Methods

We now have a number of tools for processing data and making complex decisions. Our programs can now get pretty complicated, which requires a lot of thinking to keep track of how they work. As good computer scientists, we should be starting to feel worried that we need to be lazier, we need tools to help us organize our code into small pieces so that we never have to keep the whole program in mind at once.

Also, we often write a multi-step task that we then need to apply to various different data values. We could do this by copy-pasting the code and putting in different specific values, but any time we have multiple copies of the same code, that means we have to be responsible for keeping track of those copies and updating all of them any time we need to adjust how they work… which we should be too lazy to do.

The solution to both of these problems is to divide our programs up into methods. Methods are also known as modules, functions, procedures, subroutines, etc.

Method Header

Each method begins with a header that gives it a name, says what values it needs as input from the rest of the program, and says what type of result it will give back to the rest of the program.

```
returnType methodName(type1 param1, type2 param2, …)
```

The method name follows the usual rules for naming, and should indicate what the code inside the method body will do.

The parameters are a list of zero or more local variables that the method will use to perform this code. Each parameter will start with a value that it is given by another part of the program.

The return type is the type of a value that the method will give back to another part of the program when it is finished. This can be any valid value for a variable. The return type could also be the special type *void* which means that it does not return any value at all – it just does a series of steps but does not have a single result value.

To start off with, we'll create very simple methods that have zero parameters and type void, so they get no values from the rest of the program, and give back no result.

Method Body

Like other structures, a method can have code inside it, which is called its body. We can use curly braces {} to begin and end the method's body, or else use `end` and then the method's name.

```
// curly brace format
void printHello5() {
   //body of the method
   for (int i = 0; i < 5; i++)
         print "Hello"
   endfor
} // end of the method

//end format
void countDown()
   //body of the method
   for (int i = 10; i >=0; i--)
         print i
   endfor
         print "YAY!"
end countdown
```

The body of a method can include any code we would put anywhere else including conditionals, loops, printing, etc, but in most languages you cannot create a method inside another method.

Local Variables

Variables created inside a method are local to that method – they only exist within that method ("in that method's scope"), while it is running. They do not persist afterwards. This means that multiple methods could have variables with the same name without interfering with each other. Consider readability when choosing local variable names; if two methods are doing related things to the same kind of value, it may make more sense for them to use the same variable name. If two methods are doing very different things, it may cut down on confusion if they use different variable names as well.

Calling Methods

The point of a method is that it is a tool that other parts of the code can use. Methods are used by *calling*, that is, using their name. The name of a method always includes the parentheses after it, with values for any parameters inside.

If a program has multiple methods, it should have one method called main(), a.k.a. the main method. The program automatically begins with the first line of main, and if main calls another method, that method is run at that point in main. If main does not call a method, it does not run, even if it is present in the code.

```
void square()
   int userNum = prompt "your number?"
   print "square is " + userNum * userNum
end square

void cube()
   int userNum = prompt "your number?"
   print "cube is " + userNum * userNum * userNum
end cube

void main()
   // decide whether to call the cube or square method
   String user = prompt “1 for cube, 2 for square”
   if (user == 1) then
         cube()
   else
         square()
   endif
   print “Goodbye”
end main
```

In this code, the first thing that happens is that the user is prompted to choose between cube or square. If the user chooses 1, the cube() method runs. Once the cube method finishes, we are back in main, and we print “Goodbye” and in that case, the square() method never runs. if the user chooses anything else, the square method runs, then we print “Goodbye” but the cube() method never runs.

In this example, I put square() and cube() above main(). In some languages, the order of methods is important because you cannot call a method until after you have defined what the method does. In other languages, the order of whole methods does not matter; however, the code *inside* each method is still run from top to bottom. In this example I put main at the bottom, but remember that if we have a main, the program will start with the first line inside the main, no matter where main is located in relation to other methods.

The Stack

When a method is called, a chunk of main memory space for the program is taken up by the method, including space for any local variables it has. The part of program memory used for methods is called the *stack* and each method's chunk of space is called a *stack frame*.

When the method ends, the stack frame is "given back" that is, that chunk of main memory space is marked as no longer occupied, it is open for use to store a new stack frame for another method. If one method calls another which calls another, then we could end up with multiple stackframes which take a while before they come back off the stack.

Parameters

To make methods really useful, we need a way to get input from where the method is called into the method, so we can do the same code multiple times on different starting values. We do this by passing values in to the method's parameters.

Parameter Lists

In its parentheses, each method is defined with a list of zero or more parameters. These are local variables in the method, but unlike variables created in the body of the method, they are automatically given values before the method starts. When writing the method, you don't have to give them values, but you do have to think about how to write the method so that it will work properly no matter what value they start with.

