

ACM Computing Seminar Fortran Guide

Joseph P. McKenna

October 10, 2016

Contents

1	Introduction	2
1.1	About the language	3
2	Getting started	3
2.1	Text editor	3
2.2	Compiler	3
2.3	Writing and compiling a program	4
2.3.1	Hello world	4
2.3.2	Template	4
2.4	Exercises	6
3	Data types	6
3.1	The <code>logical</code> type	7
3.2	The <code>integer</code> type	8
3.3	Floating point types	8
3.4	The <code>character</code> type	9
3.5	Casting	9
3.6	The <code>parameter</code> keyword	9
3.7	Setting the precision	10
3.8	Pointers	10
3.9	Arrays	11
3.9.1	Fixed-length arrays	11
3.9.2	Dynamic length arrays	13
4	Control structures	14
4.1	Conditionals	14
4.1.1	The <code>if</code> construct	14
4.1.2	Example: <code>if / else</code> and random number generation	15

4.1.3	Example: <code>if / else if / else</code>	16
4.2	Loops	16
4.2.1	The <code>do</code> loop	16
4.2.2	The <code>do while</code> loop	19
4.2.3	Example: the <code>exit</code> keyword	21
5	Input/Output	22
5.1	File input/output	22
5.1.1	Reading data from file	22
5.1.2	Writing data to file	23
5.2	Formatted input/output	23
5.3	Command line arguments	24
6	Functions/Subroutines	25
6.1	Writing a function	25
6.1.1	Example: <code>linspace</code> : generating a set of equally-space points	26
6.2	Writing a subroutine	27
6.2.1	Example: polar coordinates	27
6.3	Passing procedures as arguments	28
6.3.1	Example: Newton's method for rootfinding	28
6.3.2	Example: The midpoint rule for definite integrals	31
6.4	Polymorphism	32
6.4.1	Example: machine epsilon	33
6.5	Recursion	34
6.5.1	Example: factorial	34
7	Object-oriented programming	35
7.1	Derived types	35
7.2	Modules	36
7.3	Example: determinant of random matrix	37
7.4	Example: matrix module	41

1 Introduction

This guide is intended to quickly get you up-and-running in scientific computing with Fortran.

1.1 About the language

Fortran was created in the 1950s for mathematical **FOR**-mula **TRAN**-slation, and has since gone through a number of revisions (FORTRAN 66, 77, and Fortran 90, 95, 2003, 2008, 2015). The language standards are put forth by the Fortran standards committee J3 in a document (ISO 1539-1:2010) available for purchase. The language syntax and intrinsic procedures make it especially suited for scientific computing. Fortran is a **statically-typed** and **compiled** language, like C++. You must declare the **type**, i.e. integer, real number, etc. of variables in programs you write. Your programs will be translated from human-readable *source code* into an executable file by software called a **compiler**. Fortran is **not case-sensitive**, so `matrix` and `MaTriX` are translated to the same token by the compiler.

2 Getting started

The software that you need to get started comes prepackaged and ready to download on most Linux distributions. There are a few options for emulating a Linux environment in Windows or Mac OS, such as a virtual machine (VirtualBox) or package manager (MinGW or Cygwin on Windows and Brew on Mac OS).

2.1 Text editor

You will write the source code of your programs using a text editor. There are many options that have features designed for programming such as syntax highlighting and auto-completion. If you are an impossible-to-please perfectionist, you might want to check out Emacs. If you are easier to please, you might want to check out Sublime Text.

2.2 Compiler

To translate your source code into an executable, you will need a Fortran compiler. A free option is **gfortran**, part of the GNU compiler collection (gcc). The features of the Fortran language that are supported by the **gfortran** compiler are specified in the compiler manual. This is your most complete reference for the procedures intrinsic to Fortran that your programs can use. At the time of this writing, **gfortran** completely supports Fortran 95 and partially supports more recent standards.

2.3 Writing and compiling a program

A program is delimited by the `begin program / end program` keywords. A useful construct for keeping code that a program can use is called a **module**. A module is delimited by the `begin module / end module` keywords.

2.3.1 Hello world

Let's write a tiny program that prints "hello world" to the terminal screen in `hello.f90`.

```
1 program main
2   print*, 'hello world'
3 end program main
```

To compile the program, execute the following command on the command line in the same directory as `hello.f90`

```
gfortran hello.f90
```

This produces an executable file named `a.out` by default (On Windows, this is probably named `a.exe` by default). To run, execute the file.

```
./a.out
```

```
hello world
```

We could have specified a different name for the executable file during compilation with the `-o` option of `gfortran`.

```
gfortran hello.f90 -o my_executable_file
```

On Windows, you should append the `.exe` extension to `my_executable_file`.

