
Chapter 7

Recursion

What is in This Chapter ?

This chapter explains the notion of ***recursion*** as it appears in computer science. We discuss briefly the idea of using recursion for simple **math-based problems** and for simple graphical drawing, such as **drawing fractal pictures** which are often found in computer science. Recursion usually involves thinking in a different manner that you may be used to. It is a very powerful technique, however, that can simplify algorithms and make your life easier ... once you understand it. A more thorough discussion of recursive techniques is covered in further courses.




7.1 Recursion

Recall in the first chapter, we discussed the notion of "divide-and-conquer" in regards to breaking a problem down into smaller, more manageable, modular pieces called functions or procedures.

Sometimes when we break a problem down into smaller pieces of the same type of problem, we are simply reducing the amount of information that we must process, since the problem is solved in the same manner ... it is just a smaller version of the same problem.

For example, there are many real-world examples where we often break problems into smaller ones of the same type in order to solve the bigger problem:

1. Jigsaw puzzles are solved in "steps": border, interesting portion, grass, sky, etc..
2. Math problems are broken down into smaller/simpler problems
3. Even climbing stairs eventually breaks down to climbing one step at a time.



Such problems can be solved *recursively*.

Recursion is a divide-and-conquer technique in which the problem being solved is expressed (or defined) in terms of smaller problems of the same type.

The word **Recursion** actually comes from a Latin word meaning "a running back". This makes sense because recursion is the process of actually "going off" and breaking down a problem into small pieces and then bringing the solutions to those smaller pieces back together to form the complete solution. So ...

- recursion **breaks down** a complex problem into **smaller sub-problems**
- the subproblems are **smaller** instances of the **same type** of problem.

So, when using recursion, we need to consider how to:

1. break the problem down into smaller sub-problems
2. deal with each smaller sub-problem
3. merge the results of the smaller sub-problems to answer the original problem

In fact, it's actually easier than this sounds. Most of the time, we simply take the original problem and **break/bite off** a small piece that we can work with. We simply keep biting off the small pieces of the problem, solve them, and then merge the results.



Consider the simple example of computing the factorial of an integer. You may recall that the function works as follows:

```
0! = 1
1! = 1
2! = 2x1
3! = 3x2x1
4! = 4x3x2x1
5! = 5x4x3x2x1
etc..
```

So, the function is defined non-recursively as follows:

```
1
N! = N x (N-1) x (N-2) x ... x 3 x 2 x 1    if N = 0
                                           if N ≥ 1
```

If you were asked to write a factorial function, you could do so with a FOR loop as follows:

Function: Factorial

n: a positive integer for which to find the factorial

```
1.  answer ← 1
2.  for each integer i from 2 to n do {
3.      answer ← answer * i
    }
4.  return answer
```

This code is simple and straight-forward to write.

The factorial problem, however, can also be defined recursively as follows:

```
1
N! = N x (N-1)!    if N = 0
                  if N ≥ 1
```

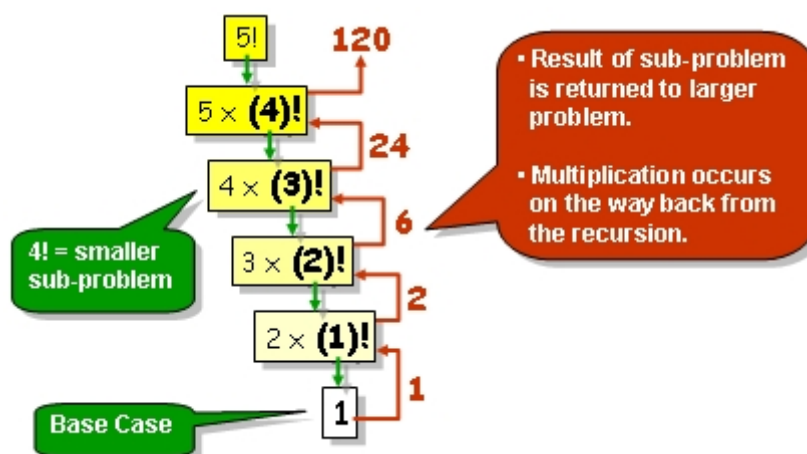
Notice how **N!** is expressed as a function of the smaller sub-problem of **(N-1)!** which is of the same type of problem (i.e., it is also a factorial problem, but for a smaller number).

