.
4. Do a web search for “Chazelle age of algorithm”. You will be rewarded with a link to an essay written by Prof. Chazelle. Read the essay and write a short commentary on it.
5. What is the fastest computer in use today? Can you find out? How fast is it compared to the computer you use? 100 times? 1000 time? 100,000 times?
6. What is/was the “Y2K” problem?



12

Fast, Cheap & Out of Control

*We suggest that within a few years it will be possible at
modest cost to invade a planet with millions of tiny robots...
With imagination and nerve we can invade the whole solar system.*

From a paper titled,
Fast, cheap and out of control: A robot invasion of the solar system,
Rodney Brooks and Anita M. Flynn,
Journal of the British Interplanetary Society, Volume 42, pp 478-485, 1989.

Opposite page: Orb Swarms

Used with permission, Copyright Phil Spitler (www.orbswarm.com)

Nearly two decades ago, Brooks and Flynn laid out a radical vision for the future that, on July 4, 1997, resulted in successful landing of *Sojourner* rover on the surface of Mars. A few years later, in 2004, two rovers *Spirit* and *Opportunity* arrived on Mars to explore the planet. In 2002, in his book, *Flesh and Machines*, a more reflective Rodney Brooks confesses to his exuberance in making the comments in his paper with Anita Flynn and describes the rovers as planetary ambassadors. The vision put forward by Brooks and Flynn was to replace a single large 1000 kilogram rover by 100 smaller 1 kilogram rovers to explore planetary surfaces. The smaller robots would move much faster and could be built much cheaper and mass produced. They will also offer much needed redundancy in case of occasional robot failure which in the case of a single larger robot would lead to total mission failure. Moreover, each individual robot would be autonomous (i.e. out of human control) operating on its own agenda as defined in its control program.

It is debatable whether Brooks actually coined the phrase, *fast, cheap and out of control*. But the credit for using it in the title of their landmark paper followed by the deployment of three rovers on Mars led to its widespread use in popular culture. The phrase became the title of a 1997 documentary movie by Errol Morris in which Brooks himself was featured. In his 2002 book, Brooks recounts how a scientist at NASA's Jet Propulsion Laboratory lambasted the idea of fast, cheap, and autonomous rovers. While the original dream of 100 small, autonomous robots remains to be realized, today one would find it hard to argue against the idea. We have used Brooks' phrase in the title of this chapter to suggest that less than twenty years from the coining of the phrase we are in the thick of exploring the full potential of small, cheap, personal robots. Given that several million personal robots have already been sold, we could conclude that our planet itself has been invaded by them. Brooks says that, "the robotics revolution is in its nascent stage, set to burst over us in the early part of the twenty-first century." In this chapter we present several examples of the ways robots are becoming a part of our everyday lives.

Robots are mechanisms guided by automated control

We accepted the above definition for robots in the beginning of this text. Using the Scribbler robot we have also learned much about its mechanisms: sensors and motors; and how to control them via C++ programs. In the course of this journey we have learned valuable lessons in building different kinds of robot brains, how to design insect-like behaviors, create sounds, images, and also ventured briefly into the realm of Artificial Intelligence. Computing lies at the heart of defining automated control. Any robot, no matter how small or large, has sensing and motor mechanisms and its behavior can be programmed to enable it to act autonomously in its environment. It is important to keep this in mind as we explore the various dimensions of the robotics spectrum.

The range of applications for robots is nearly as diverse as the world we live in. In fact, it is limited only by our own imagination. Throughout this text we have presented examples of robots in several domains: planetary rovers, vacuum cleaners, explorers in hazardous environments, toys and entertainment, education, medical surgery, manufacturing, etc. Below, we present additional interesting examples of the use of robot technology. It will be important, as we explore some of these examples, to try and use your understanding from this text to figure out what mechanisms are in use in each case. Try to think about possible other areas where similar mechanisms could be out to use.

Toys


Simple robotic toys are everywhere. Some actually look and behave like robots but most of them use computing and robotics technology in simple and innovative ways. Take for example the Tengu, designed by Crispin Jones (tengutengutengu.com). Tengu plugs into your computer's USB port. It is capable of displaying



over a dozen mouth shapes. It is designed to react to sound and music in its environment. Facial/mouth expressions change depending upon the sounds heard. If you sing a song or if there is music playing, it will appear to lip-sync to it.



Desktop robotic toys are increasing in popularity for several reasons. The enabling technology is the presence of a computer with a USB port. The USB ports are unique in that in addition to providing channels for exchange of data (like we do in communicating with the Scribbler) they can also provide power. Many toys, like Tengu, use the USB port only for the power. All the controls are built into the Tengu unit.



Some desktop toys can run independent of a computer (all the controls are present in the unit itself) and yet do not require any batteries: they run on solar power. The Flip Flap flowers and plants made by the TOMY Company (www.tomy.com) are good examples. These toys incorporate very simple mechanisms: solar sensing and power generation coupled with small motors that are activated by the electric current. These are simple, yet clever, and entertaining Braitenberg creatures. The TOMY Company has a whole line of products that employ these ideas.