2.3.2 Template

Now let's write an empty source code template for future projects. Our source code template will consist of two files in the same directory (`./source/`). In the following files, the contents of a line after a `!` symbol is a comment that is ignored by the compiler. One file `header.f90` contains a module that defines things to be used in the main program.

```
1 module header
2   implicit none
3   ! variable declarations and assignments
4 contains
5   ! function and subroutine definitions
6 end module header
```

This file should be compiled with the `-c` option of `gfortran`.

```
gfortran -c header.f90
```

This outputs the **object file** named `header.o` by default. An object file contains machine code that can be *linked* to an executable. A separate file `main.f90` contains the main program.

```
1 program main
2   use header
3   implicit none
4   ! variable declarations and assignments
5   ! function and subroutine calls
6 contains
7   ! function and subroutine definitions
8 end program main
```

On line 2 of `main.f90`, we instruct the main program to use the contents of `header.f90`, so we must link `header.o` when compiling `main.f90`.

```
gfortran main.f90 header.o -o main
```

To run the program, execute the output file `main`.

```
./main
```

As you get more experience, you may find it cumbersome to repeatedly execute `gfortran` commands with every modification to your code. A way around this is to use the `make` command-line utility. Using `make`, all the of the compilation commands for your project can be coded in a file named `makefile` in the same directory as your `.f90` source files. For example, the template above could use the following `makefile`.

```

1  COMPILER = gfortran
2  SOURCE = main.f90
3  EXECUTABLE = main
4  OBJECTS = header.o
5
6  all: $(EXECUTABLE)
7  $(EXECUTABLE): $(OBJECTS)
8      $(COMPILER) $(SOURCE) $(OBJECTS) -o $(EXECUTABLE)
9  %.o: %.f90
10     $(COMPILER) -c $<

```

Then, to recompile both `header.f90` and `main.f90` after modifying either file, execute

```
make
```

in the same directory as `makefile`. The first four lines of the `makefile` above define the compiler command, file name of the main program, file name of the executable to be created, and file name(s) of linked object file(s), respectively. If you wrote a second module in a separate file `my_second_header.f90` that you wanted to use in `main.f90`, you would modify line 4 of `makefile` to `OBJ = header.o my_second_header.o`. The remaining lines of `makefile` define instructions for compilation.

2.4 Exercises

1. Compile and run `hello.f90`.
2. Execute `man gfortran` in any directory to bring up the manual for `gfortran`. Read the description and skim through the options. Do the same for `make`.

3 Data types

In both programs and modules, variables are declared first before other procedures. A variable is declared by listing its data type followed by `::` and the variable name, i.e. `integer :: i` or `real :: x`.

We will use the `implicit none` keyword at the beginning of each program and module as in line 2 of `header.f90` and line 3 of `main.f90` in Section 2.3.2. The role of this keyword is to suppress implicit rules for interpreting undeclared variables. By including it, we force ourselves to declare each

variable we use, which should facilitate debugging when our program fails to compile. Without it, an undeclared variable with a name such as `i` is assumed to be of the `integer` data type whereas an undeclared variable with a name such as `x` is assumed to be of the `real` data type.

In addition to the most common data types presented below, Fortran has a `complex` data type and support for data types defined by the programmer (see Section 7.1).

3.1 The logical type

A logical data type can have values `.true.` or `.false.`. Logical expressions can be expressed by combining unary or binary operations.

```
1 logical :: a,b,c
2 a = .true.
3 b = .false.
4
5 ! '.not.' is the logical negation operator
6 c = .not.a ! c is false
7
8 ! '.and,' is the logical and operator
9 c = a.and.b ! c is false
10
11 ! '.or.' is the logical or operator
12 c = a.or.b ! c is true
13
14 ! '==' is the test for equality
15 c = 1 == 2 ! c is false
16
17 ! '/=' is test for inequality
18 c = 1 /= 2 ! c is true
19 print*, c
```

Other logical operators include

- `<` or `.lt.:` less than
- `<=` or `.le.:` less than or equal
- `>` or `.gt.:` greater than
- `>=` or `.ge.:` greater than or equal

Logical expressions are often used in control structures.

3.2 The integer type

An integer data type can have integer values. If a real value is assigned to an integer type, the decimal portion is chopped off.

```
1 integer :: a = 6, b = 7 ! initialize a and b to 6 and 7, resp
2 integer :: c
3
4 c = a + b ! c is 13
5 c = a - b ! c is -1
6 c = a / b ! c is 0
7 c = b / a ! c is 1
8 c = a*b ! c is 42
9 c = a**b ! c is 6^7
10 c = mod(b,a) ! c is (b mod a) = 1
11 c = a > b ! c is 0 (logical gets cast to integer)
12 c = a < b ! c is 1 (logical gets cast to integer)
```

3.3 Floating point types

The two floating point data types `real` and `double precision` correspond to IEEE 32- and 64-bit floating point data types. A constant called *machine epsilon* is the least positive number in a floating point system that when added to 1 results in a floating point number larger than 1. It is common in numerical analysis error estimates.

```
1 real :: a ! declare a single precision float
2 double precision :: b ! declare a double precision float
3
4 ! Print the min/max value and machine epsilon
5 ! for the single precision floating point system
6 print*, tiny(a), huge(a), epsilon(a)
7
8 ! Print the min/max value and machine epsilon
9 ! for the double precision floating point system
10 print*, tiny(b), huge(b), epsilon(b)
```

