

Exploring Languages  
with Interpreters  
and Functional Programming  
Chapter 2

H. Conrad Cunningham

02 April 2022

Contents

<b>2</b>	<b>Programming Paradigms</b>	<b>2</b>
2.1	Chapter Introduction . . . . .	2
2.2	Abstraction . . . . .	2
2.2.1	What is abstraction? . . . . .	2
2.2.2	Kinds of abstraction . . . . .	2
2.2.3	Procedures and functions . . . . .	3
2.3	What is a Programming Paradigm? . . . . .	4
2.4	Imperative Paradigm . . . . .	4
2.4.1	Java . . . . .	4
2.4.2	Other languages . . . . .	5
2.5	Declarative Paradigm . . . . .	5
2.5.1	Functional paradigm . . . . .	6
2.5.1.1	Haskell . . . . .	6
2.5.1.2	Other languages . . . . .	8
2.5.2	Relational (or logic) paradigm . . . . .	8
2.5.2.1	Prolog . . . . .	8
2.5.2.2	Other languages . . . . .	9
2.6	Other Programming Paradigms . . . . .	9
2.6.1	Procedural paradigm . . . . .	10
2.6.1.1	Python . . . . .	10
2.6.1.2	Other languages . . . . .	11
2.6.2	Modular paradigm . . . . .	11
2.6.2.1	Python . . . . .	11
2.6.2.2	Other languages . . . . .	16
2.6.3	Object-based paradigms . . . . .	16
2.6.4	Concurrent paradigms . . . . .	16

2.7	Motivating Functional Programming: John Backus . . . . .	16
2.7.1	Excerpts from Backus's Turing Award Address [1] . . . . .	17
2.7.2	Aside on the disorderly world of statements . . . . .	19
2.7.3	Perspective from four decades later . . . . .	19
2.8	What Next? . . . . .	20
2.9	Exercises . . . . .	20
2.10	Acknowledgements . . . . .	20
2.11	Terms and Concepts . . . . .	21
2.12	References . . . . .	22

Copyright (C) 2016, 2017, 2018, 2022, H. Conrad Cunningham  
 Professor of Computer and Information Science  
 University of Mississippi  
 214 Weir Hall  
 P.O. Box 1848  
 University, MS 38677  
 (662) 915-7396 (dept. office)

**Browser Advisory:** The HTML version of this textbook requires a browser that supports the display of MathML. A good choice as of April 2022 is a recent version of Firefox from Mozilla.

## 2 Programming Paradigms

### 2.1 Chapter Introduction

The goals of this chapter are to:

- introduce the concepts of procedural and data abstraction
- examine the characteristics and concepts the primary programming paradigms, imperative and declarative (including functional and relational)
- survey other paradigms such as procedural and modular programming

### 2.2 Abstraction

Programming concerns the construction of appropriate abstractions in a programming language. Before we examine programming paradigms, let's examine the concept of abstraction.

#### 2.2.1 What is abstraction?

As computing scientists and computer programmers, we should remember the maxim:

*Simplicity is good; complexity is bad.*

The most effective weapon that we have in the fight against complexity is *abstraction*. What is abstraction?

Abstraction is *concentrating on the essentials and ignoring the details*.

Sometimes abstraction is described as *remembering the “what” and ignoring the “how”*.

Large complex problems can only be made understandable by decomposing them into subproblems. Ideally, we should be able to solve each subproblem independently and then compose their solutions into a solution to the larger problem.

In programming, the subproblem solution is often expressed with some kind of abstraction represented in a programming notation. From the outside, each abstraction should be simple and easy for programmers to use correctly. The programmers should only need to know the abstraction's *interface* (i.e., some small number of assumptions necessary to use the abstraction correctly).<sup>4</sup>

#### 2.2.2 Kinds of abstraction

Two kinds of abstraction are of interest to computing scientists: *procedural abstraction* and *data abstraction*.

**Procedural abstraction:** the separation of the logical properties of an *action* from the details of how the action is implemented.

**Data abstraction:** the separation of the logical properties of *data* from the details of how the data are represented.

In procedural abstraction, programmers focus primarily on the actions to be carried out and secondarily on the data to be processed.

For example, in the top-down design of a sequential algorithm, a programmer first identifies a sequence of actions to solve the problem without being overly concerned about how each action will be carried out.

If an action is simple, the programmer can code it directly using a sequence of programming language statements.

If an action is complex, the programmer can abstract the action into a subprogram (e.g., a procedure or function) in the programming language. The programmer must define the subprogram's name, parameters, return value, effects, and assumptions—that is, define its interface. The programmer subsequently develops the subprogram using the same top-down design approach.

In data abstraction, programmers primarily focus on the problem's data and secondarily on its actions. Programmers first identify the key data representations and develop the programs around those and the operations needed to create and update them.

We address procedural and data abstraction further in Chapters 6 and 7.

### 2.2.3 Procedures and functions

Generally we make the following distinctions among subprograms:

- A *procedure* is (in its pure form) a subprogram that takes zero or more arguments but does not return a value. It is executed for its effects, such as changing values in a data structure within the program, modifying its reference or value-result arguments, or causing some effect outside the program (e.g., displaying text on the screen or reading from a file).
- A *function* is (in its pure form) a subprogram that takes zero or more arguments and returns a value but that does not have other effects.
- A *method* is a procedure or function often associated with an object or class in an object-oriented program. Some object-oriented languages use the metaphor of message-passing. A method is the feature of an object that receives a message. In an implementation, a method is typically a procedure or function associated with the (receiver) object; the object may be an *implicit parameter* of the method.

