



Chapter 1 Introduction to Program Design

Chapter Overview

- . What is a computer program?
- . What are the parts of a program? How are they put together?
- . What kinds of questions does a program designer ask?

In this chapter you will learn how a computer can be controlled by a set of instructions called a **program**. This chapter introduces two different aspects of computation: single-minded instruction following and coordination among instruction followers. The programs in this book involve both aspects of computation.

The first aspect of computation is as step-by-step instruction following, like the process of making a single sandwich. This kind of computation is a sequence of instructions that produce some desired result. The question that drives this part is "What do I do next?" Pieces are put together using "Next,...", "If ... then ... else ...", and "until...". This kind of computation has an end goal that execution of these instructions will accomplish. The programs in this book use short sequences of instructions, executed over and over, to create entities that can provide services

©1999 Lynn Andrea Stein. This chapter is excerpted from a draft of *Interactive Programming In Java*, a forthcoming textbook from Morgan Kaufmann Publishers. It is an element of the course materials developed as a part of Lynn Andrea Stein's Rethinking CS101 Project at the MIT AI Lab and the Department of Electrical Engineering and Computer Science at the Massachusetts Institute of Technology.

Permission is granted to copy and distribute this material for educational purposes only, provided that the following credit line is included: "©1999 Lynn Andrea Stein." In addition, if multiple copies are made, notification of this use must be sent to ipij@ai.mit.edu or ipij@mkp.com.

or respond to requests (e.g., a sandwich-maker).

The second aspect of computation involves coordinating among many of these instruction-following entities. This is like gathering the sandwich-makers (and table-waiters and others) together to run a restaurant. This kind of computation is creating (and managing) a community. The driving questions are "Who are the members of the community?", "How do they interact?", and "What is each one made of?" The members of the community -- the instruction-following entities -- are glued together through their interactions and communications. Executing this kind of computation provides an ongoing program such as your car's cruise control, a web browser, or a library's card catalog.

When you finish this chapter, you will know the basic questions to ask about every computational system. These questions will allow you to begin to design a wide variety of computer programs.

1.1 Computers and Programs

Computers provide services. A suitably equipped computer can retrieve a web page, locate the book whose author you're thinking of, fly an airplane, cook dinner, or send a message to your friend half way around the world. In order for a computer to do any one of these things, two things must happen. First, the computer must be told *how* to provide the required services. Second, the computer must be asked to do so.

The how-to instructions that enable computers to provide services are called **programs**. A computer program is simply a set of instructions in a language that a computer can (be made to) follow. When the computer actually follows the program instructions, we say that it is **executing** that program. The program is like the script for a play. It contains instructions for how the play should go. But the script itself is just a piece of paper: no actors, no costumes, no set, no action. Executing a program is like performing the play. Now there is something to watch.

This analogy goes further, too. The same script can be performed multiple times, just as the same program can be executed again and again. If audience reaction (or the director's interpretation, or the theater, or the time of day) influences the performance, two performances of the same script may be quite different. Similarly, user input, hardware, software, or other environmental circumstances may make two different executions of the same program quite different from one

another. (Think of running the same word processing program on two different occasions; the experiences are extremely different even though the computer follows the same general-purpose instructions both times.)

When you sit down at a computer, someone else has already told it how to do a lot of things. For example, when you press the power switch, it boots up, or gets started running, in the way that it has been instructed to. Personal computers typically come with a fairly sophisticated set of startup instructions already installed. Simply turning on the computer causes the computer to execute this startup program.¹ Each computer has a program that it runs automatically. The program that your desktop or laptop PC runs is called its operating system. A disk drive -- which is really a separate computer plus the electronic equivalent of a huge filing cabinet -- comes equipped with instructions for how to retrieve information from (or store information in) that filing cabinet plus how to transmit that information across the cable that connects the disk drive with your "main" computer. A microwave oven comes with a computer that follows instructions for how to tell time and how to turn on its microwave generator for specified periods. The library's card catalog provides lookup services. Your car's cruise control accelerates and decelerates to keep you car moving at a steady rate. A web browser fetches and displays information it retrieves from the hard drives (file cabinets) of computers scattered around the world (at your request) (with the assistance of the "web server" programs running on those distant computers as well as the network (transmission) services provided by a set of intervening computers.