So, if we had the solution for **(N-1)!**, say **S**, we could use this to compute the solution to **N!** as follows: **N! = N x S**

However, how do we get the solution for $(N-1)!$? We use the same formula, setting N to $N-1$ as follows:

$$(N-1)! = (N-1) \times ((N-1)-1)! = (N-1) \times (N-2)!$$

Similarly, we can recursively break down $(N-2)!$ in the same way. Eventually, as we keep reducing N by 1 each time, we end up with $N=0$ or 1 and for that simple problem we know the answer is 1. So breaking it all down we see the solution of $5!$ as follows:



If you were asked to write this **recursive** factorial function, you could do so as follows:

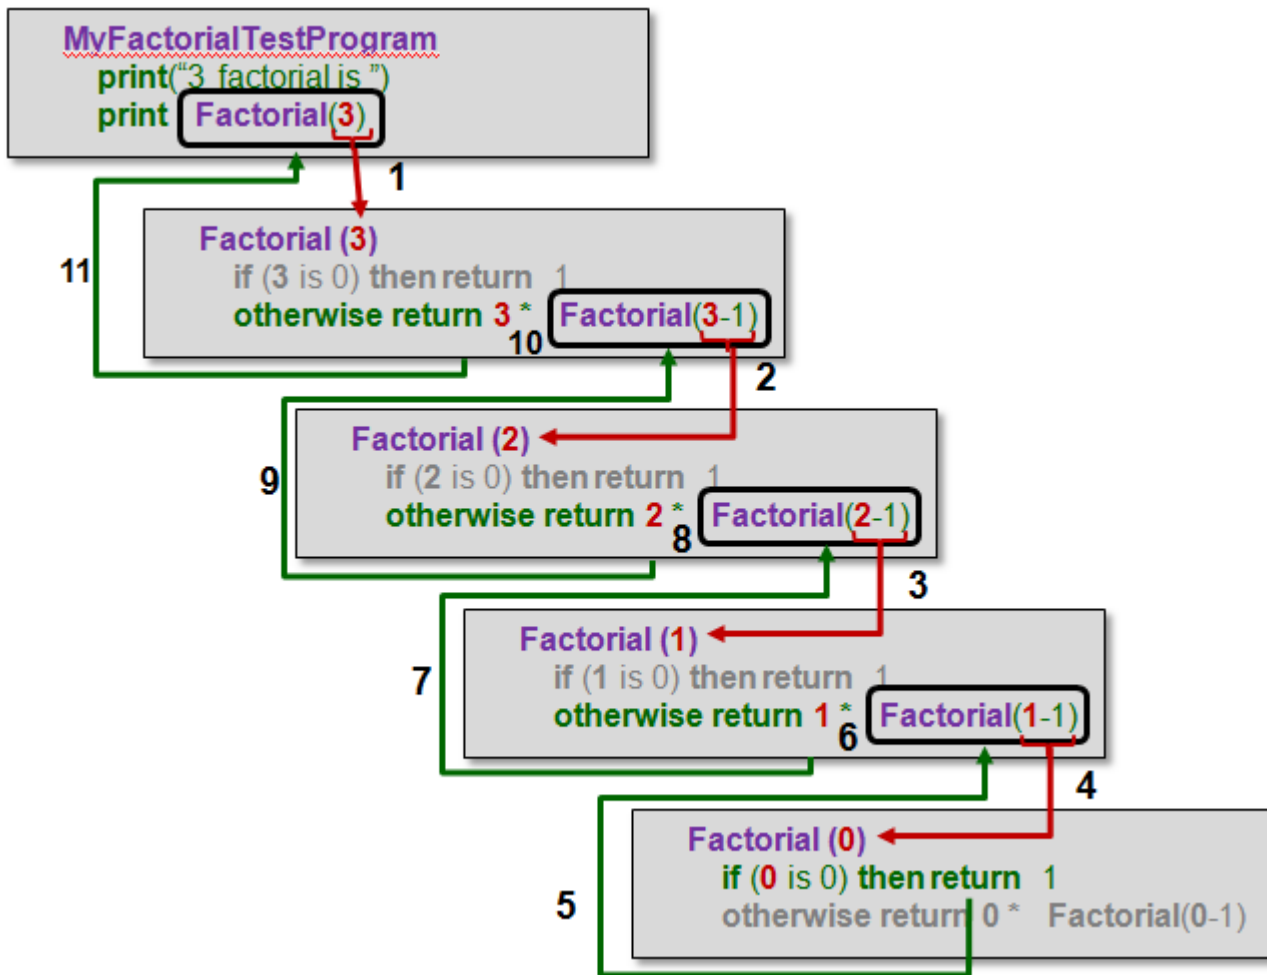
Function: FactorialRecursive

n: a positive integer for which to find the factorial

1. **if** (n is 0) **then**
2. **answer** \leftarrow 1
3. **otherwise**
4. **answer** \leftarrow n * **FactorialRecursive**($n-1$)
5. **return answer**

Notice how we no longer have a repeating loop in our code. Instead, the **FactorialRecursive** function calls itself with a smaller value of n . That is, the function repeatedly calls itself over and over again with smaller values of n until finally n has the smallest value of 0, in which case the code stops calling itself recursively. A recursive function is easily identifiable by the fact that a function or procedure calls itself. The "stopping situation" is called the **base case** of the problem. In general, a recursive function (or procedure) may have multiple base cases (i.e., multiple stopping conditions).

To understand the inner-workings of the function as it calls itself, consider the following diagram which shows how the control passes from one function call to the next when computing the factorial of 3. There are 11 main "steps" of the recursion as shown numbered in black bold font:



If you were to compare the non-recursive solution with the recursive solution ... which do you find to be simpler ?