There are also plenty of battery operated desktop robotic toys. The Facebank designed by Takada “eats” a coin when you flash it in front of its eyes. It runs on batteries and its mechanism includes IR sensors, similar to the ones on the Scribbler and a motor that drives or pushes the “skin” of the

face from behind.

Most electronic educational toys employ programmed automated control mechanisms. An interesting educational toy is the Tag Reading Pen made by Leapfrog (leapfrog.com/tag). A child can use the pen to point on words or text on specially made story books and the pen speaks out the word or the entire sentence. Designed for pre-school kids who are just getting interested in reading, such a toy can enhance a child's reading and pronunciation abilities. The pen has an optical reader and a speaker coupled with a memory that records a child's reading patterns. Parents can plug the pen into a computer to download and track their child's progress.



Robots need not be constructed out of digital mechanisms. One can also create control mechanisms using analog circuitry. The picture on the right shows devices called Thingamapoops (bleeplabs.com). What do they do? They can be used to produce or synthesize crazy sounding beeps. The beeps can also be sent as inputs to standard musical instruments for creating even more bizarre sound effects.

Why do you make this thing?

Because there are not nearly enough beeping, zapping, bixxerfouping, anthropomorphic synthesizer monsters in the world.


From: FAQ on bleeplabs.com

Art

Robots have been actively used in creating all kinds of art and robots themselves have been the subject of art. There are several organizations worldwide that are devoted to creating art with robots and robotic devices. In this book you have also experimented with the Scribbler drawings. A couple of nice examples of robots creating artwork are illustrated in the works of Swiss engineers Jürg Lehni and Uli Franke who created the Hektor graffiti drawing robot (hector.ch) at the School of Art in Lousanne, and Zefrank

(zefrank.com) who has created two versions of a robot he calls Scribbler which is different from the robot you have.

Both Hektor and Scribbler create new drawings based on an existing drawing. A drawing is first created. The robot (program) reads the drawing and then embellishes a new drawing based on the one input. Hektor is mounted on the wall and has a spray can that moves on a system of robot controlled pulleys. The graffiti shown here on left was painted by Hektor. You can visit Lehni and Franke's web site to view movies of the robot in action. Scribbler uses basic sketches as the basis for creating drawings. In the picture shown here, three sketches are shown that were created based on the one in the top left corner. The Scribbler concept is interesting in that anyone can go and use it via the web to create drawings in a web browser. The creators have also constructed a physical robot that makes actual drawings.



Do This: Write a program that samples or scans an image and creates a graphics drawing based on that. Read the details provided on the Scribbler web site and use it to create some sketches. Watch the process of drawing and think about how you might create something similar.

Show me the way

Global positioning systems (GPS) have been in use for many years now. More recently, small portable GPS devices have become available for use in consumer cars. A GPS enables you to enter a destination or a point of interest

and then it plots a route for you to take. It provides real-time map-based turn-by-turn guidance. The technology used in these devices uses mechanisms that involve satellite signals and street maps. These devices also form the core of autopilot mechanisms in an airplane. An airplane operating in autopilot mode can be considered a robot by our definition. Over 90% take-offs and landings of commercial flights these days are done using autopilot systems. You may have also read about robotic surveillance drones that are used by the army to spy on enemy territories. Unmanned guided missiles use similar devices. The technology also exists today to create autopilot systems in cars. That is, it is possible for your car to drive itself to wherever you'd like it to go. In some of these applications, the questions of technology deployment become more social and ethical: Do we really want these or not? What are the implications?

Affective & Social Robots

In Chapter 10 we mentioned that one of the challenges of AI research is to understand and/or artificially create human-level embodied intelligence. Many AI researchers work on advancing the state of understanding human intelligence while many others are working on building smarter, more intelligent models of behaviors. At the same time the field of robotics itself is rapidly moving in the direction of more capable, agile, and human-like robots. A good, fun way to find convergence of these advances can be seen in the goals of the RoboCup (www.robotcup.org). The RoboCup organization is focusing on soccer playing robots as a test bed for AI and robotics. They hold yearly robot soccer playing competitions that include, besides two-legged humanoids, four-legged and wheeled robot soccer players.


Besides, soccer playing robots, another area of AI and robotics research that is gathering momentum is Human-Robot Interaction (HRI). As the name suggests it is an area of research that is studying models of interaction

By the year 2050,
develop a team of
fully autonomous
humanoid robots
that can win
against the
human world
soccer champion
team.



between humans and robots. Given that we already have millions of robots amongst us, it is important to recognize the need for creating friendlier ways of interacting with robots. While we are of the opinion that every citizen of this planet should be well versed in programming and computation we also recognize that we are nowhere near that goal. As we mentioned several times earlier soon there will be more computers than people on this planet. Perhaps robots will also follow? Nevertheless the need for friendlier interactions with computers has always been recognized. With the rapid increase in the number of robot-based applications it would be even more imperative for robots, especially given their physical presence, to have socially relevant behavior traits. If not for any other reason, it would make it easier for people to accept them into our society and many levels.