When you load a new piece of software onto your computer -- a cool new game, for example -- what you are actually doing is giving your computer a copy of the program -- the set of instructions that tells it how to do display graphics and make appropriate sound effects or whatever it is that the particular piece of software does. Writing down these instructions was the job of the person (or people) who wrote the software, the **programmer**. Loading the software makes the instructions (the script) available to your computer. Just having these instructions lying around doesn't do you much good, though. To actually play the game (perform the play), you need to do one more thing. You need to run the program.²

¹ Starting a computer is called "booting it up", presumably from the phrase "pulling yourself up by your bootstraps". The startup program that a computer executes each time that it is turned on is called the computer's "boot sequence".

² Some computer games can be run off of removable media, like CD ROMs. In this case, you don't need to load the program onto the computer, but you do need to make sure that the disk is in the drive, i.e., that the instructions are available to the computer.

Tomorrow, if you want to play the game again, you only have to run it; you don't have to start by loading it onto your computer.

1.2 Thinking like a programmer

A computer program -- "how-to" instructions for your computer -- must be written in a language that the computer can follow. There are many languages designed for instructing computers. These languages are called **programming languages**, and they are typically quite different from the kinds of languages in which people talk to one another. One of the main differences between talking to a person and programming a computer is the increased level of precision required to tell a computer how to do things. With people, it is often possible to give very vague instructions and still get the behavior you want. A computer has no common sense. You must be very specific with it. Your instructions must be step by step, in great detail. In some ways, programming a computer can be a lot like talking to a very young child or a creature from a different planet.

Imagine teaching a Martian how to make a peanut butter and jelly sandwich. You need to give detailed, step by step instructions:

1. Get a loaf of bread.
2. Remove two slices of bread and put them on the counter.
3. Get a jar of peanut butter. Put it on the counter, too.
4. Get a jar of jelly. Put it next to the peanut butter.
5. Get a knife.
6. Open the jar of peanut butter.
7. Pick up a slice of bread.
8. Using the knife, pick up a glob of peanut butter and spread it on the top of the slice of bread.
9.

These instructions tell the Martian, in very specific terms, what to do. To follow the instructions, the Martian simply needs to perform each step, one by one, in the order given. As long as each of these instructions is one that the Martian knows how to perform, when the Martian finishes executing this program, the Martian will have a peanut butter and jelly sandwich.

If there is an instruction here that the Martian does not understand, that instruction needs to be rewritten in more detail so that the Martian will be able to execute it. For example, "pick up a glob of peanut butter" might require further explanation:

1.
 - a. Insert the knife blade half-way into the jar of peanut butter.
 - b. Remove the knife from the jar of peanut butter at a slight angle so that some peanut butter is carried out of the jar by the knife.
 - c.

An instruction that needs further explanation before the Martian (or computer) can execute it is one that we call **high level**. We can use high level steps in our programs only if we can supply additional instructions to explain how to actually execute these higher level steps.

Although we don't know what instructions Martians are likely to understand, a programmer knows what kinds of instructions are a part of the particular programming language in which s/he is developing a computer program. In this book, we will use a programming language called Java. As you read this book, you will learn how to think like a programmer and how to write instructions that computers can understand. You will also learn specifically about the kinds of instructions that are part of the Java programming language.

As a programmer, you will design sequences of steps much like the peanut butter and jelly sandwich instructions. The goal of such a sequence is to get something done, to find an answer or to create something. In order to design a program like this, you will need to repeatedly answer the question, "What do I do next?" until you have reached your desired result. In many ways, this approach makes computers seem much like sophisticated calculators. In fact, this is where computers got their start: the word "computer" used to refer to people who did (mathematical) computations, and the original mechanical computers were designed to perform these computations automatically.

When you are designing a program, you should ask yourself, "What do I do next?" You don't necessarily have to write out all of the basic steps in one long sequence. You can group them together in bigger, more abstract, higher level chunks:

- I. Assemble the ingredients.
- II. Spread the peanut butter.

- III. Spread the jelly.
- IV. Put the sandwich together.
- V. Clean up.

This is a perfectly good set of instructions. But, as in the case of the Martian who didn't know how to "pick up a glob of peanut butter", these instructions will require further elaboration. A programming language such as Java allows you to make up your own high level steps, like "Assemble the ingredients", and then to explain how to do this: "1. Get a loaf of bread...." Your program is complete only when every line is either understandable by the computer or further explained in terms that are understandable by the computer. When you are done asking yourself "What do I do next?" you must then ask "How do I do each of these things?" until every line of your program is something that the computer knows how to do.