Of course, the features of various programming languages and usual practices for their use may not follow the above pure distinctions. For example, a language may not distinguish between procedures and functions. One term or another may be used for all subprograms. Procedures may return values. Functions may

have side effects. Functions may return multiple values. The same subprogram can sometimes be called either as a function or procedure.

Nevertheless, it is good practice to maintain the distinction between functions and procedures for most cases in software design and programming.

### 2.3 What is a Programming Paradigm?

According to Timothy Budd, a *programming paradigm* is “a way of conceptualizing what it means to perform computation, of structuring and organizing how tasks are to be carried out on a computer” [3:3].

Historically, computer scientists have classified programming *languages* into one of two primary paradigms: *imperative* and *declarative*.

This imperative-declarative taxonomy categorizes programming styles and language features on how they handle state and how they execute programs.

In recent years, many imperative languages have added more declarative features, so the distinction between languages has become blurred. However, the concept of *programming* paradigm is still meaningful.

### 2.4 Imperative Paradigm

A program in the imperative paradigm has an *implicit state* (i.e., values of variables, program counters, etc.) that is modified (i.e., side-effected or mutated) by *constructs* (i.e., commands) in the source language [19].

As a result, such languages generally have an explicit notion of *sequencing* (of the commands) to permit precise and deterministic control of the state changes.

Imperative programs thus express *how* something is to be computed. They emphasize procedural abstractions.

#### 2.4.1 Java

Consider the following Java program fragment from file `Counting.java`:

```
int count = 0 ;
int maxc  = 10 ;
while (count <= maxc) {
    System.out.println(count) ;
    count = count + 1 ;
}
```

In this fragment, the program’s *state* includes at least the values of the variables `count` and `maxc`, the sequence of output lines that have been printed, and an indicator of which statement to execute next (i.e., location or program counter).

The assignment statement changes the value of `count` and the `println` statement adds a new line to the output sequence. These are *side effects* of the execution.

Similarly, Java executes these commands in sequence, causing a change in which statement will be executed next. The purpose of the `while` statement is to cause the statements between the braces to be executed zero or more times. The number of times depends upon the values of `count` and `maxc` and how the values change within the `while` loop.

We call this state *implicit* because the aspects of the state used by a particular statement are not explicitly specified; the state is assumed from the context of the statement. Sometimes a statement can modify aspects of the state that are not evident from examining the code fragment itself.

The Java variable `count` is *mutable* because its value can change. After the declaration, `count` has the value 0. At the end of the first iteration of the `while` loop, it has value 1. After the `while` loop exits, it has a value 10. So a reference to `count` yields different values depending upon the state of the program at that point.

The Java variable `maxc` is also *mutable*, but this code fragment does not change its value. So `maxc` could be replaced by an *immutable* value.

Of course, the Java fragment above must be included within a `main` method to be executed. A `main` method is the entry point of a Java program.

```
public class Counting {
    public static void main(String[] args) {
        /* Java code fragment above */
    }
}
```

Imperative languages are the “conventional” or “von Neumann languages” discussed by John Backus in his 1977 Turing Award address [1]. (See Section 2.7.) They are suited to traditional computer architectures.

Most of the languages in existence today are primarily imperative in nature. These include Fortran, C, C++, Java, Scala, C#, Python, Lua, and JavaScript.

#### 2.4.2 Other languages

The Scala [23,26] program `CountingImp.scala` is equivalent to the Java program described above. The program `CountingImp2.scala` is also equivalent, except that it makes the `maxc` variable *immutable*. That is, it can be bound to an initial value, but its binding cannot be changed subsequently.

### 2.5 Declarative Paradigm

A program in the declarative paradigm has *no implicit* state. Any needed state information must be handled explicitly [19].

A program is made up of *expressions* (or terms) that are *evaluated* rather than commands that are executed.

Repetitive execution is accomplished by *recursion* rather than by sequencing.

Declarative programs express *what* is to be computed (rather than how it is to be computed).

The declarative paradigm is often divided into two types: *functional* (or applicative) and *relational* (or logic).

### 2.5.1 Functional paradigm

In the functional paradigm the underlying model of computation is the mathematical concept of a *function* [19].

In a computation, a function is applied to zero or more arguments to compute a single result; that is, the result is deterministic (or predictable).

**2.5.1.1 Haskell** Consider the following Haskell code from file `Counting.hs`:

```
counter :: Int -> Int -> String
counter count maxc
  | count <= maxc = show count ++ "\n"
                  ++ counter (count+1) maxc
  | otherwise     = ""
```

This fragment is similar to the Java fragment above. This Haskell code defines a function `counter` (i.e., a procedural abstraction) that takes two integer arguments, `count` and `maxc`, and returns a string consisting of a sequence of lines with the integers from `count` to `maxc` such that each would be printed on a separate line. (It does not print the string, but it inserts a newline character at the end of each line.)

In the evaluation (i.e., “execution”) of a function call, Programming for the {Newton}: Software Development with {NewtonScript}`counter` references the *values* of `count` and `maxc` corresponding to the explicit arguments of the function call. These values are not changed during the evaluation of that function call. However, the values of the arguments can be changed as needed for a subsequent *recursive* call of `counter`.

We call the state of `counter` *explicit* because it is passed in arguments of the function call. These parameters are *immutable* (i.e., their values cannot change) within the body of the function. That is, any reference to `count` or `maxc` within a call gets the same value.