Within HRI researchers are studying models of emotion, gesture recognition, and other social modalities. Emotive robots have been studied in many applications that range from toys to medical therapy (we mentioned the Paro seal robot in Chapter 1). Given the interest and recent advances in the area of HRI a new field of research has emerged: *social robotics*. Social robotics studies models that take into account social norms of behavior relevant to the environment and/or application of the robot. Emotions can play an important role in this area of research. Within the relatively small community of researchers there is much debate about whether social robotics requires one to have a physical robot. Simulated agents acting socially are acceptable to some researchers but not to others. Take for



iCat is a research platform for studying human-robot interaction topics. The robot is 38 cm tall and is equipped with 13 servos that control different parts of the face, such as the eyebrows, eyes, eyelids, mouth and head position. With this setup iCat can generate many different facial expressions - happy, surprise, angry, sad - that are needed to create social human-robot interaction dialogues.

From: research.philips.com


example the simulated robotic agent Ananova (www.ananova.com). Ananova was designed to deliver news over the web just as a news anchor might on TV. While it is not a physically embodied robot, it has the simulated morphology of a human and is capable of using the same models of emotion and expression that are designed for embodied robots. In fact, given the physical limitations of robots, the simulation is far more realistic (see picture on right).



Ananova News reader
(www.ananova.com)

We do not want to leave you with a picture of the iCat as a representative image of physical robots capable of emotional expression. In the late 1990's Cynthia Breazeal and Brian Scassellati developed a sociable robot, Kismet, to explore expressions and basic human-robot interactions. Both Breazeal and Scassellati are former students of Rodney Brooks and have been active in HRI research. Breazeal's current efforts in affective computing are based on the robot, Nexi (see picture below) which was developed at MIT in collaboration with University of Massachusetts and other corporate partners.

Brian Scassellati is developing a humanoid robot at Yale University that is roughly the size of a 1-year old child. The robot is being used to explore more fundamental developmental tasks like hand-eye coordination. It is also being used to enhance the diagnosis of autism in children. Brian and



A new experimental robot from the MIT Media Lab can slant its eyebrows in anger, or raise them in surprise, and show a wide assortment of facial expressions to communicate with people in human-centric terms. Called Nexi, it is aimed at a range of applications for personal robots and human-robot teamwork. From: David Chandler, MIT News Office, April 2008.

his colleagues postulate that the social cues that a robot needs to detect and learn are the very cues that are deficient in autistic children. They also think that such robots can be used to create functional models of autistic behavior.

Autonomous robots are already in use delivering daily mail in large offices. You can easily imagine the same technology being used for building roaming vending machines. One can already configure an office espresso maker to make coffee based on each individual's preferences: the machine senses you (or devices you wear) which transmit your preferences to the machine and it gets to work. The possibilities are endless. We conclude here with another novel application: Here is problem description (from: www.koert.com/work): *In the morning paper, I can read the weather report as well as the stock quotes. But when I look out of my window I only get a weather update and no stock exchange info. Could someone please fix this bug in my environmental system? Thanks!*

The solution: Data Fountain

Designed by Koert van Minsvoort, Charles Mignot and their colleagues, the Data Fountain uses the environment to display real time data. They designed the system to map currency rates to the height of the fountains. They call this idea information decoration. You do not have to restrict yourself to



computer screens and displays for visualizing information. The actual realization of the concept employs robotics or device control technology that obtains real time currency data from the web and then transforms it into the actuation of fountain jets.

Summary

Robots have been used to clean up hazardous substances, removal of mines, and even unclog deep sewage systems in large cities. They are increasingly finding applications in unusual situations: exploring small passageways in ancient pyramids, medical surgery, and even farming. Researchers are developing intelligent agricultural vehicles to monitor crops, plant diseases, and general growing conditions. While it may be acceptable and even desirable to have a robot that vacuums your house, or even clean up streets, one has to pause and wonder if the same technology can be used for destructive purposes. Invariably, like it is with any new technology, there are benefits as well as potential for misuse. Social and ethical concerns have to be taken into consideration to make sound decisions in all circumstances.

Robots are getting fast, small, cheap, and autonomous. However, it is only in the sense of autonomy that makes them desirable to be out of control. We may not yet have sent hundreds of small robots to another planet but we surely seem to be surrounded by millions of them here on earth. The picture of orb-like robots on a landscape comes from the Orb Swarm project (orbwarm.com). They are exploring using swarms of these robots coordinating and exploring and interacting with their environment similar to the vision of Brooks and Flynn. Whether we are successful in sending armies of orb-like robots to another planet still remains to be seen. Though one thing is for sure: it makes learning computing a whole lot more fun with a personal robot!

C++

Overview

Chapter 1

Note: All commands you type (including the Myro commands listed elsewhere) are essentially C++ commands. Later, in this section we will list those commands that are a part of the C++ language.