1.3 Programming Primitives, Briefly

What kinds of things to computers know how to do? Most computers don't know how to make peanut butter and jelly sandwiches. Most computers do know how to manipulate numbers and also other kinds of information, like words. In the Java programming language, you will find tools that let you send messages to other computers on a network or create windows and buttons to communicate with people using your programs. Other computers may have special kinds of instructions. A robot control system has instructions that tell the robot when, where, and how to move. A security system may have an instruction to sound an alarm. These are the basic instructions out of which programs for each of these systems can be constructed.

These basic instructions can be combined by sequencing them, as we've already seen. They can also be grouped into mini-programs and given names, like "Assemble the ingredients". These names can then be used as new instructions. When the computer needs to execute one of these new instructions, it simply looks up the rule for how to do it. (When the Martian needs to assemble the ingredients, it uses the detailed instructions that begin "1. Get a loaf of bread....")

Instructions can also be combined in other ways. Sometimes, there is a choice to be made. For example, after spreading a glob of peanut butter on top of the bread (step 8), the next step in the peanut butter and jelly program might say:

[Number is wrong; should continue list above. Same problem above and below.]

1. If the top of the slice of bread is covered in peanut butter, go to step 10. Otherwise, go back to step 8.

This step contains a choice; the next step might be 8 or it might be 10, depending on whether the slice of bread is full. The Martian (or computer) executing this program will have to keep track of which step comes next. This kind of choice step is called a conditional, and it is a common construct in programming languages. It is especially useful when the answer to the question "What do I do next?" depends on something you won't be able to figure out until you're executing the program.

We might want to go further, replacing steps 8 and 9 with a new kind of step that says

1. Repeat the following substeps until the top of the slice of bread is completely covered in peanut butter
 - a. pick up a glob of peanut butter
 - b. spread it on the top of the slice of bread.

This step ("repeat until") is called a loop. It, too, is a common construct in programming languages. Some loops tell you to keep going until something is true (like the bread becoming full), while others tell you how many times to do the steps inside the loop. Some loops even go on forever. For example, a clock is basically a loop that moves its hand(s) (or changes its display) once a minute. Loops are especially useful when part of "What do I do next?" is to repeat (almost) the same thing several times.

Each of the techniques described above -- sequencing steps, conditionals, loops, and grouping steps into new basic steps (also called procedural abstraction) -- is an important part of building computer programs. You will learn more about how to do these things in Part 2 of this book. These are the pieces that a programmer uses to answer the questions "What do I do next?" and "How do I do each of these things?" But this is only one part of the programming problem. The second part of programming is coordinating the activities of many interdependent participants in a computational community.

1.4 Ongoing Computational Activity

Some computer programs are very much like peanut butter and jelly sandwich making instructions. They start with some ingredients and step by step calculate whatever it is they're designed to create, producing an answer or result before stopping. The original mechanical computers, which mimicked human computers performing mathematical calculations, were very much like this. Sometimes, you would bring your program to a computer operator and then come back the next day for the result!

Today, most computer programs aren't like this. Instead, computer programs today are constantly **interacting**. They may interact with people, machines, other computers, or other programs on the same computer. For example, a word processing program or spreadsheet waits for you to type at it, then rearranges things on the page or recalculates values as you type. A video game moves things around on your screen, some in response to you and others by itself. A web browser responds to your requests, but also talks to computers all across the network. The cruise control system for your car responds to road conditions, sensor readings, and your input. A robot control system interacts with the robot and, through the robot, with the robot's environment, perhaps with no human input at all.

These computations aren't concerned with solving some pre-specified problem and then stopping. Most computations of interest these days are things called servers or agents or even just applications. Most of them have some basic control loop that responds to requests or other incoming information continually. These computations are embedded in an environment and they interact with that environment: users, networks or other communication devices, physical devices (like the car), and other software that runs at the same time.

These programs are not just interacting with the things around them, either. In fact, each of these programs may itself be composed of many separate pieces that interact with each other (as well as with the world outside the program). Coordinating the activity among the many entities that make up your program -- and their interactions with the world around them -- is the second aspect of computer programming.