If you need the variable to start with a specific value, then it should not be a parameter, it should be a local variable defined in the body of the method. If you need input from the user, the variable for that should not be a parameter, it should be a local variable defined in the body of the method. The point of a parameter is that it is a way for another part of the program to feed a value to the method.

Calling Methods with Parameters

Parameters to a method are listed in a specific order. When the method is called, it must be given values in the same order, in its parentheses. In this example, two methods each have two parameters.

```
void printDo(boolean printIt, string word)
   if(printIt) then
         print word
   endIf
end printDo

void printTimes(string word, int howMany)
   for(int i = 0; i < howMany; i++)
         print word
   endFor
end printTimes

void main()
   printTimes("hello", 3)
   printDo(true, "hi")
   int times = 2
   printTimes("goodbye", times)
   printDo(true, "bye")
   string str = "this is awkward"
   printTimes(str, times - 3)
   printDo(false, "huh")
   print "END"
end main
```

output:

```
hello
hello
hello
hi
goodbye
goodbye
bye
END
```

Main calls each method three times with different values for the parameters, called *arguments*. (I will only talk about the stack for the first call, all the others are similar).

- printTimes("hello", 3)
    - just before printTimes begins, a stack frame is placed on the stack and in it are space for the parameter variable word, which is initialized to the value "hello" and the parameter variable howMany, which is initialized to the value 3
    - the method runs and "hello" is printed three times
    - the method ends, and we give back the stack frame, so this space in main memory can be re-used for other methods. Now we are back in main()
- printDo(true, "hi")

    - just before printDo begins, the parameter variable printIt is initialized to the value true and the parameter variable howMany is initialized to the value “hi”
    - the method runs, printing “hi”
    - the method ends and we are back in main()
- printTimes(“goodbye”, times)
    - just before printTimes begins, the parameter variable word is initialized to the value “goodbye” and the parameter variable howMany is initialized to the value 2 (the values of the times variable)
    - the method runs and “goodbye” is printed two times
    - the method ends, and we are back in main()
- printDo(true, “bye”)
    - just before printDo begins, the parameter variable printIt is initialized to the value true and the parameter variable howMany is initialized to the value “bye”
    - the method runs, printing “bye”
    - the method ends and we are back in main()
- printTimes(str, times - 3)
    - just before printTimes begins, the parameter variable word is initialized to the value “this is awkward” and the parameter variable howMany is initialized to the value -1 (the result of subtracting 3 from the value of the times variable)
    - the method runs and since the for loop starts at 0 and goes only as long as it is < howMany, which is -1, nothing is printed
    - the method ends, and we are back in main()
- printDo(false, “huh”)
    - just before printDo begins, the parameter variable printIt is initialized to the value false and the parameter variable howMany is initialized to the value “huh”
    - the method runs, but since it checks printIt, which is false, it doesn’t print anything
    - the method ends and we are back in main()

Notice that we were able to pass in literal values, values stored in variables, and values from expressions as our arguments.

In general, the reason we are passing values into a method is so that the method can use those values, so in most cases we wouldn’t be giving a parameter variable a new value, but we might adjust the value of a parameter in the course of the task the method does, so it is worth considering what happens when we pass a variable in as argument to a parameter that is changed by its method.

```
void countDownFrom(int top)
   // since top is already initialized
   // there is no initialization step
   // in this for loop
   for(     ; top > 0; top = top -1)
         print top
   end for
   print "YAY"
end countDownFrom

void main()
   countDownFrom(10)
   int k = 5
   countDownFrom(k)
   print "k is now " + k
end main
```

In this code, the countDownFrom method uses the top parameter as its counter, counting down from whatever value top was initialized to until top is down to 0, decreasing each time.

When we call countDownFrom(10), the top parameter variable is initialized to 10, and then it counts down from 10, changing the top variable each time through the loop. There was no variable in main, so there is nothing to change.

When we call countDownFrom(k), if we are doing the top parameter variable is initialized to 5, which is the value of k from main. The code counts down from 5, each time through the loop changing the value of top. But does this also change the value of k?

Approaches to Arguments

The most common approach is *pass by value*: a parameter is initialized with a copy of the value that is given when the method is called. In this case, any change made inside the method would not affect any variable that was passed in for that parameter.

In this case, no change made to a parameter's value inside the called method can change the value of a variable in the calling method. In our example above, top would be initialized to 5, a copy of the value of k, but the changes made to k after that would not affect k, and k would still be 5 after the countDownFrom method finished for the second time.

Another approach is *pass by reference*: a parameter is assigned the same memory address as a variable passed in as argument. So the parameter variable becomes another name for

the same storage location in main memory – two names for the same thing – and either variable name can be used to access that spot in main memory. So any changes made to the parameter variable are immediately made to the value stored by the original variable. When we learn about reference variables (a.k.a. pointers) more generally we will have a better sense for how this works.