Chapter 2

```
#include "FILE NAME"
```

Includes the contents of the file named `FILE NAME` just as though it had been typed at this location in the current file. The file is sought in the same directory as your other C++ programs.

```
void <FUNCTION NAME> (<PARAMETERS>)  
{  
    <SOMETHING>  
    ...  
    <SOMETHING>  
}
```

Defines a new command/function named `<FUNCTION NAME>`. A function name should always begin with a letter and can be followed by any sequence of letters, numbers, or underscores (`_`), and not contain any spaces. Try to choose names that appropriately describe the command being defined.

Chapter 3

Values

Values in C++ include numbers (integers or floating point numbers) and strings. Each type of value can be used in an expression by itself or using a combination of operations defined for that type (for example, +, -, *, /, % for numbers). Strings are considered sequences of characters (or letters).

Names

A name in C++ must begin with either an alphabetic letter (a-z or A-Z) or the underscore (i.e. `_`) and can be followed by any sequence of letters, digits, or underscore letters.

```
cin >> <variable>;
```

Waits for the user to type a value compatible with the type of the variable. This value becomes the new value of the variable.

```
#include "Myro.h"  
<any other includes>  
<function definitions>  
int main() {  
    connect();  
    <do something>  
    <do something>  
    ...  
  
    disconnect();  
} // end of main program
```

This is the basic structure of a robot control program in C++. Without the `#include` line and the `connect` and `disconnect` commands, it is the basic structure of all C++ programs.

```
cout << <expression> << <expression> << ... << <expression>;
```

This function prints out the values of the expressions (one or more) in the console window. If any of the expressions are `endl`, the output will go to a new line at that point.

```
<type> <variable name> = <expression>;
```

This is how C++ defines a variable of the specified type (e.g., `double`, `int`, `string`). The value generated by `<expression>` will become the initial value of `<variable name>`.

```
<variable name> = <expression>;
```

This is how C++ assigns values to variables. The value generated by `<expression>` will become the new value of `<variable name>`.

```
<expression> <= <expression>
```

This condition is true if the first expression is less than or equal to the second. Other relations include `>=` (greater or equal), `<`, `>`, `==` (equal) and `!=` (not equal).

```
<variable>++
```

Increments `<variable>`, that is, increases the value of `<variable>` by 1.

Repetition

```
for (<initialization>; <continuation>; <incrementation>) {  
    <do something>  
    <do something>  
    ...  
}
```

```
while (timeRemaining(<seconds>)) {  
    <do something>  
    <do something>  
    ...  
}
```

```
while (true) {  
    <do something>  
    <do something>  
    ...  
}
```

These are different ways of doing repetition in C++. The first version will initialize a variable and repeat the loop body, executing the `<incrementation>` at the end of each repetition, and continue repeating so long as the continuation condition remains true. The second version will carry out the body for `<seconds>` amount of time. `timeRemaining` is a Myro function (see above). The last version specifies an un-ending repetition.

Chapter 4

`true, false`

These are Boolean or logical values in C++. C++ also defines true as 1 and false as 0 and they can be used interchangeably.

`<, <=, >, >=, ==, !=`

These are relational operations in C++. They can be used to compare values. See text for details on these operations.

`&&, ||, !`

These are logical operations. They can be used to combine any expression that yields Boolean values.

`%`

The modulus operator. The expression `<dividend> % <divisor>` returns the remainder after dividing integer `<dividend>` by integer `<divisor>`.

`rand()`

Returns a random integer between 0 and `RAND_MAX` (inclusive). This function is a part of the `cstdlib` library in C++.

RAND_MAX

A constant, which is the maximum random integer that can be returned by the `rand()` function in C++.

Chapter 5

```
if (<CONDITION>) {
    <statement-1>
    ...
    <statement-N>
```

```
}
```

If the condition evaluates to `true`, all the statements are performed.

Otherwise, all the statements are skipped.

```
return <expression>;
```

Can be used inside any function to return the result of the function.

```
<TYPE> <NAME> [] = { <VALUES> };
```

Declares a C++ array called `<NAME>` with elements of type `<TYPE>`. The array is initialized to the specified `<VALUES>`.

```
char <NAME>;
```

Declares character variable `<NAME>` (able to hold one character).

Vectors and Lists:

```
#include <vector>
#include <list>
```

Includes definitions for C++ vectors or lists, which are both examples sequence containers.

```
vector< <TYPE> > <NAME>;
```

Defines an initially empty vector called `<NAME>` with elements of type `<TYPE>`.

```
list< <TYPE> > <NAME>;
```

Defines an initially empty list called <NAME> with elements of type <TYPE>.

```
vector< <TYPE> > <NAME> (<ARRAY START>, <ARRAY END>);
```

```
list< <TYPE> > <NAME> (<ARRAY START>, <ARRAY END>);
```

Defines a vector or list called <NAME> with elements of type <TYPE> initialized with values from location <ARRAY START> up to, but not including, <ARRAY END>. These may be of the form <ARRAY NAME> + <CONSTANT>.

```
<seq>.size()
```

Returns the current size (length) of the vector, list, or string <seq>.

```
<vector>[i]
```

```
<string>[i]
```