This is kind of like taking a group of Martians and organizing them to run a restaurant. Some of the Martians will take orders from and serve food to the customers. Other Martians will need to cook food for the customers. Still others will need to check on supplies, make change, or coordinate other aspects of the restaurant's operation. Each of these Martians will provide services to and make

request of other Martians (or to the restaurant's customers or suppliers or other parts of the environment in which the restaurant is embedded). Coordinating the interactions among these Martians (and between the Martian restaurant and its environment) involves different kinds of questions from the instruction-following "What do I do next?"

Before we turn to the coordination of activity, though, let's look closely for a moment at one of the Martians who will staff our restaurant. We will see that, deep down, peanut butter and jelly programming still has an important role to play in creating computational activity. Keep in mind that this Martian represents just one of the many things going on in our restaurant.

The instructions that a Martian chef follows might look very much like this:

1. Pick up a new food order.
2. Find the instructions for the dish ordered and follow them.
3. Put the completed dish and the order information on the counter for pickup.
4. Go back to step 1.

Step 2 of this program is the kind of "higher level" step that we described above. It is not itself complete; instead, it refers to other, more detailed instructions to be followed. For example, if an order comes in for a peanut butter and jelly sandwich, the Martian chef will need to use the instructions developed above for how to make a peanut butter and jelly sandwich. A computer still follows simple sequenced steps written in a language that it can execute. But while this Martian is making a peanut butter and jelly sandwich, another Martian is asking the customer at table 3 whether she would like some more water. Later, the Martian waiter will come into the kitchen and pick up the sandwich that the Martian chef just made. And when the Martian chef is done making the peanut butter and jelly sandwich, the Martian will turn to the next food order, continuing its ongoing interaction.

The peanut butter and jelly style of program instructions is an important part of how the Martian chef does its job. But the Martian chef's instructions are not simply the steps of the peanut butter and jelly program. The basic structure of the Martian chef program is an **infinite loop** -- a loop that goes on forever. This program accepts requests (in the form of new food orders) and provides services (in the form of the completed dishes) over and over again. We sometimes call this kind of loop -- one that provides the main behavior for a participant in the interactive program community -- its **control loop**. Many program community

participants take this form, and we will look more closely at control loops in Part 3 of this book. Programs with ongoing central control loops like this are the members of our interactive computational community.

1.5 Coordinating a Computational Community

At its most basic level, every computer program is made of instructions that are followed, one by one. But a single computer program may have many instruction-followers inside it, just as our restaurant is run by many individual Martians. When you look at the whole program -- like the whole restaurant -- you don't necessarily see the individual instruction steps. Instead, you see coordinated activity among a group of interacting entities. The behavior of this community -- providing customers with hot meals -- is not the responsibility of any particular member of the community. Instead, it is the result of many community members working together in a coordinated fashion.

Building modern interactive software involves something very much like organizational design. We call this part of programming "constituting a community of interacting entities". The programmer's job to figure out how to tell the computer what to do, and no matter what the specific problem to be solved may be, there are fundamental questions that each programmer must ask. Designing a computation which is a community of interacting entities involves figuring out who the members of this community are, how each one works, and how they interact. This is like setting the cast of a play, or deciding what the sub-units of your business will be, as well as how they should interrelate. In planning the organizational structure of your business (or program), you also have to figure out how each unit works and what -- and how -- they are supposed to communicate. These are the big questions of this second aspect of programming.

When you are designing this kind of activity, you ask yourself several questions:

- . What is the desired behavior of the program?
- . Who are the entities who interact to produce this behavior?
- . How does each one work?
- . How do these entities interact?

In the remainder of this section, we will expand these questions and begin to explore them in somewhat greater detail. Understanding these questions and their

ramifications is the theme of this entire book. Coordinating communities is a special focus of Part 4.

1.5.1 What is the desired behavior of the program?

Before you can design a system to solve your problem, you must know what your problem is. This involves knowing not only what you want, but how it should work or fail to work under a variety of different circumstances.

Some questions that you ought to be able to answer about your desired program include:

- . What services should your program provide?
- . What guarantees does your program make about these services?
- . Under what assumptions (circumstances, conditions) does your program make these guarantees?

Consider the restaurant of the previous section. What can we say about its behavior? In answering this question, we consider both the experiences of individual customers and the ongoing properties that the restaurant must maintain, such as remaining solvent. A basic specification of the service provided by the restaurant might be: Each customer is seated at a clean table, the order is taken, food is served, a bill presented, and payment collected.