In a pure functional language like Haskell, the names like `count` and `maxc` are said to be *referentially transparent*. In the same context (such as the body of the function), they always have the same value. A name must be defined before it is used, but otherwise the order of evaluation of the expressions within a function body does not matter; they can even be evaluated in parallel.

There are no “loops”. The functional paradigm uses recursive calls to carry out a task repeatedly.

As we see in later chapters, referential transparency is probably the most important property of functional programming languages. It underlies Haskell’s evaluation model (Chapter 8). It also underlies the ability to state and prove “laws” about Haskell programs (e.g., Chapters 25 and 26). Haskell programmers and Haskell compilers can use the “mathematical” properties of the programs to transform programs that are more efficient.

The above Haskell fragment does not really carry out any actions; it just defines a mapping between the arguments and the return value. We can “execute” the `counter` function above with the arguments 0 and 10 with the following **IO** program.

```
main = do
    putStrLn (counter 0 10)
```

By calling the `main` function from the `ghci` interpreter, we get the same displayed output as the Java program.

Haskell separates pure computation (as illustrated by function `counter`) from computation that has effects on the environment such as input/output (as illustrated by **IO** function `main`).

In most programming languages that support functional programming, functions are treated as *first-class* values. That is, like other data types, functions can be stored in data structures, passed as arguments to functions, and returned as the results of functions. (The implementation technique for first-order functions usually involves creation of a *lexical closure* holding the function and its environment.)

In some sense, functional languages such as Haskell merge the concepts of procedural and functional abstraction. Functions are procedural abstractions, but they are also data.

A function that can take functions as arguments or return functions in the result is called a *higher-order function*. A function that does not take or return functions is thus a *first-order function*. Most imperative languages do not fully support higher-order functions.

The higher-order functions in functional programming languages enable regular and powerful abstractions and operations to be constructed. By taking advantage of a library of higher-order functions that capture common patterns of computation, we can quickly construct concise, yet powerful, programs.

Purely functional languages include Haskell, Idris, Miranda, Hope, Elm, and Backus’s FP.

Hybrid functional languages with significant functional subsets include Scala, F#, OCaml, SML, Erlang, Elixir, Lisp, Clojure, and Scheme.



Mainstream imperative languages such as Java (beginning with version 8), C#, Python, Ruby, Groovy, Rust, and Swift have recent feature extensions that make them hybrid languages as well.

**2.5.1.2 Other languages** The Scala [23,26] program `CountingFun.scala` is equivalent to the above Haskell program.

## 2.5.2 Relational (or logic) paradigm

In the relational (logic) paradigm, the underlying model of computation is the mathematical concept of a *relation* (or a *predicate*) [19].

A computation is the (nondeterministic) association of a group of values—with backtracking to resolve additional values.

**2.5.2.1 Prolog** Consider the following Prolog [6] code from file `Counting.pl`. In particular, this code runs on the SWI-Prolog interpreter [27].

```
counter(X,Y,S) :- count(X,Y,R), atomics_to_string(R,'\n',S).

count(X,X,[X]).
count(X,Y,[]) :- X > Y.
count(X,Y,[X|Rs]) :- X < Y, NX is X+1, count(NX,Y,Rs).
```

This fragment is somewhat similar to the Java and Haskell fragments above. It can be used to generate a string with the integers from `X` to `Y` where each integer would be printed on a separate line. (As with the Haskell fragment, it does not print the string.)

This program fragment defines a *database* consisting of four *clauses*.

The clause

```
count(X,X,[X]).
```

defines a *fact*. For any *variable* value `X` and list `[X]` consisting of the single value `X`, `count(X,X,[X])` is asserted to be true.

The other three clauses are *rules*. The left-hand-side of `:-` is true if the right-hand-side is also true. For example,

```
count(X,Y,[]) :- X > Y.
```

asserts that

```
count(X,Y,[])
```

is true when `X > Y`. The empty brackets denote an empty list of values.

As a logic or relational language, we can *query* the database for any missing components. For example,

```
count(1,1,Z).
```

yields the value `Z = [1]`. However,

```
count(X,1,[1]).
```

yields the value `X = 1`. If more than one answer is possible, the program can generate all of them in some nondeterministic order.

So, in some sense, where imperative and functional languages only run a computation in one direction and give a single answer, Prolog can potentially run a computation in multiple directions and give multiple answers.

As with Haskell, the above Prolog fragment does not really carry out any computational actions; it just adds facts to the database and defines general relationships among facts. We can “execute” the query `counter(0,10,S)` above and print the value of `S` using the following rule.

```
main :- counter(0,10,S), write(S).
```

Example relational languages include Prolog, Parlog, and miniKanren.

Most Prolog implementations have imperative features such as the “cut” and the ability to assert and retract clauses.

**2.5.2.2 Other languages** TODO: Perhaps add a new example using miniKanren [4,5,14] in some reasonable base language—preferably Java, Python, or Scala.

## 2.6 Other Programming Paradigms

As we noted, the imperative-declarative taxonomy described above divides programming styles and language features on how they handle state and how they are executed.

The computing community often speaks of other paradigms—procedural, modular, object-oriented, concurrent, parallel, language-oriented, scripting, reactive, and so forth. The definitions of these “paradigms” may be quite fuzzy and vary significantly from one writer to another.

Sometimes a term is chosen for “marketing” reasons—to associate a language with some trend even though the language may be quite different from others in that paradigm—or to make a language seem different and new even though it may not be significantly different.

These paradigms tend to divide up programming styles and language features along different dimensions than the primary taxonomy described in Sections 2.4 and 2.5. Often the languages we are speaking of are subsets of the imperative paradigm.