Function: Factorial	Function: Factorial
n: a positive integer	n: a positive integer
1. answer \leftarrow 1	1. if (n is 0) then
2. for each integer i from 2 to n do {	2. answer \leftarrow 1
3. answer \leftarrow answer * i	3. otherwise
4. }	4. answer \leftarrow n * Factorial (n-1)
5. return answer	5. return answer

You may feel that the non-recursive version is simpler. In this particular example, you are probably right. So the question that arises is ... "Why should we use recursion ?".

Here are 4 reasons:

1. Some problems are **naturally** recursive and are easier to express recursively than non-recursively. For example, some math problems are defined recursively.
2. Sometimes, using recursion results in a simpler, **more elegant solution**.
3. Recursive solutions may actually be **easier to understand**.
4. In some cases, a recursive solution may be the **only way to approach** a seemingly **overwhelming problem**.

As we do various examples, the advantages of using recursion should become clear to you.

How do we write our own recursive functions and procedures ? It is important to remember the following two very important facts about the sub-problems:

- must be an instance of the **same kind** of problem
- must be **smaller** than the original problem



The trick is to understand how to "bite off" small pieces of the problem and knowing when to stop doing so. When we write the code, we will usually start with the base cases since they are the ones that we know how to handle. For example, if we think of our "real world" examples mentioned earlier, here are the base cases:



1. For the **jigsaw puzzle**, we divide up the pieces until we have just a few (maybe 5 to 10) pieces that form an interesting part of our picture. We stop dividing at that time and simply solve (by putting together) the simple base case problem of this small sub picture. So the base case is the problem in which there are only a few pieces all part of some identifiable portion of the puzzle.



2. For the **math problem**, we simply keep breaking down the problem until either (a) there is a simple expression remaining (e.g., $2 + 3$) or (b) we have a single number with no operations (e.g., 7). These are simple base cases in which there is either one operation to do or none.



3. For the **stair climbing problem**, each stair that we climb brings us closer to solving the problem of climbing all the stairs. Each time, we take a step up, the problem is smaller, and of the same type. The base case is our simplest case when there is only one stair to climb. We simply climb that stair and there is no recursion to do at that point.

So then to write recursive methods, we first need to understand "where" the recursion will fit in. That is, we need to understand how to reduce the problem and then merge the results. It is important to fully understand the problem and the kind of parameters that you will need as the function/procedure gets called repeatedly. Make sure you understand how the problem gets smaller (i.e., usually the parameters change in some way). Then, you can implement the simple base cases and add the recursion. Let us look at a few examples.

7.2 Math Examples

Recursive math examples are easy to program because we are often given the recursive definition directly and then it easily transfers into a program.

Example:

For example, consider computing how much money per month a person would have to pay to pay back a mortgage on his/her home. Consider the following notation:

- **amount** = mortgage amount borrowed (e.g., \$130,000)
- **rate** = annual interest rate as a percentage (e.g., 3.5%)
- **months** = length of time for the mortgage to be paid off (e.g., 300 months for a 25 year mortgage)
- **payment(amount, rate, months)** = monthly payment to pay off mortgage (e.g., \$647.57)



We would like to calculate the value for the **payment** function. Here is the function and its parameters as defined:

Function: **Payment**

amount: the principal amount of the mortgage
rate: annual interest rate to be paid
months: mortgage term (in months)

If we borrow money for 0 months, then the whole point of borrowing is silly. Assume then that **months** > 0. What if **months** = 1 ... that is ... we want to pay off the entire mortgage one month after we borrowed the money? Then we must pay one month interest and so the value of the **payment** function should return:

$$\text{amount} + \text{amount} * (\text{rate}/12) \quad \text{or} \quad \text{amount} * (1 + (\text{rate}/12))$$

which is one month of interest on the original loan amount. We divide by 12 since the rate is per year and we want only 1 month of interest. For a 2-month term, the payment would be:

$$[\text{amount} * (1 + (\text{rate}/12))] * (1 + (\text{rate}/12))$$

and for a 3-month term:

$$\{[\text{amount} * (1 + (\text{rate}/12))] * (1 + (\text{rate}/12))\} * (1 + (\text{rate}/12))$$

So you can see, the interest to be paid each month is compounded monthly...so we pay interest on interest owed and this grows larger each month.

There are a few ways to compute mortgage payments, but we will use the following formula so that we can practice our recursion. The formula for computing the payment is actually recursive and is defined as follows (showing first letter only for each parameter name):

$$p(a, r, m) = \frac{a}{\left[\frac{a}{p(a, r, m-1)} + \frac{1}{(1 + r/12)^m} \right]}$$

Looking at the formula, do you see that the problem is recursive? The result of the payment function $p(a, r, m)$ depends on the result of the payment function $p(a, r, m-1)$ for one less month. It is easy to see that the problem for one less month is the same problem (i.e., same function) and that the problem is smaller.