There are a number of guarantees we want to make about these services. For example, customers should not have to wait for an unduly long time. Different parts of the restaurant must communicate; customers should not be charged for food that they were not served, etc. Over time, the restaurant should take in at least enough revenue to cover its operating expense. Supplies should not run out, nor should they rot.

We will make certain assumptions in order to be able to provide these guarantees. For example, the "timely service" guarantee will only be possible if the load on the restaurant is reasonable. We might decide that we will only be able to uphold this guarantee if the number of people wanting to eat in the restaurant at one time never exceeds its capacity, and if the rate of arrival of these people doesn't exceed the rate at which the restaurant can serve them.³ These assumptions should be

³ How many customers the restaurant can handle is called its bandwidth. How quickly each one can be served is called its latency. The number of customers per hour that the restaurant can handle is its throughput. These quantities -- bandwidth, latency, throughput -- are common

made explicit, and we will also need to say what happens when they are violated. (In this case, the timely service guarantee won't be upheld, but how slow the service gets should be related to how overloaded the restaurant is.)

There are other assumptions we do not make about our program, and we can articulate these as well.. We do not assume that only one customer will be served at a time. Instead, we expect that multiple tables must be handled (roughly) simultaneously. It certainly won't do to wait until the first has eaten, paid, and left before addressing the second. We also permit different interactions with each table to be handled simultaneously or at least overlapped; food may be cooking while checks are being written up.

This description is still fairly general, and we can imagine making it more specific. (For example, are customers constrained to ordering off of a menu?) In general, the more detail you can give of what your program ought to do, the easier your task will be in designing and building it.

1.5.2 Who are the members of the community?

This question can't be answered in isolation, because any and every decision you make about *who* the entities are is also at least a partial commitment to *what* they are and *how* they work. So answering this question is in many ways like solving the whole problem. The trick is to answer this question in fairly high-level, general terms, then to sit down and try to hash out the answers to all of the what and how questions. In answering those, you'll almost certainly have to return to this question and rearrange your answer a few times. This is fine; it's even typical enough to have a name: **incremental program design**.

In the restaurant, an appropriate high level division of labor might have a wait staff unit (the people who deal directly with the customers), a kitchen staff unit (the people who cook the food), and a financial unit (who keep track of how much which things cost, collect money, and buy supplies). At this point, we haven't committed to whether these are three roles played by a single Martian, three separate Martians, or even three groups of several Martians each.

1.5.3 What goes inside each one?

To answer this question requires knowing a bit about how each entity will interact with the other members of its community. This means that answering "what goes

measures of program performance.

inside" is closely related to "how do they interact." After all, specifying what interactions each entity needs to support goes a far way towards telling you whether the "what goes inside" meets the requirements of the community.

Some subsidiary questions to ask about how each of the entities is constituted include:

- . What responsibilities does it have?
- . What guarantees (promises, commitments) does it make? Under what assumptions?
- . What resources does it control?
- . How does it work?
- . Is it a community, too?

For example, the restaurant's wait staff might be responsible for greeting the customers in a timely fashion, supplying each one with a menu (a structure that the program will have to provide and keep updated!), taking the order, delivering it to the kitchen staff, picking up and serving the cooked meal, obtaining a price from the accounting entity, and obtaining payment for that amount from the customer. The wait staff might guarantee to communicate with (most of) the customers within minutes, provided the total number of customers is limited and the maximum time spent with each is under a certain amount. It might also promise to deliver food within some small amount of time after it's done cooking, provided that the kitchen staff notifies the wait staff in a timely manner. The wait staff controls menus, knows which food items were ordered by which customers, and is the only part of the restaurant that deals directly with the customers. And so on.

When it comes to "How does it work?", there are two kinds of answers. One answer is that the behavior of the entity is accomplished by a single rule-follower running an interactive control loop. We saw an example of this when we considered the Martian chef earlier. In this case, we ask "What does the Martian do next?" over and over, until we wind up with a well-defined set of instructions for this Martian to follow.

The other possible answer to the question "How does this entity work?" is that this entity is itself a community. (The wait staff might be further divided into the person who takes the order, the person who clears the table, and the person who serves the wine.) In this case, we need to figure out how to build each of these entities, asking again "What goes inside each one?" The problem of figuring out

how to coordinate the activity of a community continues until each community member is a single (rule-follower) Martian. Then we ask about the instructions this Martian follows.