This section briefly discusses some of these paradigms. We discuss the prominent object-based paradigms in the next chapter.

### 2.6.1 Procedural paradigm

The *procedural paradigm* is a subcategory of the imperative paradigm. It organizes programs primarily using procedural abstractions. A procedural program consists of a sequence of steps that access and modify the program’s state.

Some of the steps are abstracted out as subprograms—procedures or functions—that can be reused. In some cases, subprograms may be nested inside other subprograms, thus limiting the part of the program in which the nested subprogram can be called.

The procedural programming approach arose in programming languages such as Fortran, Algol, PL/I, Pascal, and C from the 1950’s to the 1970’s and beyond. In this chapter, we use the Python programming language to illustrate its features.

**2.6.1.1 Python** Consider the following Python [25] code from file `CountingProc.py`:

```
# File CountingProc.py
def counter(count,maxc):
    def has_more(count,maxc): # new variables
        return count <= maxc
    def adv():
        nonlocal count      # from counter
        count = count + 1
    while has_more(count,maxc):
        print(f'{count}') # Python 3.6+ string interpolation
        adv()
```

When called as

```
counter(0,10)
```

this imperative Python “procedure” executes similarly to the Java program fragment we examined in Section 2.4.

Python does not distinguish between procedures and functions as we have defined them. It uses the term “function” for both. Both return values and can have side-effects. The value returned may be the special default value `None`.

This Python code uses procedural abstraction more extensively than the earlier Java fragment. The Python procedure encloses the `while` loop in procedure `counter` and abstracts the loop test and incrementing operation into function `has_more` and procedure `adv`, respectively.

Like many procedural languages, Python uses *lexical scope* for variable, procedure, and function names. That is, the *scope* of a name (i.e., range of code in which it can be accessed) begins at the point it is defined and ends at the end of that block of code (e.g., function, class, or module).

Function `has_more` and procedure `adv` are encapsulated within `counter`. They can only be accessed inside the body of `counter` after their definitions.

Parameters `count` and `maxc` of procedure `counter` can be accessed throughout the body of `counter` unless hidden by another variable or parameter with the same name. They are hidden within the function `has_more`, which reuses the names for its parameters, but are accessible within procedure `adv`.

But to allow assignment to `count` within the nested procedure `adv`, the variable must be declared as `nonlocal` in the inner procedure. Otherwise, the assignment would have created a new variable with the name `count` within the body of procedure `adv`.

Languages like Python, C, Fortran, Pascal, and Lua are primarily procedural languages, although most have evolved to support other styles.

**2.6.1.2 Other languages** Scala [23,26] is a hybrid object-functional language that enables function definitions to be nested inside other function definitions. The procedural Scala program `CountingProc.scala` is equivalent to the Python program above.

## 2.6.2 Modular paradigm

*Modular programming* refers more to a design method for programs and program libraries than to languages.

Modular programming means to decompose a program into units of functionality (i.e., modules) that can be developed separately and then recomposed. These *modules* can hide (i.e., encapsulate) key design and implementation details within the modu

The module's *public* features can be accessed through its interface; its *private* features cannot be accessed from outside the module. Thus a module supports the principle of *information hiding*. This method also keeps the interactions among modules at a minimum, maintaining a low degree of coupling.

We discuss modular programming in more depth in Chapters 6 and 7.

A language that provides constructs for defining modules, packages, namespaces, or separate compilation units can assist in writing modular programs.

In this chapter, we examine some aspects of the modular paradigm using the imperative language Python. We examine modular programming in the purely functional programming language Haskell on Chapters 6 and 7 and later chapters.

### 2.6.2.1 Python

**2.6.2.1.1 Using one module** First, let's consider the following Python [25] code from file `CountingMod.py` to illustrate use of modules in Python programs. This module is similar to the procedural program in the previous section.

This modular program, however, has all the functions and procedures at the same level of the Python module (file) instead of most being nested within procedure `counter`. The modular program also uses module-level variables instead of local variables of procedure `counter`.

```
# File CountingMod.py
count = 0
maxc = 10

def has_more():
    return count <= maxc

def adv():
    global count
    count = count + 1

def counter():
    while has_more():
        print(f'{count}')
        adv()
```

This module creates two module-level *global* variables `count` and `maxc` and defines three module-level Python functions `has_more`, `adv`, and `counter`.

The module assigns initial values to the variables. Their values can be accessed anywhere later in the module unless hidden by parameters or local variables with the same name.

Function `has_more()` tests module-level variables `count` and `maxc` to determine whether there are more items in the sequence.

Procedure `adv()` assigns a new value to the module-level variable `count`. It must declare `count` as `global` so that a new local variable is not created.

Variable `maxc` is also mutable, but this module does not modify its value.

Each module is a separate file that can be imported by other Python code. It introduces a separate name space for variables, functions, and other features.

For example, we can import the module above and execute `counter` with the following Python code from file `CountingModTest1.py`:

```
from CountingMod import counter
counter()
```

The `from-import` statement imports feature `counter` (a Python function) from the module in file `CountingMod.py`. The imported name `counter` can be used without qualifying it. The other features of `CountingMod` (e.g., `count` and `adv`) cannot be accessed.

As an alternative, we can import the module from file `CountingModTest2.py` as follows:

```
import CountingMod

CountingMod.count = 10
CountingMod.maxc = 20
CountingMod.counter()
```

This code imports all the features of the module. It requires the variables and functions to be accessed with the name prefix `CountingMod.` (i.e., the module name followed by a period). This approach enables the importing code to modify the values of global variables in the imported module.