Returns the *i*th element in the <vector> or <string>. Indexing starts from 0. (Strings are a lot like vectors of characters.)

```
<seq>.push_back(<value>);
```

Appends the <value> at the back (end) of vector, list, or string <seq>.

```
<list>.push_front(<value>);
```

Appends the <value> at the front (beginning) of <list>.

```
<list>.sort();
```

Sorts the <list> in ascending order.

```
<list>.reverse();
```

Reverses the elements in the list.

```
<seq>.begin()
```

Returns an iterator referring to the first element of the vector, list, or string

```
<seq>.
```

```
<seq>.end()
```

Returns an iterator referring the last element of the vector, list, or string

```
<seq>.
```

```
vector< <TYPE> >::const_iterator <NAME>;
vector< <TYPE> >::iterator <NAME>;
list< <TYPE> >::const_iterator <NAME>;
list< <TYPE> >::iterator <NAME>;
string::const_iterator <NAME>;
string::iterator <NAME>;
```

Declares an iterator called `<NAME>`, which can refer to locations in a vector/list with elements of type `<TYPE>` or in a string. A `const` (constant) iterator does not allow the elements of the vector or list to be modified, whereas a non-`const` iterator does.

```
<iterator>++
```

Increments `<iterator>` to refer to the next element of a vector, list, or string.

```
* <iterator>
```

Returns the value stored in the vector, list, or string element referred to by `<iterator>`.

```
<seq>.insert( <it1>, <it2-begin>, <it2-end> );
```

Inserts elements from one vector, list, or string into the vector, list, or string `<seq>`. The new elements are inserted at a location specified by iterator `<it1>`. For example, use `<seq>.begin()` to insert at the beginning of `<seq>` or use `<seq>.end()` to insert at its end. The elements to be inserted are specified by iterators `<it2-begin>` and `<it2-end>`. For example, to insert all the elements from vector, list, or string `<seq2>`, use `<seq2>.begin()` and `<seq2>.end()`.

```
for ( <it> = <seq>.begin(); <it> != <seq>.end(); <it>++) {
    ... STATEMENTS using <it> or *<it> ...
}
```

Executes the loop body with iterator `<it>` referring to the elements of vector or list `<seq>` from its beginning up to its end.

Streams:

```
#include <fstream>
```

Includes definitions for C++ streams for file input/output.

```
ifstream <name> (<string>);
```

Declares <name> to be an input file stream connected to file name <string>.

```
ofstream <name> (<string>);
```

Declares <name> to be an output file stream connected to file name <string>.

```
<ifstream> >> <var>;
```

Reads input value from stream <ifstream> and stores it in variable <var>.

```
<ofstream> << <expression>
```

Outputs values of <expression> to output stream <ofstream>.

```
<stream>.close();
```

Closes (input or output) file stream <stream>, thus completing operation on it. You should close files when you are done with them.

```
<ifstream>.get (<chvar>);
```

Puts the next character from input stream <ifstream> into character variable <chvar>.

```
<ifstream>.eof()
```

Returns true if input stream <ifstream> is at the end of the file, and false if not.

Chapter 6

The `if`-statement in C++ has the following forms:

```
if (<condition>) {  
    <this>  
}
```

```
if (<condition>) {  
    <this>  
} else {  
    <that>  
}
```

```
if (<condition-1>) {  
    <this>  
} else if (<condition-2>) {  
    <that>  
} else if (<condition-3>) {  
    <something else>  
...  
...  
} else {  
    <other>  
}
```

The conditions can be any expression that results in a `true`, `false`, `1`, or `0` value. Review Chapter 4 for details on writing conditional expressions.

Chapter 7

```
struct <structure name> { <declaration>; <declaration>; ... };
```

Defines a structure called `<structure name>` with *members* declared by the `<declarations>`s (usually variable and function declarations).

```
<structure name> <variable name>, <variable name>, ...;
```

Defines variables belonging to a structured type. Each variable has the members declared in the preceding `<structure name>` definition.

```
<structure>.<variable name>  
<structure>.<function name> ( <arguments> )
```

Accesses named member-variable (field) or member-function of a `<structure>`, which must have a structured type that declares the referenced member variable or function.

```
printf (<format string>, <expression>, ... );
```

Prints the `<expression>`s according to the `<format string>`. Within it, `%W.Df` prints a floating-point number in a field of minimum width `W` with `D` digits after the decimal point, `%Wd` prints an integer in a field of minimum width `W`, and `%Ws` prints a string in a field of minimum width `W`. The value is right-justified in the field, unless a minus sign follows the `%` (e.g., `%-W.Df`, `%-Wd`, `%-Ws`), in which case it's left-justified. There are many other format specifications besides these. In a `<format string>`, `%%` prints a single percent sign.

The math library module provides several useful mathematics functions. Some of the commonly used functions are listed below:

ceil(x) Returns the ceiling of `x` as a float, the smallest integer value greater than or equal to `x`.

floor(x) Returns the floor of x as a float, the largest integer value less than or equal to x.

exp(x) Returns e^x .

log(x[, base]) Returns the logarithm of x to the given base. If the base is not specified, return the natural logarithm of x (i.e., $\log_e x$).

log10(x) Returns the base-10 logarithm of x (i.e. $\log_{10} x$).

pow(x, y) Returns x^y .

sqrt(x) Returns the square root of x (\sqrt{x}).