1.5.4 How do they interact?

This question concerns coordination and communication among two or more entities. Some of the questions that you should ask about how these entities interact include:

- What are the entities' interfaces?
 - What promises does each one make?
 - What contracts does it fulfill?
 - What services does it provide?
- How do they communicate?
 - What mechanisms do they use?
 - What interaction patterns do they use?
 - How do they preserve liveness, i.e., make sure that things keep moving?
- What interaction patterns are possible?
- What happens when something goes wrong?

A **protocol** is the specification for an interaction between two entities. For example, a common protocol for the interaction between the wait staff and kitchen staff of a restaurant involves a slip of paper with the customer's order written on it. The waiter hangs this piece of paper in the window over the kitchen's food pickup counter, a place where it will be easy to find when someone from the kitchen is ready for a new job. When a member of the kitchen staff is ready to process the order, the piece of paper is removed and used to guide the food preparation. When the order is ready, it is placed on the food pickup counter together with the original order slip. This identifies the food with the original request when the waiter returns to retrieve it. The slip of paper serves as a crucial reminder of several associated pieces of information: what was ordered, by whom, and where they are seated.

Protocols can also address temporal issues. For example, the wait staff/kitchen

staff interaction described in the preceding paragraph needs to happen in real time, meaning that the protocol itself can't introduce significant delays. There must also be guarantees made about the frequency with which the wait staff checks for completed dishes (or the kitchen staff for incoming orders). If assumptions such as these are built into protocols, they must be documented so that they are maintained in the behavior of participant entities.

In contrast, the wait staff interacts with the financial unit by obtaining prices for food and turning over any moneys collected. These interactions could happen in batch, meaning that it is OK for the wait staff to get the price list at the beginning of the week or for money to be handed over at the end of the day.⁴ The difference between real time and batch interactions is only one dimension that must be determined in order to coordinate the activities of the members of your computational community.

A protocol specifies the interface, or meeting, between various entities in the community that constitutes your program. Once the interfaces have been thoroughly fleshed out, each entity can in theory be implemented by a separate programmer (or team of programmers) provided that it is built to spec, i.e., that it meets the specifications of the agreed-upon interface.

In practice, the task of implementing an entity to match a given specification often results in questions about or revision of that interface. Programming is not so neat a task as students of computer science would often like to believe; there's a cycle of specification and implementation, debugging and testing, usage and revision, that characterizes almost all real-world software. The later stages of this process are sometimes called the software life cycle; but the repeated revision that characterizes those later stages start before a piece of software is even born.

1.6 The Development Cycle

The sections above concern the design of a computer program. Typically, you will be given a set of specifications and some components that need to be integrated into the system you build. Perhaps you will only be asked to build a single entity or to modify existing entities to facilitate coordination. Regardless of your particular design problem, you will find it useful to situate your task in the context

⁴ Batch processing is like the old-fashioned computations in which you handed your program to a computer operator and came back the next day for your results.

of these six questions:

- . What is the behavior of this program?
- . Who are the entities that combine to produce this behavior?
- . How do they interact?
- . What is each one made of? (A community of entities or a single instruction-following control loop?)

And, when we get down to instruction-followers,

- . What does it do next?
- . How does it do each one of these things?

Once you have the answers to all of these questions, you can start to build your program. Of course, you will already have found that you needed to go back to earlier parts of the design process to modify or flesh out various decisions. You may also have shown your completed design to other programmers -- or, perhaps more importantly, to the users or customers for whom you are creating this service -- and revised your design specification in response to their feedback.

The implementation phase of the project is no different. In building a program that is supposed to meet your specification, you will often find that you need to go back and change that specification. When this happens, you need to be careful to consider all of the interdependencies that led you to your original design. That is, the development of software is cyclic, beginning with design but often returning to it. It will not always be desirable (or even possible) to change your design, but it is quite common to discover additional assumptions or nuances that must be percolated through the design during later phases of development.

When you begin to build your program, it is often advisable to implement only a small piece of your system first. This may mean implementing only some of the entities, or it may mean implementing all of the entities but only simple, basic versions of each. In large scale system development, this initial phase is called prototyping. Even in most of the smaller scale programs that you will encounter in your early coursework, it is a good idea to utilize this approach of incremental program development. Part of developing good programming skills involves learning to consciously and explicitly design a staged development plan in which smaller simpler programs are constructed and debugged, then gradually expanded until the desired functionality is obtained.