1.17549435E-38 3.40282347E+38 1.19209290E-07
2.2250738585072014E-308 1.7976931348623157E+308 2.2204460492503131E-016

3.4 The character type

A `character` data type can have character values, i.e. letters or symbols. A character string is declared with a positive `integer` specifying its maximum possible length.

```
1 ! declare a character variable s at most 32 characters
2 character(32) :: s
3
4 ! assign value to s
5 s = 'file_name'
6
7 ! trim trailing spaces from s and
8 ! append a character literal '.txt'
9 print*, trim(s) // '.txt'

file_name.txt
```

3.5 Casting

An `integer` can be cast to a `real` and vice versa.

```
1 integer :: a = 1, b
2 real :: c, PI = 3.14159
3
4 ! explicit cast real to integer
5 b = int(PI) ! b is 3
6
7 ! explicit cast integer to real then divide
8 c = a/real(b) ! c is .3333...
9
10 ! divide then implicit cast real to integer
11 c = a/b ! c is 0
```

3.6 The parameter keyword

The `parameter` keyword is used to declare constants. A constant must be assigned a value at declaration and cannot be reassigned a value. The following code is not valid because of an attempt to reassign a constant.

```
1 ! declare constant variable
2 real, parameter :: PI = 2.*asin(1.) ! 'asin' is arcsine
```

```
3
4 PI = 3 ! not valid
```

The compiler produces an error like `Error: Named constant 'pi' in variable definition context (assignment)`.

3.7 Setting the precision

The `kind` function returns an `integer` for each data type. The precision of a floating point number can be specified at declaration by a literal or constant `integer` of the desired kind.

```
1 ! declare a single precision
2 real :: r
3 ! declare a double precision
4 double precision :: d
5 ! store single precision and double precision kinds
6 integer, parameter :: sp = kind(r), dp = kind(d)
7 ! set current kind
8 integer, parameter :: rp = sp
9
10 ! declare real b in double precision
11 real(dp) :: b
12
13 ! declare real a with precision kind rp
14 real(rp) :: a
15
16 ! cast 1 to real with precision kind rp and assign to a
17 a = 1.0_rp
18
19 ! cast b to real with precision kind rp and assign to a
20 a = real(b,rp)
```

To switch the precision of each variable above with kind `rp`, we would only need to modify the declaration of `rp` on line 8.

3.8 Pointers

Pointers have the same meaning in Fortran as in C++. A pointer is a variable that holds the **memory address** of a variable. The implementation of pointers is qualitatively different in Fortran than in C++. In Fortran,

the user cannot view the memory address that a pointer stores. A pointer variable is declared with the `pointer` modifier, and a variable that it points to is declared with the `target` modifier. The types of a `pointer` and its `target` must match.

```
1 ! declare pointer
2 integer, pointer :: p
3 ! declare targets
4 integer, target :: a = 1, b = 2
5
6 p => a ! p has same memory address as a
7 p = 2 ! modify value at address
8 print*, a==2 ! a is 2
9
10 p => b ! p has same memory address as b
11 p = 1 ! modify value at address
12 print*, b==1 ! b is 1
13
14 ! is p associated with a target?
15 print*, associated(p)
16
17 ! is p associated with the target a?
18 print*, associated(p, a)
19
20 ! point to nowhere
21 nullify(p)

T
T
T
F
```

3.9 Arrays

The length of an array can be fixed or dynamic. The index of an array starts at 1 by default, but any index range can be specified.

3.9.1 Fixed-length arrays

An array can be declared with a single `integer` specifying its length in which case the first index of the array is 1. An array can also be declared with an

integer range specifying its first and last index.

Here's a one-dimensional array example.

```
1 ! declare array of length 5
2 ! index range is 1 to 5 (inclusive)
3 real :: a(5)
4
5 ! you can work with each component individually
6 ! set the first component to 1
7 a(1) = 1.0
8
9 ! or you can work with the whole array
10 ! set the whole array to 2
11 a = 2.0
12
13 ! or you can work with slices of the array
14 ! set elements 2 to 4 (inclusive) to 3
15 a(2:4) = 3.0
```

And, here's a two-dimensional array example.

```
1 ! declare 5x5 array
2 ! index range is 1 to 5 (inclusive) in both axes
3 real :: a(5,5)
4
5 ! you can work with each component individually
6 ! set upper left component to 1
7 a(1,1) = 1.0
8
9 ! or you can work with the whole array
10 ! set the whole array to 2
11 a = 2.0
12
13 ! or you can work with slices of the array
14 ! set a submatrix to 3
15 a(2:4, 1:2) = 3.0
```

Fortran includes intrinsic functions to operate on an array `a` such as

- `size(a)`: number of elements of `a`
- `minval(a)`: minimum value of `a`

- `maxval(a)`: maximum value of `a`
- `sum(a)`: sum of elements in `a`
- `product(a)`: product of elements in `a`

See the `gfortran` documentation for more.

3.9.2 Dynamic length arrays

Dynamic arrays are declared with the `allocatable` modifier. Before storing values in such an array, you must `allocate` memory for the array. After you are finished the array, you ought to `deallocate` the memory that it occupies.