Now let us determine the base case (i.e., the simplest case). Well, the **amount** and **rate** do not change throughout the computation. Only the **month** value changes ... by getting smaller. So the base case must be the stopping condition for the months being decreased. This base case is the simplest case which follows from the definition.

Function: Payment

amount: the principal amount of the mortgage
rate: annual interest rate to be paid
months: mortgage term (in months), assumed to be > 0

1. **if (months is 1) then**
2. **return amount * (1 + rate/12)**

That wasn't so bad. Now what about the recursive part? It also follows from the formula:

Function: Payment

amount: the principal amount of the mortgage
rate: annual interest rate to be paid
months: mortgage term (in months), assumed to be > 0

1. **if (months is 1) then**
2. **return amount * (1 + rate/12)**
3. **return amount / ((amount / Payment(amount, rate, months-1)) + (1 + rate/12)^{months})**

As you can see, the code is easy to write as long as we have a recursive definition to begin with.

Example:

Imagine now that we have a teacher who wants to choose 10 students from a class of 25, to participate in a special project. Assuming that each student is unique in his/her own way ... how many different groups of students can be chosen ?



This is a classic problem using binomial coefficients that is often encountered in combinatorics for determining the amount of variety in a solution, giving insight as to the complexity of a problem. In general, if we have n items and we want to choose k items, we are looking for the solution of how many groups of k -elements can be chosen from the n elements.

The expression is often written as follows and pronounced "n choose k":

$$\binom{n}{k}$$

Assuming that n is fixed, we can vary k to obtain a different answer. The simplest solutions for this problem are when $k=0$ or $k=n$. If $k=0$, we are asking how many groups of zero can be chosen from n items. The answer is 1 ... there is only one way to choose no items. Similarly, if $k=n$ then we want to choose all items ... and there is only one way to do that as well. Also, if $k>n$, then there is no way to do this and the answer is zero (e.g., cannot choose 10 students from a group of 6).

Otherwise, we can express the problem as a recursive solution by examining what happens when we select a single item from the n items. Imagine therefore that we select one particular student that we definitely want to participate in the project (e.g., the teacher's "pet").

We need to then express the solution to the problem, taking into account that we are going to use that one student. In the remainder of the classroom, we therefore have $n-1$ students left and we still need to select $k-1$ students. Here is what is left →

$$\binom{n-1}{k-1}$$

The other situation is that we decide that we definitely DO NOT want a particular student to participate in the project (e.g., he's been a "bad boy" lately). In this case,

$$\binom{n-1}{k}$$

if we eliminate that student, we still need to select **k** students, from the **n-1** remaining. Here is what is left →

Now, I'm sure you will agree that if we examine one particular student ... then we will either select that student for the project or not. These two cases represent all possible solutions to the problem. Therefore, we can express the entire solution as:

$$\binom{n}{k} = \binom{n-1}{k-1} + \binom{n-1}{k}$$

This is a recursive definition. The problem is of the same type and gets smaller each time (since **n** gets smaller and **k** does also in one case). This problem is interesting because there are two recursive calls to the same problem. That is, the problem branches off twice recursively.

So then, to write the code, we start with the function:

Function: StudentCombinations

n: total number of students

k: number of students to choose

What about the base case(s) ? Well, we discussed the situations where **k=n**, **k=0**, **k>n**, so these are the simplest cases:

Function: StudentCombinations

n: total number of students

k: number of students to choose

1. **if (k > n) then**
2. **return 0**
3. **if (k is 0) then**
4. **return 1**
5. **if (k is n) then**
6. **return 1**

As before, the recursive case is easy, since we have the definition. This time, the code will be interesting as the function will call itself twice:

Function: StudentCombinations**n:** total number of students**k:** number of students to choose

```
1.  if (k > n) then
2.    return 0
3.  if (k is 0) then
4.    return 1
5.  if (k is n) then
6.    return 1
7.  return StudentCombinations(n-1,k-1) + StudentCombinations(n-1,k)
```

Not too bad ... was it ? Do you understand how the recursion will eventually stop ?

7.3 Graphical Examples

Recursion also appears in many graphical applications. For example, in computer games, trees are often created and displayed using recursion. Fractals are often used in computer graphics.


*A **fractal** is a rough or fragmented geometric shape that can be split into parts, each of which is (at least approximately) a reduced-size copy of the whole. (wikipedia)*


So, a fractal is often drawn by repeatedly re-drawing a specific pattern in a way that the pattern becomes smaller and smaller, but is of the same type.

Example:


Consider how we can draw a snowflake. See if you can detect the recursion in the following snowflake (called the **Koch Snowflake**):

It may not be easy to see. However ... if we were to examine the process of creating this snowflake recursively, notice how it can be done by starting with an equilateral triangle:






The recursion happens on each side of the triangle. Consider a single horizontal edge. We can recursively alter the edge into 4 segments as follows:



To accomplish this, we need to understand how a single edge is broken down into 4 pieces. Then we need to understand how this gets done recursively. Notice above how each recursive edge gets smaller. In fact, during each step of the recursion, the edge gets smaller by a factor of 3. Eventually, the edge will be very small in length (e.g., 2) and the recursion can stop, since breaking it down any further will not likely make a difference graphically.