Building a simpler version of your system gives you an opportunity to test your basic approach before you have built up too much complexity. It also means that your **bugs**, or program errors, will be easier to find. Bugs come in many flavors, ranging from simple syntactic errors such as spelling mistakes, to programming errors such as incorrect variable scoping, to conceptual design problems such as impossible-to-meet but critical guarantees.

Even after you've found the bugs that keep your program from running, you will need to subject your code to rigorous testing. This means trying out not only the "normal" expected behavior, but also checking how your program handles unexpected or anomalous behavior. Think of your program as an opponent you're trying to trick; see if you can get it to misbehave. This testing -- when done right -- will lead you to modify your code or even your design.

This repeated cycling through and between the various stages of specification (or design) development, implementation, and testing is a crucial skill for any good programmer. Classroom programs are too often written once and tested on obvious cases. Most of the time and money spent on real-world software is spent on revision and maintenance rather than on initial development. Acquainting yourself with this cycle -- and with writing clean, easy-to-read, reusable code -- may be the most important part of becoming a skilled programmer. These issues -- together with a tour through the development cycle -- are the topic of the next chapter.

1.7 The Interactive Control Loop

This book focuses on the problem of designing interactive software. At the heart of our approach is the idea of an **interactive control loop**. This is a simple program that repeatedly receives an input -- a new request, a set of sensor readings, or some other information -- and responds appropriately. In the general case, the response may involve initiating a series of other activities, so this kind of program can in principle become almost arbitrarily complex. The basic idea is rather simple, though.

To conclude this chapter, we present an extremely simple interactive control loop. This example will be used as a motivator for the development of the next part of the book. The interactive control loop idea is a theme that runs through this entire book. In a way, it might be thought of as the "atomic unit" or basic vocabulary element of this kind of computation.

Perhaps the simplest interactive control loop is an **echo** program. When run, this program waits for the user to type something. When the user finishes typing, the program simply repeats back what it has been given. That is, it's a loop that gets some input, processes that input (in this case trivially), and then spits out its result.

Although the echo program seems too trivial to be of much use, a minor variant of it runs in almost every program you type to: it's what makes the characters appear on the screen. Far more importantly, the basic structure of this program underlies essentially every interactive computation. And it demonstrates many of the important properties of an interactive computation:

- . It is embedded in an environment (in this case involving a user's typing and a display that the user can see).
- . It is interactive (with that user, but we could have it talk to another program or over a network instead).
- . It is concurrent: other things happen at the same time that the program is running. (In this case, the user might be typing the next line even while the echo program is producing its output.)

The idea of an interactive control loop is the root of this approach to programming. By putting together interactive control loops, you constitute a community of interacting entities. Interactive control loops are what goes inside; communication between them is how they interact. In other words, as they say, all the rest is corollary....

Chapter Summary

- Computers follow special instructions, called a program, which is written in a special programming language.
- Computation results when a computer has access to these instructions and executes them.
- Each set of instructions must answer:
 - What should the program do next?

- How should it do it?
 - Groups of steps can be combined to make a "higher order" step.
 - Steps can involve choices or decisions.
 - Steps can be executed over and over again using a loop.
 - Most modern programs combine many separate looping instruction-followers into an interacting community.
 - Every computation is embedded in an environment and interacts with the other (computational and physical) entities around it.
 - The programmer's job is to figure out:
 - What services (behavior) does my program provide?
 - Who are the entities that together provide this behavior?
 - How does each one work?
 - How do they interact?
 - Program construction is a cycle of designing, building, testing, and then designing again.
-

Exercises

1. Give step by step instructions for how to tie shoelaces.
2. Select your favorite recipe and give step by step instructions for how to cook it.
3. Give detailed directions for how to get from your classroom to where you live. Include indications that will tell whether you've gone too far and how to get back on track.
4. Specify the expected behavior for each of the following interrelated services provided by a bank account:
 1. A deposit.
 2. A withdrawal request.

3. Checking your balance.

Does your specification permit overdrafts?

5. You are at a fruit market. Describe the protocol by which you purchase a piece of fruit from the fruit seller.

6. Describe the division of responsibility and coordination of activities among the players on a soccer team.