Trigonometric functions

acos(x) Returns the arc cosine of x, in radians.

asin(x) Returns the arc sine of x, in radians.

atan(x) Returns the arc tangent of x, in radians.

cos(x) Returns the cosine of x radians.

sin(x) Returns the sine of x radians.

tan(x) Returns the tangent of x radians.

cosh(x) Returns the hyperbolic cosine of x.

sinh(x) Returns the hyperbolic sine of x.

tanh(x) Returns the hyperbolic tangent of x.

Chapter 8

In this chapter we presented informal *scope rules* for names in C++ programs. While these can get fairly complicated, for our purposes you need to know the distinction between a *local name* that is local within the scope of a function versus a *global name* defined outside of the function. The text ordering defines what is accessible.

Chapter 9

There were no new C++ features introduced in this chapter.

Chapter 10

`TYPE & NAME`

Indicates that the specified parameter is to be passed by reference (rather than passed by value), which allows it to be modified from within the function.

`typedef TYPE NAME`

Allows `NAME` to be used as an abbreviation for `TYPE`.

Chapter 11

`exit (EXIT_MODE) ;`

May be invoked anywhere to exit the program immediately. If `EXIT_MODE` is `EXIT_SUCCESS` then the program exits normally (equivalent to invoking `return (0)` from `main ()`); if it is `EXIT_FAILURE` then the program signals abnormal termination to the operating system.

`#include "FILENAME.h"`

Inserts the contents of `FILENAME.h` at the above point in your code, which is a simple way of including function and other definitions in your program.

Myro

Overview

Below is a chapter by chapter summary of all the Myro features introduced in this text. For a more comprehensive listing of all the Myro features you should consult the C++/Myro Reference Manual.

Chapter 1

```
#include "Myro.h"
```

This preprocessor directive inserts into your program all the robot commands available in the C++/Myro library. We will use this whenever we intend to write programs that use the robot.

```
connect ();
```

This command establishes a wireless communication connection with the robot and reports when the connection has been made.

```
disconnect ();
```

This command disconnects the wireless connection to your robot.

```
robot.beep (<TIME>, <FREQUENCY>);
```

Makes the robot beep for <TIME> seconds at frequency specified by <FREQUENCY>.

```
robot.getName ();
```

Returns the name of the robot.

```
robot.setName (<NEW NAME>);
```

Sets the name of the robot to <NEW NAME>. The new name should be enclosed in double quotes, no spaces, and not more than 16 characters long. For example: setName ("Bender").

```
robot.gamepad();
```

Enables manual control of several robot functions and can be used to move the robot around.¹

Chapter 2

```
robot.backward(SPEED);
```

Move backwards at SPEED (value in the range -1.0...1.0).

```
robot.backward(SPEED, SECONDS);
```

Move backwards at SPEED (value in the range -1.0...1.0) for a time given in SECONDS, then stop.

```
robot.forward(SPEED);
```

Move forward at SPEED (value in the range -1.0..1.0).

```
robot.forward(SPEED, TIME);
```

Move forward at SPEED (value in the range -1.0...1.0) for a time given in seconds, then stop.

```
robot.motors(LEFT, RIGHT);
```

Turn the left motor at LEFT speed and right motor at RIGHT speed (value in the range -1.0...1.0).

```
robot.move(TRANSLATE, ROTATE);
```

Move at the TRANSLATE and ROTATE speeds (value in the range -1.0...1.0).

¹ The gamepad is not supported by the current C++/Myro interface.

```
robot.rotate(SPEED);
```

Rotates at `SPEED` (value in the range -1.0...1.0). Negative values rotate right (clockwise) and positive values rotate left (counter-clockwise).

```
robot.stop();
```

Stops the robot.

```
robot.translate(SPEED);
```

Move in a straight line at `SPEED` (value in the range -1.0...1.0). Negative values specify backward movement and positive values specify forward movement.

```
robot.turnLeft(SPEED);
```

Turn left at `SPEED` (value in the range -1.0...1.0)

```
robot.turnLeft(SPEED, SECONDS);
```

Turn left at `SPEED` (value in the range -1.0..1.0) for a time given in seconds, then stops.

```
robot.turnRight(SPEED);
```

Turn right at `SPEED` (value in the range -1.0..1.0)

```
robot.turnRight(SPEED, SECONDS);
```

Turn right at `SPEED` (value in the range -1.0..1.0) for a time given in seconds, then stops.

```
wait(TIME)
```

Pause for the given amount of `TIME` seconds. `TIME` can be a decimal number.