Here's a one-dimensional array example.

```

1  ! declare a one-dim. dynamic length array
2  real, allocatable :: a(:)
3
4  ! allocate memory for a
5  allocate(a(5))
6
7  ! now you can treat a like a normal array
8  a(1) = 1.0
9  ! etc...
10
11 ! deallocate memory occupied by a
12 deallocate(a)
13
14 ! we can change the size and index range of a
15 allocate(a(0:10))
16
17 a(0) = 1.0
18 ! etc...
19
20 deallocate(a)

```

Without the last `deallocate` statement on line 20 the code above is valid, but the memory that is allocated for `a` will not be freed. That memory then cannot be allocated to other resources.

Here's a two-dimensional array example.

```

1  ! declare a two-dim. dynamic length array

```

```

2  real, allocatable :: a(:, :)
3
4  ! allocate memory for a
5  allocate(a(5,5))
6
7  ! now you can treat a like a normal array
8  a(1,1) = 1.0
9  ! etc...
10
11 ! deallocate memory occupied by a
12 deallocate(a)
13
14 ! we can change the size and index range of a
15 allocate(a(0:10,0:10))
16
17 a(0,0) = 1.0
18 ! etc...
19
20 deallocate(a)

```

4 Control structures

Control structures are used to direct the flow of code execution.

4.1 Conditionals

4.1.1 The if construct

The `if` construct controls execution of a single block of code. If the block of code is more than one line, it should be delimited by an `if / end if` pair. If the block of code is one line, it can be written on one line. A common typo is to forget the `then` keyword following the logical in an `if / end if` pair.

```

1  real :: num = 0.75
2
3  if (num < .5) then
4      print*, 'num: ', num
5      print*, 'num is less than 0.5'
6  end if
7
8  if (num > .5) print*, 'num is greater than 0.5'

```

num is greater than 0.5

4.1.2 Example: if / else and random number generation

The if / else construct controls with mutually exclusive logic the execution of two blocks of code.

The following code generates a random number between 0 and 1, then prints the number and whether or not the number is greater than 0.5

```
1 real :: num
2
3 ! seed random number generator
4 call srand(789)
5
6 ! rand() returns a random number between 0 and 1
7 num = rand()
8
9 print*, 'num: ', num
10
11 if (num < 0.5) then
12     print*, 'num is less than 0.5'
13 else
14     print*, 'num is greater than 0.5'
15 end if
16
17 ! do it again
18 num = rand()
19
20 print*, 'num: ', num
21
22 if (num < 0.5) then
23     print*, 'num is less than 0.5'
24 else
25     print*, 'num is greater than 0.5'
26 end if
```

```
num:    6.17480278E-03
num is less than 0.5
num:    0.783314705
num is greater than 0.5
```

Since the random number generator was seeded with a literal integer, the above code will produce the *same* output each time it is run.

4.1.3 Example: if / else if / else

The if / else if / else construct controls with mutually exclusive logic the execution of three or more blocks of code. The following code generates a random number between 0 and 1, then prints the number and which quarter of the interval [0, 1] that the number is in.

```
1 real :: num
2
3 ! seed random number generator with current time
4 call srand(time())
5
6 ! rand() returns a random number between 0 and 1
7 num = rand()
8
9 print*, 'num:', num
10
11 if (num > 0.75) then
12     print*, 'num is between 0.75 and 1'
13 else if (num > 0.5) then
14     print*, 'num is between 0.5 and 0.75'
15 else if (num > 0.25) then
16     print*, 'num is between 0.25 and 0.5'
17 else
18     print*, 'num is between 0 and 0.25'
19 end if

num: 0.679201365
num is between 0.5 and 0.75
```

Since the random number generator was seeded with the current time, the above code will produce a *different* output each time it is run.

4.2 Loops

4.2.1 The do loop

A do loop iterates a block of code over a range of integers. It takes two integer arguments specifying the minimum and maximum (inclusive) of the

range and takes an optional third `integer` argument specifying the iteration stride in the form `do i=min,max,stride`. If omitted, the stride is 1.

The following code assigns a value to each component of an array then prints it.

```
1 integer :: max = 10, i
2 real, allocatable :: x(:)
3
4 allocate(x(0:max))
5
6 do i = 0,max
7     ! assign to each array component
8     x(i) = i / real(max)
9
10    ! print current component
11    print "('x(', i0, ') = ', f3.1)", i, x(i)
12 end do
13
14 deallocate(x)

x(0) = 0.0
x(1) = 0.1
x(2) = 0.2
x(3) = 0.3
x(4) = 0.4
x(5) = 0.5
x(6) = 0.6
x(7) = 0.7
x(8) = 0.8
x(9) = 0.9
x(10) = 1.0
```

An *implicit do loop* can be used for formulaic array assignments. The following code creates the same array as the last example.

```
1 integer :: max = 10
2 real, allocatable :: x(:)
3
4 allocate(x(0:max))
5
```

```

6 ! implicit do loop for formulaic array assignment
7 x = [(i / real(max), i=0, max)]
8
9 deallocate(x)