In this second example, the importing code can both access and modify the global variables of the imported module.

Python does not enforce the encapsulation of module-level variable or function names. All names are public (i.e., can be imported to other modules). However, programmers can, by convention, designate module-level names as private by beginning the name with a single underscore character `_`. The alternative `import` above will not automatically import such names.

For example, good modular programming practice might suggest that the names `_count`, `_maxc`, `_has_more()`, and `_adv()` be used in the `CountingMod` module above. This naming convention would designate those as private and leave only `counter()` as public.

Most modern languages support “modules” in some way. Other languages (e.g., Standard ML) provide advanced support for modules with the ability to encapsulate features and provide multiple implementations of common interfaces.

**2.6.2.1.2 Using multiple modules** To see the flexibility of modular programs, let’s consider a variant of the above that uses two modules.

The first module—`CountingModA` from file `CountingModA.py`—is shown below.

```
# File CountingModA.py
from Arith import reset, adv, get_count, has_more

def counter():
    while has_more():
        count = get_count()
        print(f'{count}')
        adv()
```

`CountingModA` has similar overall functionality to the `CountingMod` module in the previous example. However, its `counter` procedure uses a `has_more` function, an `adv` procedure, and a new `get_counter` function implemented in a separate

module named `Arith`. The `CountingModA` module has no module-level variables and its `counter` procedure has no local variables.

The second module—`Arith` from file `Arith.py`—is shown below.

```
# File Arith.py
_start = 0
_stop = 10
_change = 1
_count = _start

def reset(new_start, new_stop, new_change):
    global _start, _stop, _change, _count
    _start = new_start
    _stop = new_stop
    _count = _start
    if new_change == 0:
        print('Error: Attempt to reset increment to 0; not reset.')
    else:
        _change = new_change

def adv():
    global _count
    _count = _count + _change

def get_count():
    return _count

def has_more():
    if _change > 0:
        return _count <= _stop
    else:
        return _count >= _stop
```

This module makes the module-level variables private to the module by convention.

By default, module `Arith` generates the same arithmetic sequence as `CountingMod` in the previous modular programming example. However, it generalizes `CountingMod` in the following ways:

- renaming variable `count` to be `_count` and variable `maxc` to be `_stop`
- replacing the constant `0` in the initialization of variable `_count` by a new variable `_start`, which is itself initialized to `0`
- replacing the constant `1` in the increment of variable `_count` by a new variable `_change`, which is itself initialized to `1`

- adding a new function `get_count` that enables a user module (e.g., `CountingModA`) to get the current value of the `_count` variable

This is called an *accessor* or *getter* function.

- implementing the function `has_more()` and the procedure `adv()` used by module `CountingModA`

These argumentless public functions operate on `Arith`'s private module-level variables `_start`, `_stop`, `_change`, and `_count`.

- adding a new procedure `reset()` that enables the values of `_start`, `_stop`, `_change`, and `_count` to be reinitialized to new values

Now let's consider an alternative to `Arith`, the second module. Module `Geom` from file `Geom.py` is shown below.

```
# File Geom.py
_start = 1
_stop = 100
_change = 2
_count = _start

def reset(new_start, new_stop, new_change):
    global _start, _stop, _change, _count
    _start = new_start
    _stop = new_stop
    _count = start
    if abs(new_change) <= 1:
        print('Error: Attempt to set abs(_change) <= 1; not reset.')
    else:
        _change = new_change

def adv():
    global _count
    _count = _count * _change

def get_count():
    return _count

def has_more():
    return _count <= _stop
```

Module `Geom` has essentially the same interface as `Arith`, but it generates a geometric sequence instead of an arithmetic sequence.

To use this module, the only change needed to `CountingModA.py` is to import the module `Geom` instead of `Arith`. This alternative is in module `CountingModG` in file `CountingModG.py`.



This two-level example illustrates the additional flexibility that modular programming can enable.

**2.6.2.2 Other languages** The modular Scala [23,26] program `CountingMod.scala` is equivalent to the first Python program above. The similar Scala program `CountingMod2.scala` uses a Scala *trait* to define the interface of the module. It is used in manner similar to the second Python program above.

TODO: Probably should show a Java 8+ example for this. Also the Scala might need more update to be similar to new modular Python examples.

### 2.6.3 Object-based paradigms

The dominant paradigm since the early 1990s has been the *object-oriented paradigm*. Because this paradigm is likely familiar with most readers, we examine it and related object-based paradigms in the next chapter.

### 2.6.4 Concurrent paradigms

TODO: Perhaps describe a paradigm like actors and give an example in Elixir [13,28].

## 2.7 Motivating Functional Programming: John Backus

In this book we focus primarily on the functional paradigm—on the programming language Haskell in particular. Although languages that enable or emphasize the functional paradigm have been around since the early days of computing, much of the later interest in functional programming grew from the 1977 Turing Award lecture.

John W. Backus (December 3, 1924 – March 17, 2007) was a pioneer in research and development of programming languages. He was the primary developer of Fortran while a programmer at IBM in the mid-1950s. Fortran is the first widely used high-level language. Backus was also a participant in the international team that designed the influential languages Algol 58 and Algol 60 a few years later. The notation used to describe the Algol 58 language syntax—Backus-Naur Form (BNF)—bears his name. This notation continues to be used to this day.