Chapter 3

```
speak(<something>);
```

The computer converts the text in `<something>` to speech and speaks it out. `<something>` is also simultaneously printed on the screen. Speech generation is done synchronously. That is, anything following the `speak` command is done only after the entire thing is spoken.

```
speak(<something>, 0);
```

The computer converts the text in `<something>` to speech and speaks it out. `<something>` is also simultaneously printed on the screen. Speech generation is done asynchronously. That is, execution of subsequent commands can be done prior to the text being spoken.

```
timeRemaining(<seconds>)
```

This is used to specify timed repetitions in a while-loop (see below).

Chapter 4

```
randomNumber()
```

Returns a random number in the range 0.0 and 1.0. This is an alternative Myro function that works just like `(double) rand() / RAND_MAX`, which uses the `rand()` function from the C++ `cstdlib` library (see below).

```
askQuestion(MESSAGE-STRING)
```

A dialog window with `MESSAGE-STRING` is displayed with choices: "Yes" and "No". Returns "Yes" or "No" depending on what the user selects.

```
askQuestion(MESSAGE-STRING, LIST-OF-OPTIONS)2
```

A dialog window with `MESSAGE-STRING` is displayed with choices indicated in `LIST-OF-OPTIONS`, which is string of labels separated by commas. Returns option string depending on what the user selects.

```
currentTime()
```

The current time, in seconds from an arbitrary starting point in time, many years ago.

```
robot.getStall()
```

Returns `true` if the robot is stalled when trying to move, `false` otherwise.

² This description may not reflect the final C++/Myro API.

```
robot.getBattery()
```

Returns the current battery power level (in volts). It can be a double number between 0.0 and 9.0 with 0.0 indicating no power and 9.0 being the highest. There are also LED power indicators present on the robot. The robot behavior becomes erratic when batteries run low. It is then time to replace all batteries.

Chapter 5

```
robot.getBright(<POSITION>)
```

Returns the current value in the <POSITION> light sensor. <POSITION> can either be one of "left", "center", "right" or one of the numbers 0, 1, 2.

```
robot.getGamepad(<device>)
robot.getGamepadNow(<device>)
```

Returns a double vector indicating the status of the specified <device>. <device> can be "axis" or "button". The getGamepad function waits for an event before returning values. getGamepadNow immediately returns the current status of the device.

```
robot.getIR(<POSITION>)
```

Returns the current value in the <POSITION> IR sensor. <POSITION> can either be one of "left" or "right" or one of the numbers 0, 1.

```
robot.getLight(<POSITION>)
```

Returns the current value in the <POSITION> light sensor. <POSITION> can either be one of "left", "center", "right" or one of the numbers 0, 1, 2. The positions 0, 1, and 2 correspond to the left, center, and right sensors.

```
robot.getObstacle(<POSITION>)
```

Returns the current value in the <POSITION> IR sensor. <POSITION> can either be one of "left", "center", "right" or one of the numbers 0, 1, or 2.

```
savePicture(<picture>, <file>);
savePicture(<picture vector>, <file>);
```

Saves the picture in the file specified. The extension of the file should be ".gif" or ".jpg". If the first parameter is a vector of pictures, the file name

should have an extension “.gif” and an animated GIF file is created using the pictures provided.

```
robot.senses();
```

Displays Scribbler’s sensor values in a window. The display is updated every second.

```
show(<picture>);
```

Displays the picture in a window. You can click the left mouse anywhere in the window to display the (x, y) and (r, g, b) values of the point in the window’s status bar.

```
robot.takePicture()  
robot.takePicture("color")  
robot.takePicture("gray")
```

Takes a picture and returns a `picture` object. When no parameters are specified, the picture is in color.

```
cout << <vector>  
cout << <list>
```

Print vector `<vector>` or list `<list>`.

Chapters 6 & 7

No new Myro features were introduced in these chapters.

Chapter 8³

```
GraphWin()  
GraphWin(<title>, <width>, <height>)
```

Returns a `graphics window` object. It creates a graphics window with title,

³ The graphics and sound capabilities of the C++/Myro interface are under development, and so this description might not reflect the final interface. Consult the documentation for more information.

`<title>` and dimensions `<width> x <height>`. If no parameters are specified, the window created is 200x200 pixels.

```
<window>.close();
```

Closes the displayed graphics window `<window>`.

```
<window>.setBackground(<color>);
```

Sets the background color of the window to be the specified color. `<color>` can be a named color (Google: color names list), or a new color created using the `color_rgb` command (see below).

```
color_rgb(<red>, <green>, <blue>)
```

Creates a new RGB color object using the specified `<red>`, `<green>`, and `<blue>` values. The values can be in the range 0..255.

```
Point(<x>, <y>)
```

Creates a `Point` object at `(<x>, <y>)` location in the window.

```
<point>.getX()
```

```
<point>.getY()
```

Returns the x and y coordinates of the `Point` object `<point>`.

```
Line(<start point>, <end point>)
```

Creates a `Line` object starting at `<start point>` and ending at `<end point>`.

```
<line>.getP1()
```

```
<line>.getP2()
```

Returns an endpoint (`Point` object) of a `Line` object.

```
Circle(<center point>, <radius>)
```

Creates a `Circle` object centered at `<center point>` with radius `<radius>` pixels.

```
Rectangle(<point1>, <point2>)
```

Creates a `Rectangle` object with opposite corners located at `<point1>` and `<point2>`.