Consider an edge going from a starting location **(Sx, Sy)** and going in direction **angle** towards edge **(Ex, Ey)**. Here is how we break the edge down into 4 pieces:



So the code for doing this, is as follows:

Function: drawKochEdge

Sx: x location to start drawing at
Sy: y location to start drawing at
Ex: x location to stop drawing at
Ey: y location to stop drawing at
angle: direction to draw towards

1. **length** \leftarrow distance from **(Sx, Sy)** to **(Ex, Ey)**
2. **Px** \leftarrow **Sx** + **length** * **cos**(**angle**)
3. **Py** \leftarrow **Sy** + **length** * **sin**(**angle**+180°)
4. **Qx** \leftarrow **Px** + **length** * **cos**(**angle** + 60°)
5. **Qy** \leftarrow **Py** + **length** * **sin**(**angle** + 60°+180°)
6. **Rx** \leftarrow **Qx** + **length** * **cos**(**angle** - 60°)
7. **Ry** \leftarrow **Qy** + **length** * **sin**(**angle** - 60°+180°)
8. draw line from **(Sx, Sy)** to **(Px, Py)**
9. draw line from **(Px, Py)** to **(Qx, Qy)**
10. draw line from **(Qx, Qy)** to **(Rx, Ry)**
11. draw line from **(Rx, Ry)** to **(Ex, Ey)**

The additional value of **180°** in the **cosine** function is used because of the reverse-**y**-coordinate system that is used in computer graphics. In mathematics, the origin **(0,0)** of the coordinate system is in the **bottom** left corner. In computer graphics, however, the origin **(0,0)** is at the **top** left corner. By adding **180°** to our angles we are compensating for the different origin, essentially flipping the **y**-coordinate.

The above code will work for any start/end locations. The angle is also passed in as a parameter, but alternatively it can be computed from the start and end points of the edge. It is important that the angle passed is actually the correct angle from the start to the end location.


Now, how do we draw the entire snowflake ? We need to draw three edges to represent the starting triangle. Assume that we want the snowflake centered at some location **(Cx, Cy)** and that the snowflake's size is defined by the length of the three triangle edges called **size**.

We can draw the triangle with simple straight line edges as follows:

Function: drawSnowflake

Cx, Cy: location of center of snowflake
size: size of snowflake (i.e., length of inner triangle edge)

1. $P_x \leftarrow C_x + \text{size} * \cos(90^\circ)$
2. $P_y \leftarrow C_y + \text{size} * \sin(90^\circ + 180^\circ)$
3. $Q_x \leftarrow C_x + \text{size} * \cos(-30^\circ)$
4. $Q_y \leftarrow C_y + \text{size} * \sin(-30^\circ + 180^\circ)$
5. $R_x \leftarrow C_x + \text{size} * \cos(210^\circ)$
6. $R_y \leftarrow C_y + \text{size} * \sin(210^\circ + 180^\circ)$
7. draw line from (P_x, P_y) to (Q_x, Q_y)
8. draw line from (Q_x, Q_y) to (R_x, R_y)
9. draw line from (R_x, R_y) to (P_x, P_y)




We can draw the triangle with a single Koch edge as follows:

Function: drawSnowflake

Cx, Cy: location of center of snowflake
size: size of snowflake (i.e., length of inner triangle edge)

1. $P_x \leftarrow C_x + \text{size} * \cos(90^\circ)$
2. $P_y \leftarrow C_y + \text{size} * \sin(90^\circ + 180^\circ)$
3. $Q_x \leftarrow C_x + \text{size} * \cos(-30^\circ)$
4. $Q_y \leftarrow C_y + \text{size} * \sin(-30^\circ + 180^\circ)$
5. $R_x \leftarrow C_x + \text{size} * \cos(210^\circ)$
6. $R_y \leftarrow C_y + \text{size} * \sin(210^\circ + 180^\circ)$
7. **drawKochEdge**($P_x, P_y, Q_x, Q_y, -60^\circ$)
8. **drawKochEdge**($Q_x, Q_y, R_x, R_y, 180^\circ$)
9. **drawKochEdge**($R_x, R_y, P_x, P_y, 60^\circ$)



The code above will draw the shape with one level of recursion. But how do we complete the code so that it does many levels of recursion ? We need to come up with a base case for stopping the recursion. This can be done in the **drawKochEdge** procedure by noticing when

the length of the edge is very small (e.g., 2). In this case, we draw a regular line as opposed to breaking it up into 4 and recursively drawing. Notice the code:

Function: **drawKochEdge**

Sx, Sy: location to start drawing at
Ex, Ey: location to stop drawing at
angle: direction to draw towards

```


1.  length ← distance from (Sx, Sy) to (Ex, Ey)
2.  if (length < 2) then
3.      draw line from (Sx, Sy) to (Ex, Ey)           // just draw a straight line
    otherwise {
4.      Px ← Sx + length * cos(angle)
5.      Py ← Sy + length * sin(angle+180°)
6.      Qx ← Px + length * cos(angle + 60°)
7.      Qy ← Py + length * sin(angle + 60°+180°)
8.      Rx ← Qx + length * cos(angle - 60°)
9.      Ry ← Qy + length * sin(angle - 60°+180°)
10.     drawKochEdge(Sx, Sy, Px, Py, angle)
11.     drawKochEdge(Px, Py, Qx, Qy, angle + 60°)
12.     drawKochEdge(Qx, Qy, Rx, Ry, angle - 60°)
13.     drawKochEdge(Rx, Ry, Ex, Ey, angle)
    }
```

This code will produce the desired snowflake pattern.

Example:

Consider drawing a tree. We can draw the stem of a tree as a branch and then repeatedly draw smaller and smaller branches of the tree until we reach the end of the tree. Here, for example, is a fractal tree →



How can we draw this? Well, imagine drawing just one portion of this tree starting at the bottom facing upwards. We can travel for some number of pixels (e.g., 100) and then bend to the right, perhaps 40° as shown here. Then, we can draw for a smaller distance, perhaps 60% of the previous branch. We can do this repeatedly, turning 40° and drawing a new smaller branch. We can stop when the branch is small enough (perhaps 2 or 3 pixels in length).



Consider doing this for just one portion of branches in the tree (e.g, the right sequence of branches shown here in bold/red).

Assume that the **bendAngle** is set at 40° and that the **shrinkFactor** is 0.6 (i.e., 60%). Consider the code to draw the rightmost set of branches of the tree. Assume that the initial **(x,y)** location is centered at the bottom of the window and that the drawing direction is north (i.e., upwards at 90°).

Each time we can draw a single branch, then turn **bendAngle** and then recursively draw the remaining branches starting with branch length reduced by the **shrinkFactor**.




Here is the code:

Function: **drawSingleTreePath**

length: length of tree branch
x, y: location to start drawing at
dir: direction to draw towards

1. **if** (**length** < 2)
2. **return**
3. $\text{endX} \leftarrow x + \text{length} * \cos(\text{dir} + \text{bendAngle})$
4. $\text{endY} \leftarrow y + \text{length} * \sin(\text{dir} + \text{bendAngle} + 180^\circ)$
5. draw line from **(x, y)** to **(endX, endY)**
6. **drawTree**(**length** * **shrinkFactor**, **endX**,
 endY, **dir** + **bendAngle**)



Assume that we begin the tree by drawing a 100 pixel length vertical line from the bottom center of the screen (**windowWidth/2, windowHeight**) facing upwards (i.e. North at 90°) to (**windowWidth /2, windowHeight -100**) We would complete the drawing simply by calling the function with initial values corresponding to facing North and a 100***shrinkFactor** = 60 pixel length as follows:

drawSingleTreePath (60, **windowWidth /2, windowHeight -100, 90°**)

The recursion stops when the tree branch size reaches 2, which is somewhat arbitrary. To draw the remainder of the tree, we simply add code to draw left as well:

Function: drawTree


length: length of tree branch
x, y: location to start drawing at
dir: direction to draw towards

```




1.  if (length < 2)
2.      return
3.  endX ← x + length * cos(dir + bendAngle)
4.  endY ← y + length * sin(dir + bendAngle+180°)
5.  draw line from (x, y) to (endX, endY)
6.  drawTree(length * shrinkFactor, endX, endY, dir + bendAngle)


7.  endX ← x + length * cos(dir - bendAngle)
8.  endY ← y + length * sin(dir - bendAngle+180°)
9.  draw line from (x, y) to (endX, endY)
10. drawTree(length * shrinkFactor, endX, endY, dir - bendAngle)

```





Notice how the tree takes shape because each time we draw a single branch, we branch out 40° left and 40° right. Notice how the trees differ as we change the **bendAngle** to other angles other than 40°:





Also, as we adjust the **shrinkFactor** to other values, notice how the tree changes (here the **bendAngle** is fixed at 45°):




Example:

Now, consider drawing a more realistic tree. A real tree does not have branches that grow at exact angles. Let us begin by adjusting the **bendAngle** to be random instead of fixed at say **40°**. That way, each branch of the tree can bend randomly. Of course, we want to have some limits as to how random they should bend.

We can change the **bendAngle** so that it bend at least some amount each time (e.g., $\pi/32 = 5.6^\circ$) plus some random amount (e.g., 0 to $\pi/4$... or 0° to 45°):

$$\text{bendAngle} \leftarrow 5.6^\circ + (\text{random value from } 0^\circ \text{ to } 45^\circ)$$

The result is that our trees will look less symmetrical and will bend more naturally:



Another factor that we can adjust is the length of each branch. Currently, they shrink by a constant factor of **shrinkFactor** which is 0.6 in the examples above. However, we can shrink them by a random amount (within reason, of course). So, we can add or subtract a random amount from the **shrinkFactor**. It is easiest to simply multiply the **shrinkFactor** by some value, perhaps in the range from 1.0 to 1.5. That will allow some branches to be larger than 60% of their "parent" branch.