In 1977, ACM bestowed its Turing Award on Backus in recognition of his career of accomplishments. (This award is sometimes described as the “Nobel Prize for computer science”.) The annual recipient of the award gives an address to a major computer science conference. Backus’s address was titled “Can Programming Be Liberated from the von Neumann Style? A Functional Style and Its Algebra of Programs”.

Although functional languages like Lisp go back to the late 1950’s, Backus’s address did much to stimulate research community’s interest in functional programming languages and functional programming over the past four decades.

The next subsection gives excerpts from Backus's Turing Award address published as the article "Can Programming Be Liberated from the von Neumann Style? A Functional Style and Its Algebra of Programs" [1].

### 2.7.1 Excerpts from Backus's Turing Award Address [1]

Programming languages appear to be in trouble. Each successive language incorporates, with little cleaning up, all the features of its predecessors plus a few more. Some languages have manuals exceeding 500 pages; others cram a complex description into shorter manuals by using dense formalisms. . . . Each new language claims new and fashionable features, such as strong typing or structured control statements, but the plain fact is that few languages make programming sufficiently cheaper or more reliable to justify the cost of producing and learning to use them.

Since large increases in size bring only small increases in power, smaller, more elegant languages such as Pascal continue to be popular. But there is a desperate need for a powerful methodology to help us think about programs, and no conventional language even begins to meet that need. In fact, conventional languages create unnecessary confusion in the way we think about programs. . . . In order to understand the problems of conventional programming languages, we must first examine their intellectual parent, the von Neumann computer. What is a von Neumann computer? When von Neumann and others conceived of it . . . [in the 1940's], it was an elegant, practical, and unifying idea that simplified a number of engineering and programming problems that existed then. Although the conditions that produced its architecture have changed radically, we nevertheless still identify the notion of "computer" with this . . . concept.

In its simplest form a von Neumann computer has three parts: a central processing unit (or CPU), a store, and a connecting tube that can transmit a single word between the CPU and the store (and send an address to the store). I propose to call this tube the *von Neumann bottleneck*. The task of a program is to change the contents of the store in some major way; when one considers that this task must be accomplished entirely by pumping single words back and forth through the von Neumann bottleneck, the reason for its name becomes clear.

Ironically, a large part of the traffic in the bottleneck is not useful data but merely names of data, as well as operations and data used only to compute such names. Before a word can be sent through the tube its address must be in the CPU; hence it must either be sent through the tube from the store or be generated by some CPU operation. If the address is sent from the store, then its address must either have been sent from the store or generated in the CPU, and so on. If, on the other hand, the address is generated in the CPU, it must either be generated by a fixed rule (e.g., "add 1 to the program counter") or by an instruction that was sent through the tube, in which case its address must have been sent, and so on.

Surely there must be a less primitive way of making big changes in the store than

by pushing vast numbers of words back and forth through the von Neumann bottleneck. Not only is this tube a literal bottleneck for the data traffic of a problem, but, more importantly, it is an intellectual bottleneck that has kept us tied to word-at-a-time thinking instead of encouraging us to think in terms of the larger conceptual units of the task at hand. . . .

Conventional programming languages are basically high level, complex versions of the von Neumann computer. Our . . . old belief that there is only one kind of computer is the basis of our belief that there is only one kind of programming language, the conventional—von Neumann—language. The differences between Fortran and Algol 68, although considerable, are less significant than the fact that both are based on the programming style of the von Neumann computer. Although I refer to conventional languages as “von Neumann languages” to take note of their origin and style, I do not, of course, blame the great mathematician for their complexity. In fact, some might say that I bear some responsibility for that problem.

Von Neumann programming languages use variables to imitate the computer’s storage cells; control statements elaborate its jump and test instructions; and assignment statements imitate its fetching, storing, and arithmetic. The assignment statement is the von Neumann bottleneck of programming languages and keeps us thinking in word-at-a-time terms in much the same way the computer’s bottleneck does.

Consider a typical program; at its center are a number of assignment statements containing some subscripted variables. Each assignment statement produces a one-word result. The program must cause these statements to be executed many times, while altering subscript values, in order to make the desired overall change in the store, since it must be done one word at a time. The programmer is thus concerned with the flow of words through the assignment bottleneck as he designs the nest of control statements to cause the necessary repetitions.

Moreover, the assignment statement splits programming into two worlds. The first world comprises the right sides of assignment statements. This is an orderly world of expressions, a world that has useful algebraic properties (except that those properties are often destroyed by side effects). It is the world in which most useful computation takes place.

The second world of conventional programming languages is the world of statements. The primary statement in that world is the assignment statement itself. All the other statements in the language exist in order to make it possible to perform a computation that must be based on this primitive construct: the assignment statement.

This world of statements is a disorderly one, with few useful mathematical properties. Structured programming can be seen as a modest effort to introduce some order into this chaotic world, but it accomplishes little in attacking the fundamental problems created by the word-at-a-time von Neumann style of

programming, with its primitive use of loops, subscripts, and branching flow of control.

Our fixation on von Neumann languages has continued the primacy of the von Neumann computer, and our dependency on *it* has made non-von Neumann languages uneconomical and has limited their development. The absence of full scale, effective programming styles founded on non-von Neumann principles has deprived designers of an intellectual foundation for new computer architectures.

...

### 2.7.2 Aside on the disorderly world of statements

Backus states that “the world of statements is a disorderly one, with few mathematical properties”. Even in 1977 this was a bit overstated since work by Hoare on *axiomatic semantics* [16], by Dijkstra on the *weakest precondition (wp) calculus* [12], and by others had already appeared.

However, because of the referential transparency property of purely functional languages, reasoning can often be done in an equational manner within the context of the language itself. We examine this convenient approach later in this book.