For the moment, we will assume that arrays and only arrays are passed by reference. Remember that an array variable is actually the first address in a block of addresses in main memory and the index is an offset saying how far down in that block to go. So if we pass an array into a method, the method will have access to that same block of main memory. This means that changes made to the elements of an array during a method *will* persist after the method ends, because it is the same array, not just a copy.

```
void addList(int list[], int add)
   for (int i = 0; i < list.size; i = i + 1)
         list[i] = list[i] + add
   end for
end addList

void main()
   int numList[5] = {2, 4, 6, 8, 10}
   addList(numList, 1)
   print "0th element is now " + numlist[0]
end main
```

Since numList is an array, when it is passed into addList, the parameter variable list isn't just a copy of the array, it is the same array. So, when the addList method adds a value to each element of the list array, numList's values are changed (because they are the same) and when we get back to main, we will see that numList's 0^{th} value is now 3, not 2.

There are also less common approaches to passing variables that are used in some languages, but pass by value and pass by reference are how most languages handle parameters.

Returning from Methods

Just as we use parameters to get values from the caller into a method, a method can give a value back to a caller by returning it.

Return Type

Any type that is valid for a variable is also valid for a method's return type, including arrays. We also have the special type void to indicate that a method does not return anything. If a method specifies a return type, the method must return a value of that type by the end of the method. A method can return any value of its type, including values like zero or null. Only a single value can be returned, although we could return an array that contains several individual values that are of the same type.

Return

The keyword return is followed by the value (of the return type) that this method will evaluate to. It could return any expression that evaluates to the correct type, including a literal value, a variable, or the result of doing operations

```
double sumDouble(double a, double b)
   double sd = (a + b) * 2
   // return statement
   return sd
end sumDouble
```

In the example above, we did the math in the expression and put the result into a variable, then returned the variable. It is somewhat more common, in a case like this, just to put the expression after the keyword return:

```
double sumDouble(double a, double b)
   return (a + b) * 2
end sumDouble
```

A return statement ends the method, so a method can't have any statements after a return. For this reason, we should never have a return inside a loop unless it is also in a conditional that means it only happens under certain circumstances (which mean the method should end at that point).

A method could have more statements below a return statement if the return statement were in a structure like a conditional.

```
int findPosition(String [] namelist, String find)
   for(int i = 0; i < namelist.size; i = i + 1)
         // if we find the name at the position
         // return the position
         if (namelist[i] == find) then
```

```
            return i
        endif
    endfor
    // if we got all the way to the end without finding it
    // use -1 to indicate not found
    return -1
end findPosition

string pickName()
    string user = prompt "Enter your name:"
    if (user != "") then
        return user
    else
        return "No Name Given"
    endif
end pickName
```

The code above has two returns in each method, because we have different approaches to determining the return value depending on the situation. Many people feel that code is easier to read and debug if there is only one return statement per method, so they would use a variable to store the eventual value to return, change that variable as needed with conditionals, and then return that variable at the end:

```
int findPosition(String [] namelist, String find)
    int position = -1 // we will return this
    for(int i = 0; i < namelist.size && position==-1; i = i + 1)
        // if we find the name at the position
        // return the position
        if (namelist[i] == find) then
            position = i
        endif
    endfor
    return position
end findPosition

string pickName()
    // user input but also return variable
    string user = prompt "Enter your name:"
```

```
    if (user == "")
          user = "No Name Given"
    endif
    return user
end pickName
```

Even if we write a method using multiple return statements, only one of them can happen.

Calling Methods that Return

When we call a method that returns, that method call evaluates to the returned value – imagine the method call turning into the return value. Most of the time, if we called a method that returns, it is because we want to do something with that value. So although it is usually legal to just call a method returns as a statement of its own, this is usually a bad idea. We should be storing the return value in a variable or else using it in another expression, maybe even passing it to another method.

If we fail to do this, we sometimes say that we "dropped the return value on the floor" imagining that the method tossed its return value back to our code and we failed to catch it.

```
double sumDouble(double a, double b)
    return (a + b) * 2
end sumDouble

void main()
    // useless, we drop 14 on the floor
    sumDouble(5.1, 1.9)

    // catching 14 in a variable
    double result = sumDouble(5.1, 1.9)

    // printing 14
    print "your result is " + sumDouble(5.1, 1.9)

    // using return values as arguments
    double bigMath =
            sumDouble(sumDouble(2.2,3.3), sumDouble(1.1,5.5))
          // evaluated as:
          // sumDouble(11.0, sumDouble(1.1, 5.5))
          // sumDouble(11.0, 13.2)
```

```
        // 48.8
   print "result of big math was " + bigMath
end main
```

So at the end, we print "result of big math was 48.8"

Recursion

Suppose we have two methods, methodA and methodB, and into methodA we put a call to methodB, but then into methodB we put a call to methodA. Then methodA would call methodB which would call methodA which would call methodB which would call method…

This is called *mutual recursion*.