```

Example: row-major matrix The following code stores matrix data in a one-dimensional array named `matrix` in **row-major** order. This means the first `n_cols` elements of the array will contain the first row of the matrix, the next `n_cols` of the array will contain the second row of the matrix, etc.

```

1 integer :: n_rows = 4, n_cols = 3
2 real, allocatable :: matrix(:)
3 ! temporary indices
4 integer :: i,j,k
5
6 ! index range is 1 to 12 (inclusive)
7 allocate(matrix(1:n_rows*n_cols))
8
9 ! assign 0 to all elements of matrix
10 matrix = 0.0
11
12 do i = 1,n_rows
13     do j = 1,n_cols
14         ! convert (i,j) matrix index to "flat" row-major index
15         k = (i-1)*n_cols + j
16
17         ! assign 1 to diagonal, 2 to sub/super-diagonal
18         if (i==j) then
19             matrix(k) = 1.0
20         else if ((i==j-1).or.(i==j+1)) then
21             matrix(k) = 2.0
22         end if
23     end do
24 end do
25
26 ! print matrix row by row
27 do i = 1,n_rows
28     print "(3(f5.1))", matrix(1+(i-1)*n_cols:i*n_cols)
29 end do
30

```

```
31 deallocate(matrix)
```

```
1.0  2.0  0.0
2.0  1.0  2.0
0.0  2.0  1.0
0.0  0.0  2.0
```

4.2.2 The do while loop

A `do while` loop iterates while a logical condition evaluates to `.true..`

Example: truncated sum The following code approximates the geometric series

$$\sum_{n=1}^{\infty} \left(\frac{1}{2}\right)^n = 1.$$

The `do while` loop begins with $n = 1$ and exits when the current summand does not increase the current sum. It prints the iteration number, current sum, and absolute error

$$E = 1 - \sum_{n=1}^{\infty} \left(\frac{1}{2}\right)^n .$$

```
1 real :: sum = 0.0, base = 0.5, tol = 1e-4
2 real :: pow = 0.5
3 integer :: iter = 1
4
5 do while (sum+pow > sum)
6     ! add pow to sum
7     sum = sum+pow
8     ! update pow by one power of base
9     pow = pow*base
10
11     print "('Iter: ', i3, ', Sum: ', f0.10, ', Abs Err: ', f0.10)", iter, sum, 1-sum
12
13     ! update iter by 1
14     iter = iter+1
15 end do
```

```
Iter:   1, Sum: .5000000000, Abs Err: .5000000000
Iter:   2, Sum: .7500000000, Abs Err: .2500000000
```

```

Iter: 3, Sum: .875000000, Abs Err: .125000000
Iter: 4, Sum: .937500000, Abs Err: .062500000
Iter: 5, Sum: .968750000, Abs Err: .031250000
Iter: 6, Sum: .984375000, Abs Err: .015625000
Iter: 7, Sum: .992187500, Abs Err: .007812500
Iter: 8, Sum: .996093750, Abs Err: .003906250
Iter: 9, Sum: .998046875, Abs Err: .001953125
Iter: 10, Sum: .9990234375, Abs Err: .0009765625
Iter: 11, Sum: .9995117188, Abs Err: .0004882812
Iter: 12, Sum: .9997558594, Abs Err: .0002441406
Iter: 13, Sum: .9998779297, Abs Err: .0001220703
Iter: 14, Sum: .9999389648, Abs Err: .0000610352
Iter: 15, Sum: .9999694824, Abs Err: .0000305176
Iter: 16, Sum: .9999847412, Abs Err: .0000152588
Iter: 17, Sum: .9999923706, Abs Err: .0000076294
Iter: 18, Sum: .9999961853, Abs Err: .0000038147
Iter: 19, Sum: .9999980927, Abs Err: .0000019073
Iter: 20, Sum: .9999990463, Abs Err: .0000009537
Iter: 21, Sum: .9999995232, Abs Err: .0000004768
Iter: 22, Sum: .9999997616, Abs Err: .0000002384
Iter: 23, Sum: .9999998808, Abs Err: .0000001192
Iter: 24, Sum: .9999999404, Abs Err: .0000000596
Iter: 25, Sum: 1.000000000, Abs Err: .000000000

```

Example: estimating machine epsilon The following code finds machine epsilon by shifting the rightmost bit of a binary number rightward until it falls off. Think about how it does this. Could you write an algorithm that finds machine epsilon using the function `rshift` that shifts the bits of float rightward?

```

1 double precision :: eps
2 integer, parameter :: dp = kind(eps)
3 integer :: count = 1
4
5 eps = 1.0_dp
6 do while (1.0_dp + eps*0.5 > 1.0_dp)
7     eps = eps*0.5
8     count = count+1
9 end do
10

```

```

11 print*, eps, epsilon(eps)
12 print*, count, digits(eps)

2.2204460492503131E-016   2.2204460492503131E-016
      53                   53

```

4.2.3 Example: the exit keyword

The `exit` keyword stops execution of code within the current scope.