In contrast, the *wp*-calculus and other axiomatic semantic approaches must project the problem from the world of programming language statements into the world of predicate calculus, which is much more orderly. We leave this study to courses on program derivation and programming language semantics.

Note: For this author’s take on this formal methods topic, see my materials for University of Mississippi course Program Semantics and Derivation (CSci 550) [7,8].

### 2.7.3 Perspective from four decades later

In his Turing Award Address, Backus went on to describe FP, his proposal for a functional programming language. He argued that languages like FP would allow programmers to break out of the von Neumann bottleneck and find new ways of thinking about programming.

FP itself did not catch on, but the widespread attention given to Backus’ address and paper stimulated new interest in functional programming to develop by researchers around the world. Modern languages like Haskell developed partly from the interest generated.

In the 21st Century, the software industry has become more interested in functional programming. Some functional programming features now appear in most mainstream programming languages (e.g., in Java 8+). This interest seems to be driven primarily by two concerns:

- managing the complexity of large software systems effectively

- exploiting multicore processors conveniently and safely

The functional programming paradigm is able to address these concerns because of such properties such as referential transparency, immutable data structures, and composability of components. We look at these aspects in later chapters.

## 2.8 What Next?

This chapter (2) introduced the concepts of abstraction and programming paradigm and surveyed the imperative, declarative, functional, and other paradigms.

Chapter 3 continues the discussion of programming paradigms by examining the object-oriented and related object-based paradigms.

The subsequent chapters use the functional programming language Haskell to illustrate general programming concepts and explore programming language design and implementation using interpreters.

## 2.9 Exercises

1. This chapter used Haskell (and Scala) to illustrate the functional paradigm. Choose a language such as Java, Python, or C#. Describe how it can be used to write programs in the functional paradigm. Consider how well the language supports tail recursion.

TODO: Modify question if more examples are given in chapter.

2. This chapter used Python (and Scala) to illustrate the procedural paradigm. Choose a different language such as Java, C, C++, or C#. Describe how it can be used to write programs in the procedural paradigm.

TODO: Modify question if more examples are given in chapter.

3. This chapter used Python (and Scala) to illustrate the modular paradigm. For the same language chosen for previous exercise, describe how it can be used to write programs in the modular paradigm.

TODO: Modify question if more examples are given in chapter.

4. Repeat the previous two exercises with a different language.

## 2.10 Acknowledgements

In Summer and Fall 2016, I adapted and revised much of this work from my previous materials:

- Abstraction (Section 2.2) from the “What is Abstraction?” section of my Data Abstraction notes [10], which I wrote originally for the first C++ (CSci 490) and Java-based (CSci 211) classes at UM in 1996 but expanded and adapted for other courses in later years. In the mid-to-late 1990s, the Data Abstraction notes drew on my study of a variety of sources (e.g.,

Bird and Wadler [2], Dale [11], Gries [15]; Horstmann [17,18], Liskov [20], Meyer [21], Mossenbock [22], Parnas [24], and Thomas [29])

- Discussion of the primary programming paradigms (Sections 2.3-2.6) from Chapter 1 of my *Notes on Functional Programming with Haskell* [9], which drew on the taxonomy in Hudak's survey paper [19]. In 2016, I expanded the discussion of the paradigms and included examples. This drew in part from my use and/or teaching of a variety of programming languages since my first programming course in 1974 (e.g., Fortran, Cobol, PL/I, C, Snobol, Jovial, Ada, Pascal, Haskell, C++, Java, Ruby, Scala, Lua, Elixir, and Python).
- Motivating Functional Programming (Section 2.7) from Chapter 1 of my *Notes on Functional Programming with Haskell* [9]. This includes a long excerpt from the influential Turing Award lecture by John Backus [1].

In 2017, I continued to develop this material as a part of Chapter 1, Fundamentals, of my 2017 Haskell-based programming languages textbook.

In Spring and Summer 2018, I reorganized and expanded the previous Fundamentals chapter into four chapters for the 2018 version of the textbook, now titled *Exploring Languages with Interpreters and Functional Programming*. These are Chapter 1, Evolution of Programming Languages; Chapter 2, Programming Paradigms (this chapter); Chapter 3, Object-based Paradigms; and Chapter 80, Review of Relevant Mathematics. I added the examples on procedural and modular programming.

I retired from the full-time faculty in May 2019. As one of my post-retirement projects, I am continuing work on this textbook. In January 2022, I began refining the existing content, integrating additional separately developed materials, reformatting the document (e.g., using CSS), constructing a bibliography (e.g., using citeproc), adding cross-references, and improving the build workflow and use of Pandoc.

In 2022, I also revised and expanded the modular programming example

I maintain this chapter as text in Pandoc's dialect of Markdown using embedded LaTeX markup for the mathematical formulas and then translate the document to HTML, PDF, and other forms as needed.

## 2.11 Terms and Concepts

TODO: Update

Abstraction, procedural abstraction, data abstraction, interface, procedures, functions, methods; programming language paradigm, primary paradigms (imperative, declarative, functional, relational or logic language); other paradigms (procedural, modular, object-oriented, concurrent); program state, implicit versus explicit state, execution of commands versus evaluation of expressions, mutable versus immutable data structures, side effects, sequencing, recursion, referential

transparency, first-class values, first-order and higher-order functions, lexical scope, global versus local variables, public versus private features, information hiding, encapsulation, lexical closure; von Neumann computer, von Neumann language, worlds of expressions and statements, axiomatic semantics, weakest precondition calculus.