Oval(<point1>, <point2>)

Creates an Oval object in the bounding box defined by the corner points <point1> and <point2>.

Polygon(<point1>, <point2>, <point3>,...)

Polygon([<point1>, <point2>, ...])

Creates a Polygon object with the given points as vertices.

Text(<anchor point>, <string>)

Creates a Text object anchored (bottom-left corner of text) at <anchor point>. The text itself is defined by <string>.

Image(<centerPoint>, <file name>)

Creates an Image centered at <center point> from the image file <file name>. The image can be in GIF, JPEG, or PNG format.

All of the graphics objects respond to the following commands:

<object>.draw(<window>);

Draws the <object> in the specified graphics window <window>.

<object>.undraw();

Undraws <object>.

<line-object>.getP1()

<line-object>.getP2()

Returns the end points of the <line-object>.

<object>.getCenter()

Returns the center point of the <object>.

<object>.setOutline(<color>);

<object>.setFill(<color>);

Sets the outline and the fill color of the <object> to the specified <color>.

<object>.setWidth(<pixels>);

Sets the thickness of the outline of the <object> to <pixels>.

```
<object>.move (<dx>, <dy>);
```

Moves the object <dx>, <dy> from its current position.

The following sound-related functions were presented in this chapter.

```
<robot/computer object>.beep (<seconds>, <frequency>);
```

```
<robot/computer object>.beep (<seconds>, <f1>, <f2>);
```

Makes the robot or computer beep for <seconds> time at frequency specified.

You can either specify a single frequency <frequency> or a mix of two: <f1> and <f2>.

```
robot.playSong (<song>);
```

Plays the <song> on the robot.

```
readSong (<filename>)
```

Reads a song object from <filename>.

```
song2text (song)
```

Converts a song to text format.

```
makeSong (<text>)
```

```
text2song (<text>)
```

Converts <text> to a song format.

Chapter 9⁴

```
getBlob ()
```

Return a three-element vector of integers of the number of pixels that matched the blob, the average x and y locations of blob pixels.

```
getHeight (<picture>)
```

```
getWidth (<picture>)
```

Returns the height and width of the <picture> object (in pixels).

⁴ The graphics capabilities of the C++/Myro interface are under development, and so this description might not reflect the final interface.

`getPixel(<picture>, x, y)`

Returns the pixel object at `x`, `y` in the `<picture>`.

`getPixels(<picture>)`

Returns an iterator of type `Pixels`, which returns one `Pixel` at a time from `<picture>`.

`getRGB(pixel)`

`getRed(<pixel>)`

`getGreen(<pixel>)`

`getBlue(<pixel>)`

Returns the RGB values of the `<pixel>`. `getRGB` returns a `color_rgb` object.

`getBlob()`

Returns a three-element `vector<int>` containing the number of on-pixels and the average `x`- and `y`-locations of blob pixels. See chapter text for example.

`makeColor(<red>, <green>, <blue>)`

Creates a `color_rgb` object with the given `<red>`, `<green>`, and `<blue>` values (all of which are in the range `[0..255]`).

`makePicture(<file>)`

`makePicture(<width>, <height>)`

`makePicture(<width>, <height>, <color>)`

Creates a `picture` object either by reading a picture from a `<file>`, or of the given `<width>` and `<height>`. If `<color>` is not specified, the picture created has a white background.

`pickAColor()`

Creates an interactive dialog window to select a color visually. Returns the color object corresponding to the selected color.

`pickAFile()`

Creates an interactive dialog window that allows user to navigate to a folder and select a file to open. Note: it cannot be used to create new files.

```
repaint();  
repaint(<picture>);
```

Refreshes the displayed <picture>.

```
savePicture(<picture>, <file>);  
savePicture(<picture vector>, <gif file>);
```

Saves the <picture> **in the specified file (a GIF or JPEG as determined by the extension of the** <file>: .gif or .jpg). <picture vector> **is saved as an animated GIF file.**

```
setColor(<pixel>, <color_rgb>);  
setRed(<pixel>, <value>);  
setGreen(<pixel>, <value>);  
setBlue(<Pixel>, <value>);
```

Sets the color of <pixel> **to specified** <color_rgb> **or** <value>.

```
setPixel(<picture>, x, y, <color_rgb>);  
setPixel(<picture>, x, y, <pixel>);
```

Sets the pixel at x, y **in** <picture> **to specified** <color_rgb> **or color of** <pixel>.

```
show(<picture>);  
show(<picture>, <name>);
```

Displays the <picture> **on the screen in a window named** <name> (string).

```
takePicture()  
takePicture("gray")  
takePicture("blob")
```

Takes a picture from the Scribbler camera. It is a color picture by default, or grayscale ("gray"), or a filtered image based on the defined blob ("blob"). See chapter text for examples.

Chapters 10 - 12

No new Myro features were introduced in these chapters.

Scribblr: Myro Reference