The following code finds the *hailstone sequence* of $a_1 = 6$ defined recursively by

$$a_{n+1} = \begin{cases} a_n/2 & \text{if } a_n \text{ is even} \\ 3a_n + 1 & \text{if } a_n \text{ is odd} \end{cases}$$

for $n \geq 1$. It is an open conjecture that the hailstone sequence of any initial value a_1 converges to the periodic sequence $4, 2, 1, 4, 2, 1, \dots$. Luckily, it does for $a_1 = 6$ and the following infinite do loop exits.

```

1 integer :: a = 6, count = 1
2
3 ! infinite loop
4 do
5     ! if a is even, divide by 2
6     ! otherwise multiply by 3 and add 1
7     if (mod(a,2)==0) then
8         a = a/2
9     else
10        a = 3*a+1
11    end if
12
13    ! if a is 4, exit infinite loop
14    if (a==4) then
15        exit
16    end if
17
18    ! print count and a
19    print "('count: ', i2, ', a: ', i2)", count, a
20
21    ! increment count
22    count = count + 1
23 end do

```

```
count: 1, a: 3
count: 2, a: 10
count: 3, a: 5
count: 4, a: 16
count: 5, a: 8
```

5 Input/Output

5.1 File input/output

5.1.1 Reading data from file

The contents of a data file can be read into an array using `read`. Suppose you have a file `./data/array.txt` that contains two columns of data

```
1 1.23
2 2.34
3 3.45
4 4.56
5 5.67
```

This file can be opened with the `open` command. The required first argument of `open` is an integer that specifies a *file unit* for `array.txt`. Choose any number that is not in use. The unit numbers 0, 5, and 6 are reserved for system files and should not be used accidentally. Data are read in **row-major** format, i.e. across the first row, then across the second row, etc.

The following code reads the contents of `./data/array.txt` into an array called `array`.

```
1 ! declare array
2 real :: array(5,2)
3 integer :: row
4
5 ! open file and assign file unit 10
6 open (10, file='./data/array.txt', action='read')
7
8 ! read data from file unit 10 into array
9 do row = 1,5
10     read(10,*) array(row,:)
11 end do
```

```

12
13 ! close file
14 close(10)

```

5.1.2 Writing data to file

Data can be written to a file with the `write` command.

```

1 real :: x
2 integer :: i, max = 5
3
4 ! open file, specify unit 10, overwrite if exists
5 open(10, file='./data/sine.txt', action='write', status='replace')
6
7 do i = 0,max
8     x = i / real(max)
9
10     ! write to file unit 10
11     write(10,*) x, sin(x)
12 end do

```

This produces a file `sine.txt` in the directory `data` containing

```

0.00000000    0.00000000
0.200000003    0.198669329
0.400000006    0.389418334
0.600000024    0.564642489
0.800000012    0.717356086
1.00000000    0.841470957

```

5.2 Formatted input/output

The format of a `print`, `write`, or `read` statement can be specified with a `character` string. A format character string replaces the `*` symbol in `print*` and the second `*` symbol in `read(*,*)` or `write(*,*)`. A format string is a list of literal character strings or character descriptors from

- `a`: character string
- `iW`: integer
- `fW.D`: float point

- `esW.DeE`: scientific notation
- `Wx`: space

where `W`, `D`, and `E` should be replaced by numbers specifying width, number of digits, or number of exponent digits, resp. The width of a formatted integer or float defaults to the width of the number when `W` is 0.

```

1  character(32) :: fmt, a = 'word'
2  integer :: b = 1
3  real :: c = 2.0, d = 3.0
4
5  ! character string and 4 space-delimited values
6  print "('four values: ', a, 1x i0, 1x f0.1, 1x, es6.1e1)", trim(a), b, c, d
7
8  ! character string and 2 space-delimited values
9  fmt = '(a, 2(f0.1, 1x))'
10 print fmt, 'two values: ', c, d

```

```

four values: word 1 2.0 3.0E+0
two values: 2.0 3.0

```

5.3 Command line arguments

Arguments can be passed to a program from the command line using `get_command_argument`. The first argument received by `get_command_argument` is the program executable file name and the remaining arguments are passed by the user. The following program accepts any number of arguments, each at most 32 characters, and prints them.

```

1  program main
2    implicit none
3
4    character(32) :: arg
5    integer :: n_arg = 0
6
7    do
8      ! get next command line argument
9      call get_command_argument(n_arg, arg)
10
11     ! if it is empty, exit

```



```

12     if (len_trim(arg) == 0) exit
13
14     ! print argument to screen
15     print('argument ', i0, ': ', a)", n_arg, trim(arg)
16
17     ! increment count
18     n_arg = n_arg+1
19 end do
20
21 ! print total number of arguments
22 print "('number of arguments: ', i0)", n_arg
23
24 end program main

```

After compiling to `a.out`, you can pass arguments in the executing command.

```

./a.out 1 2 34

argument 0: ./a.out
argument 1: 1
argument 2: 2
argument 3: 34
number of arguments: 4

```

6 Functions/Subroutines

Functions and subroutines are callable blocks of code. A **function** returns a value from a set of arguments. A **subroutine** executes a block of code from a set of arguments but does not explicitly return a value. Changes to arguments made within a **function** are not returned whereas changes to arguments made within a **subroutine** can be returned to the calling program. Both functions and subroutines are defined after the **contains** keyword in a **module** or **program**.

6.1 Writing a function

The definition of a function starts with the name of the function followed by a list of arguments and return variable. The data types of the arguments and return variable are defined within the **function** body.