If into methodA we put a call to methodA itself, then this is just *recursion*.

Certain problems have recursive definitions, in which case it makes sense to write recursive methods to solve them. For example, the nth Fibonacci number is defined as the sum of the n-1st and n-2nd Fibonacci numbers. To write this as code:

```
int fibonacci(int n)
       return fibonacci(n-1) + fibonacci (n-2)
end fibonacci
```

There is a problem here, however. This would just keep running forever; we never told it where to stop. Actually, the Fibonacci numbers are only defined this way as far back as the 2nd Fibonacci number. The 0th and 1st Fibonacci numbers are just defined as 1.

For recursion to be useful, like a loop, it has to have an ending point. This is called the *base case*. So a recursive method should start with a conditional that checks whether we are in a base case, in which case it does whatever the base case does.

Otherwise, we can do the *recursive call* which means calling the method itself, but on a new argument that is closer to the base case than the argument we were given.

```
int fibonacci(int n)
   // check for base case
   if (n == 0 OR n == 1) then
        return 1
   endif
   // recursive call (n-1 and n-2 are closer to the base case)
  return fibonacci(n-1) + fibonacci (n-2)
end fibonacci
```

Remember that every time we call a method, that method gets space on the stack. Each time we do this takes up time and space. We end up getting a *stack overflow* from putting too many method calls on the stack; and even if not, it may take a very long time to work through all those method calls, even if each one is very simple.

So, in most languages it is much faster and more space efficient to use loops instead of recursion to repeat code, except when we need the form of the code to reflect the inherent recursion in our understanding of a task.

Overloading

Choosing good names for methods is an important programming skill; we want the name to clearly communicate what the method does so that when we read a single statement that calls the method, we have good expectations for that many lines of code inside the method that this may stand for.

Sometimes we need to write several methods that do the same task, but the parameters are different in number or type. For readability, it may actually be better to give all these methods the same name. When we have multiple names with the same name but different parameters, this is called *overloading* a method.

Operators like + or AND are a tool in a program that act on values they are given and result in a new value. So operators are very much like methods. In some languages, operators *are* methods (with unusual syntax) and they can be overloaded, so we can say, for instance, that in addition to what + means for numbers (addition) and what + means for strings (concatenation) we might define what + means for other types; this is primarily useful in languages where programmers can define their own types. We may also be able to define our own operators outside the basics.

Global Variables

So far, our variables have always been local. Up until methods, we just assumed all our code was in a method we didn't bother to name. Now each of our methods can define its own variables, but each such local variable only exists in that method while it is running.

It is often tempting to say that we would like to create a single variable that is accessible in all our methods, so they can all share it without having to constantly pass it in as parameter and return it. Such a variable is called a *global variable*.

```
double balance = 0
```

```
void deposit()
   double user = prompt "how much?"
   balance = balance + user
end deposit

void withdraw()
   double user = prompt "how much?"
   balance = balance - user
end withdraw

void menu()
   int user = -1
   constant int QUIT = 0;
   while (user != QUIT)
         print “1 to see balance”
         print “2 to deposit”
         print “3 to withdraw”
         print QUIT + “ to quit”
         user = prompt “Choice?”
         case user
               1: print “$” + balance
               2: deposit()
               3: withdraw()
         end case
   end while
end menu

void main()
   menu()
   print "Final Balance: $" + balance
end main
```

All the methods in this code had access to the balance variable. Some changed it. These methods were designed to be used with the balance. But every other method we added to the code would also have access to the balance, and would be able to change it. The more methods, the more complex the situation gets.

Global variables are one of those coding practices that make code superficially easier to write (especially if you are intimidated by passing in parameters and returning values) but in

practice often makes code harder to read, debug, and maintain. To use them without problems takes careful design.

In most cases, people limit global variables to constants, which gives us the advantage of creating named constants for readability that only have to be defined once and are available throughout code – so we are consistent about using the same value everywhere, but since they are constant, we never have to try to chase down which part of the code keeps messing up their value.

In object oriented design we have the idea of an instance variable, which gets us get some of the advantages of global variables while maintaining our future laziness by putting a limit on complexity. An instance variable is available to all methods within a specific code scope called a class. The instance variable represents a permanent characteristic that influences and may be changed by this specific set of methods, which represent behaviors of whatever object has that characteristic. The methods and the instance variable are designed to work together. But the instance variable is not available outside that scope and we have structured ways to control access to it.