## 2.12 References

- [1] John Backus. 1978. Can programming be liberated from the von Neumann style? A functional style and its algebra of programs (1977 Turing Award address). *Communications of the ACM* 21, 8 (1978), 613–641.
- [2] Richard Bird. 1998. *Introduction to functional programming using Haskell* (Second ed.). Prentice Hall, Englewood Cliffs, New Jersey, USA.
- [3] Timothy Budd. 1995. *Multiparadigm programming in Leda*. Addison-Wesley, Boston, Massachusetts, USA.
- [4] William E. Byrd. 2009. Relational programming in miniKanren: Techniques, applications, and implementations. PhD thesis. Indiana University, Bloomington, Indiana, USA.
- [5] William E. Byrd and Contributors. 2022. miniKanren.org. Retrieved from <http://minikanren.org/>
- [6] William F. Clocksin and Christopher S. Mellish. 2012. *Programming in Prolog: Using the ISO standard* (Fifth ed.). Springer, Berlin, Germany.
- [7] H. Conrad Cunningham. 2006. *A programmer’s introduction to predicate logic*. University of Mississippi, Department of Computer and Information Science, University, Mississippi, USA. Retrieved from [https://john.cs.olemiss.edu/~hcc/csci450/notes/haskell\\_notes.pdf](https://john.cs.olemiss.edu/~hcc/csci450/notes/haskell_notes.pdf)
- [8] H. Conrad Cunningham. 2006. *Notes on program semantics and derivation*. University of Mississippi, Department of Computer and Information Science, University, Mississippi, USA. Retrieved from <https://john.cs.olemiss.edu/~hcc/reports/umcis-1994-02.pdf>
- [9] H. Conrad Cunningham. 2014. *Notes on functional programming with Haskell*. University of Mississippi, Department of Computer and Information Science, University, Mississippi, USA. Retrieved from [https://john.cs.olemiss.edu/~hcc/docs/Notes\\_FP\\_Haskell/Notes\\_on\\_Functional\\_Programming\\_with\\_Haskell.pdf](https://john.cs.olemiss.edu/~hcc/docs/Notes_FP_Haskell/Notes_on_Functional_Programming_with_Haskell.pdf)
- [10] H. Conrad Cunningham. 2019. *Notes on data abstraction*. University of Mississippi, Department of Computer and Information Science, University, Mississippi, USA. Retrieved from <https://john.cs.olemiss.edu/~hcc/docs/DataAbstraction/DataAbstraction.html>
- [11] Nell Dale and Henry M. Walker. 1996. *Abstract data types: Specifications, implementations, and applications*. D. C. Heath, Lexington, Massachusetts, USA.

- [12] Edsger W. Dijkstra. 1975. Guarded commands, nondeterminacy and formal derivation of programs. *Communications of the ACM* 18, 8 (1975), 453–457.
- [13] Elixir Team. 2022. Elixir. Retrieved from <https://elixir-lang.org>
- [14] Daniel P. Friedman, William E. Byrd, Oleg Kiselyov, and Jason Hemenn. 2018. *The reasoned schemer* (Second ed.). MIT Press, Cambridge, Massachusetts, USA.
- [15] David Gries. 1981. *Science of programming*. Springer, New York, New York, USA.
- [16] Charles Antony Richard Hoare. 1969. An axiomatic basis for computer programming. *Communications of the ACM* 12, 10 (1969), 576–580.
- [17] Cay S. Horstmann. 1995. *Mastering object-oriented design in C++*. Wiley, Indianapolis, Indiana, USA.
- [18] Cay S. Horstmann and Gary Cornell. 1999. *Core Java 1.2: Volume I—Fundamentals*. Prentice Hall, Englewood Cliffs, New Jersey, USA.
- [19] Paul Hudak. 1989. Conception, evolution, and application of functional programming languages. *ACM Computing Surveys* 21, 3 (1989), 359–411.
- [20] Barbara Liskov. 1987. Keynote address—Data abstraction and hierarchy. In *Proceedings on object-oriented programming systems, languages, and applications (OOPSLA '87): addendum*, ACM, Orlando, Florida, USA, 17–34.
- [21] Bertrand Meyer. 1997. *Object-oriented program construction* (Second ed.). Prentice Hall, Englewood Cliffs, New Jersey, USA.
- [22] Hanspeter Mossenbock. 1995. *Object-oriented programming in Oberon-2*. Springer, Berlin, Germany.
- [23] Martin Odersky, Lex Spoon, and Bill Venners. 2021. *Programming in Scala* (Fifth ed.). Artima, Inc., Walnut Creek, California, USA.
- [24] David L. Parnas. 1972. On the criteria to be used in decomposing systems into modules. *Communications of the ACM* 15, 12 (December 1972), 1053–1058.
- [25] Python Software Foundation. 2022. Python. Retrieved from <https://www.python.org/>
- [26] Scala Language Organization. 2022. The Scala programming language. Retrieved from <https://www.scala-lang.org/>
- [27] SWI-Prolog Organization. 2022. SWI-Prolog: Robust, mature, free Prolog for the real world. Retrieved from <https://www.swi-prolog.org/>
- [28] Dave Thomas. 2018. *Programming Elixir >= 1.6: Functional > concurrent > pragmatic > fun*. Pragmatic Bookshelf, Raleigh, North Carolina, USA.
- [29] Pete Thomas and Ray Weedon. 1995. *Object-oriented programming in Eiffel*. Addison-Wesley, Boston, Massachusetts, USA.