6.1.1 Example: linspace: generating a set of equally-space points

The following program defines a function `linspace` that returns a set of equidistant points on an interval. The main function makes a call to the function.

```
1  program main
2    implicit none
3
4    real :: xs(10)
5
6    ! call function linspace to set values in xs
7    xs = linspace(0.0, 1.0, 10)
8
9    ! print returned value of xs
10   print "(10(f0.1, 1x))" , xs
11
12 contains
13
14   ! linspace: return a set of equidistant points on an interval
15   ! min: minimum value of interval
16   ! max: maximum value of interval
17   ! n_points: number of points in returned set
18   ! xs: set of points
19   function linspace(min, max, n_points) result(xs)
20     real :: min, max, dx
21     integer :: n_points
22     integer :: i
23     real :: xs(n_points)
24
25     ! calculate width of subintervals
26     dx = (max-min) / real(n_points-1)
27
28     ! fill xs with points
29     do i = 1,n_points
30       xs(i) = min + (i-1)*dx
31     end do
32
33   end function linspace
34
35 end program main
```

.0 .1 .2 .3 .4 .6 .7 .8 .9 1.0

6.2 Writing a subroutine

The definition of a subroutine begins with the name of the subroutine and list of arguments. Arguments are defined within the `subroutine` body with one of the following intents

- `intent(in)`: changes to the argument are not returned
- `intent(inout)`: changes to the argument are returned
- `intent(out)`: the initial value of the argument is ignored and changes to the argument are returned.

Subroutines are called using the `call` keyword followed by the subroutine name.

6.2.1 Example: polar coordinates

The following code defines a subroutine `polar_coord` that returns the polar coordinates (r, θ) defined by $r = \sqrt{x^2 + y^2}$ and $\theta = \arctan(y/x)$ from the rectangular coordinate pair (x, y) .

```
1 program main
2
3   real :: x = 1.0, y = 1.0, rad, theta
4
5   ! call subroutine that returns polar coords
6   call polar_coord(x, y, rad, theta)
7   print*, rad, theta
8
9 contains
10
11  ! polar_coord: return the polar coordinates of a rect coord pair
12  ! x,y: rectangular coord
13  ! rad,theta: polar coord
14  subroutine polar_coord(x, y, rad, theta)
15     real, intent(in) :: x, y
16     real, intent(out) :: rad, theta
17
18     ! compute polar coord
```

```

19      ! hypot = sqrt(x**2+y**2) is an intrinsic function
20      ! atan2 = arctan with correct sign is an intrinsic function
21      rad = hypot(x, y)
22      theta = atan2(y, x)
23
24      end subroutine polar_coord
25
26 end program main

1.41421354      0.785398185

```

6.3 Passing procedures as arguments

An `interface` can be used to pass a function or subroutine to another function or a subroutine. For this purpose, an `interface` is defined in the receiving procedure essentially the same way as the passed procedure itself but with only declarations and not the implementation.

6.3.1 Example: Newton's method for rootfinding

Newton's method for finding the root of a function $f : \mathbb{R} \rightarrow \mathbb{R}$ refines an initial guess x_0 according to the iteration rule

$$x_{n+1} = x_n - \frac{f(x_n)}{f'(x_n)}$$

for $n \geq 1$ until $f(x)$ is less than a chosen tolerance or a maximum number of iterations.

The following code defines a subroutine `newton_root` that returns a root of an input function as well as the number of iterations of Newton's method used to find the root. It is called by the main program to approximate the positive root of $f(x) = x^2 - 2$ from an initial guess $x_0 = 1$.

```

1  program main
2    implicit none
3
4    character(64) :: fmt
5    real :: x = 1.0
6    integer :: iter = 1000
7
8    ! call newton rootfinding function

```

```

9     call newton_root(f, df, x, iter, 1e-6, .true.)
10
11     ! print found root and number of iterations used
12     fmt = "('number of iterations: ', i0, ', x: ', f0.7, ', f(x): ', f0.7)"
13     print fmt, iter, x, f(x)
14
15 contains
16
17     ! function f(x) = x^2 - 2
18     function f(x) result(y)
19         real :: x, y
20         y = x*x - 2
21     end function f
22
23     ! function df(x) = 2x
24     function df(x) result(dy)
25         real :: x, dy
26         dy = 2*x
27     end function df
28
29     ! newton_root: newtons method for rootfinding
30     ! f: function with root
31     ! df: derivative of f
32     ! x: sequence iterate
33     ! iter: max number of iterations at call, number of iterations at return
34     ! tol: absolute tolerance
35     ! print_iters: boolean to toggle verbosity
36     subroutine newton_root(f, df, x, iter, tol, print_iters)
37
38         ! interface to function f
39         interface
40             function f(x) result(y)
41                 real :: x, y
42             end function f
43         end interface
44
45         ! interface to function df
46         interface
47             function df(x) result(dy)
48                 real :: x, dy

```

```

49     end function df
50 end interface
51
52 real, intent(inout) :: x
53 real, intent(in) :: tol
54 integer, intent(inout) :: iter
55 logical, intent(in) :: print_iters
56 integer :: max_iters
57
58 max_iters = iter
59 iter = 0
60
61 ! while f(x) greater than absolute tolerance
62 ! and max number of iterations not exceeded
63 do while (abs(f(x))>tol.and.iter<max_iters)
64     ! print current x and f(x)
65     if (print_iters) print "('f(', f0.7, ') = ', f0.7)", x, f(x)
66
67     ! Newton's update rule
68     x = x - f(x)/df(x)
69
70     ! increment number of iterations
71     iter = iter + 1
72 end do
73
74 end subroutine newton_root
75
76 end program main

f(1.0000000) = -1.0000000
f(1.5000000) = .2500000
f(1.4166666) = .0069444
f(1.4142157) = .0000060
number of iterations: 4, x: 1.4142135, f(x): -.0000001