Here is the code as we now have it using a variable **bendAngle** and a new variable called **stretch** to allow randomly longer branches:

Function: drawNaturalTree

length:	length of tree branch
x, y	location to start drawing at
dir:	direction to draw towards

... continued on next page ...

```

1.  if (length < 2)
2.      return

3.  bendAngle ← 5.6° + (random value from 0° to 45°)
4.  endX ← x + length * cos(dir + bendAngle)
5.  endY ← y + length * sin(dir + bendAngle+180°)
6.  draw line from (x, y) to (endX, endY)
7.  stretch ← 1 + (random value from 0 to 0.5)
8.  drawNaturalTree (length * shrinkFactor * stretch, endX, endY, dir + bendAngle)


9.  bendAngle ← 5.6° + (random value from 0° to 45°)
10. endX ← x + length * cos(dir - bendAngle)
11. endY ← y + length * sin(dir - bendAngle+180°)
12. draw line from (x, y) to (endX, endY)
13. stretch ← 1 + (random value from 0 to 0.5)
14. drawNaturalTree(length * shrinkFactor * stretch, endX, endY, dir - bendAngle)

```


Here are two examples of results due to this change:



Notice how tree is more "bushy". That is because the recursion goes further before it stops (i.e., it draws more branches). This is because the base case of the recursion stops only when the branch reaches a specific small size. Due to the randomness of the parameters, we may end up with some lopsided trees:



Of course, we can make everything even *bushier* if we increase the **shrinkFactor** to 65%:



All that remains now is to make the branches thicker and color them brown. For very small branches, we color them green:

Function: **drawColorTree**


length: length of tree branch
x: x location to start drawing at
y: y location to start drawing at
dir: direction to draw towards

```

1.  if (length < 2)
2.      return
2.  for b from -1 to 1 by 2 do {
3.      set line drawing width to length / 7
4.      if (length > 10) then
5.          set line drawing color to brown
6.      otherwise {
7.          set line drawing color to a random amount of green (e.g., 80 to 130)
8.          set line drawing width to 2 + (length / 7)
9.      }
9.      bendAngle ← 5.6° + (random value from 0° to 45°)
10.     endX ← x + length * cos(dir + b*bendAngle)
11.     endY ← y + length * sin(dir + b*bendAngle+180°)
12.     draw line from (x, y) to (endX, endY)
13.     stretch ← 1 + (random value from 0 to 0.5)
14.     drawColorTree (length * shrinkFactor * stretch, endX, endY, dir + b*bendAngle)
    }
```

Notice the **FOR** loop now which repeats exactly 2 times. This reduces our code since the code for drawing the left branch is the same as that for drawing the right branch, except that the angle offset is negative (represented by the value of **b**).

Here is the result (with **shrinkFactor** at 67% and **bendAngle** of $(\pi/32 + \text{random } 0 - \pi/6)$):



7.4 Search Examples

Sometimes, we use recursion to search for an answer to a problem. Searching for an answer recursively can often lead to a simpler, more elegant solution than writing search code that uses many loops and IF statements.

Example:

Consider writing code to solve a rubik's cube. The task may seem daunting at first ... but there are many systematic ways of solving the cube in steps or phases. Any non-recursive code that you would write could be quite complex, tedious and/or cumbersome. However, a recursive solution is quite simple.

Assume that the cube can be solved by making at most 100 quarter turns of one of the cube's sides (a reasonable assumption). Here is the idea behind the recursive solution:

Try turning 1 of the 6 sides 1/4 of a turn, then try to solve the cube in 99 more turns, recursively. If the recursion comes back with a solution from the 99 additional moves, then the 1/4 turn that we made was a good one, hence part of the solution. If however, the recursion came back after 99 moves with an unsolved cube, then the 1/4 turn that we made was a bad choice ... undo it and then try turning one of the other 5 sides.



Here is the pseudocode for a function to solve the cube recursively in **n** steps, returning **true** if it has been solved, or **false** otherwise:

Function: SolveRubiksCube

cube: the cube data structure
n: maximum number of steps remaining to solve the cube

```

1.  if (cube is solved) then
2.      return true
3.  if (n is 0) then
4.      return false
5.  for each side of the cube s from 1 to 6 do {
6.      turn side s 1/4 turn clockwise
7.      recursiveAnswer ← SolveRubiksCube(cube, n-1)
8.      if (recursiveAnswer is true) then
9.          return true
10.     turn side s 1/4 turn counter-clockwise    // i.e., undo the turn
11. }
12. return false

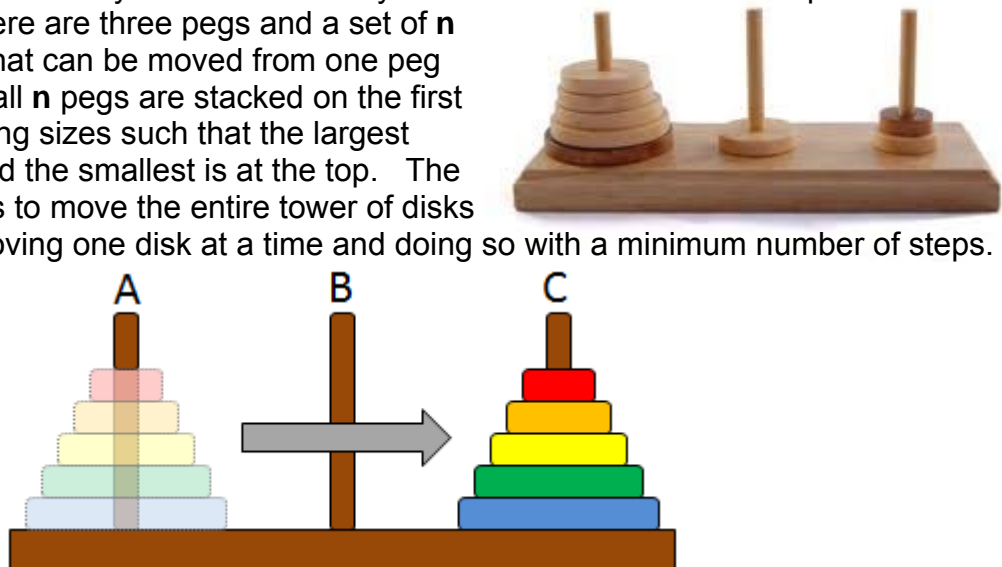
```

That is the solution! It is quite simple. Do you see the power of recursion now? Of course, this particular solution can be quite slow, as it blindly tries all possible turns without any particular strategy. Nevertheless, the solution is simple to understand. It would be difficult to produce such a simple function to solve this problem without using recursion.

Example:

Another problem that is naturally solved recursively is that the **Towers of Hanoi** puzzle.

In this puzzle game, there are three pegs and a set of **n** disks of various sizes that can be moved from one peg to another. To begin, all **n** disks are stacked on the first peg in order of increasing sizes such that the largest disk is at the bottom and the smallest is at the top. The objective of the game is to move the entire tower of disks to a different peg by moving one disk at a time and doing so with a minimum number of steps.




There are just two rules:

1. disks must be moved from one peg to another, not placed aside on the table
2. a disk can never be placed onto a smaller-sized disk.


To approach this problem recursively, we must understand how to break the problem down into a smaller problem of the same type. Since all we are doing is moving disks, then a "smaller" sub-problem would logically just be moving less disks. For example, if 1 disk is moved, then only $n-1$ disks remain to be moved ... which is a smaller problem.

So what then would be the base case? Well ... what is the simplest number of disks to move? Zero. A simple problem is also the case where $n = 1$. In that case, we simply pick up the disk and move it. For values of n greater than 1, we now have to worry about moving some disks out of the way. Here is the solution when $n = 2 \rightarrow$



Notice that after the small red disk is moved to peg B, then we have the smaller/simpler problem of moving the one larger disk from peg A to peg C. Finally, we have the simpler/smaller problem of moving the red disk from peg B to peg C and we are done.

It may be hard to see the recursion, at this point, so let us look at the solution for $n = 3$:




You may notice that steps 1 to 3 are basically the same as the problem where $n = 2$, except that we are moving the disks from peg **A** to peg **B** instead of to peg **C**. Then step 4 simply moves the 1 disk from peg **A** to peg **C**. Finally, we are left with the recursive sub-problem again of moving 2 disks from peg **B** to peg **C**.

So, the recursive solution of moving n disks is based on this idea:

1. Move the $n-1$ smaller disks out of the way to the spare peg **B**
2. Move the large disk to peg **C**
3. Move the $n-1$ smaller disks to peg **C**

As you can see above, there are two recursive sub-problems corresponding to moving $n-1$ disks out of the way first, and then solving the $n-1$ problem once the largest bottom disk is in place.

So, we can now write the code by defining a procedure that indicates the peg to move from, the peg to move to, the extra peg and the number of disks to move.



Procedure: TowersOfHanoi

pegA: the peg to move from
pegC: the peg to move to
pegB: the extra peg
n: number of disks to move

1. **if** (n is 1) **then**
2. **Move 1 disk from pegA to pegC**
3. **otherwise** {
4. **TowersOfHanoi**(pegA, pegB, pegC, $n-1$)
5. **Move 1 disk from pegA to pegC**
6. **TowersOfHanoi**(pegB, pegC, pegA, $n-1$)
7. }

As you can see, the problem is expressed in a simple manner with only about 5 lines. Notice that there was no need to check if $n = 0$ since that would represent a problem that is already solved.