```

6.3.2 Example: The midpoint rule for definite integrals

The midpoint rule approximates the definite integral $\int_a^b f(x) dx$ with integrand $f : \mathbb{R} \rightarrow \mathbb{R}$ by

$$\Delta x \sum_{i=1}^n f(\bar{x}_i) \quad (1)$$

where $\Delta x = (b - a)/n$, $x_i = a + i\Delta x$ and $\bar{x}_i = (x_{i-1} + x_i)/2$.

The following code defines a function `midpoint` that computes the approximation eq. 1 given a , b , and n . The main program calls `midpoint` to approximate the definite integral of $f(x) = 1/x$ on $[1, e]$ for a range of n .

```
1 program main
2   implicit none
3
4   real, parameter :: E = exp(1.)
5   integer :: n
6   real :: integral
7
8   ! Approximate the integral of 1/x from 1 to e
9   ! with the midpoint rule for a range of number of subintervals
10  do n = 2,20,2
11    print "('n: ', i0, ', M_n: ', f0.6)", n, midpoint(f, 1.0, E, n)
12  end do
13
14 contains
15
16   ! function f(x) = 1/x
17   function f(x) result(y)
18     real :: x, y
19     y = 1.0/x
20   end function f
21
22   ! midpoint: midpoint rule for definite integral
23   ! f: integrand
24   ! a: left endpoint of interval of integration
25   ! b: right endpoint of interval of integration
26   ! n: number of subintervals
27   ! sum: approximate definite integral
28   function midpoint(f, a, b, n) result(sum)
29
```

```

30     ! interface to f
31     interface
32         function f(x)
33             real :: x, y
34         end function f
35     end interface
36
37     real :: a, b, min, xi, dx, sum
38     integer :: n, i
39
40     ! subinterval increment
41     dx = (b-a)/real(n)
42     ! minimum to increment from
43     min = a - dx/2.0
44
45     ! midpoint rule
46     do i = 1,n
47         xi = min + i*dx
48         sum = sum + f(xi)
49     end do
50     sum = sum*dx
51 end function midpoint
52
53 end program main

```

```

n: 2, M_n: .976360
n: 4, M_n: .993575
n: 6, M_n: .997091
n: 8, M_n: .998353
n: 10, M_n: .998942
n: 12, M_n: .999264
n: 14, M_n: .999459
n: 16, M_n: .999585
n: 18, M_n: .999672
n: 20, M_n: .999735

```

6.4 Polymorphism

An interface can be used as an entry into two different implementations of a subroutine or function with the same name so long as the different im-

plementations have different argument signatures. This may be particularly useful for defining both a single precision and double precision version of a function or subroutine.

6.4.1 Example: machine epsilon

The following code implements two versions of a function that computes machine epsilon in either single or double precision. The different implementations are distinguished by their arguments. The single precision version `mach_eps_sp` accepts one single precision float and the double precision version `mach_eps_dp` accepts one double precision float. Both functions are listed in the `interface` and can be called by its name `mach_eps`.

```
1  program main
2    implicit none
3
4    integer, parameter :: sp = kind(0.0)
5    integer, parameter :: dp = kind(0.d0)
6
7    interface mach_eps
8      procedure mach_eps_sp, mach_eps_dp
9    end interface mach_eps
10
11   print*, mach_eps(0.0_sp), epsilon(0.0_sp)
12   print*, mach_eps(0.0_dp), epsilon(0.0_dp)
13
14   contains
15
16   function mach_eps_sp(x) result(eps)
17     real(sp) :: x, eps
18     integer :: count = 0
19
20     eps = 1.0_sp
21     do while (1.0_sp + eps*0.5 > 1.0_sp)
22       eps = eps*0.5
23       count = count+1
24     end do
25   end function mach_eps_sp
26
27   function mach_eps_dp(x) result(eps)
```

```

28     real(dp) :: x, eps
29     integer :: count = 0
30
31     eps = 1.0_dp
32     do while (1.0_dp + eps*0.5 > 1.0_dp)
33         eps = eps*0.5
34         count = count+1
35     end do
36 end function mach_eps_dp
37
38 end program main

1.19209290E-07   1.19209290E-07
2.2204460492503131E-016   2.2204460492503131E-016

```

6.5 Recursion

A function or subroutine that calls itself must be defined with the `recursive` keyword preceding the construct name.

6.5.1 Example: factorial

The following code defines a recursive function `factorial` that computes $n!$. If $n > 1$, the function call itself to return $n(n - 1)!$, otherwise the function returns 1. The main program calls `factorial` to compute $5!$.

```

1  program main
2  implicit none
3
4  ! print 5 factorial
5  print*, factorial(5)
6
7  contains
8
9  ! factorial(n): product of natural numbers up to n
10 ! n: integer argument
11 recursive function factorial(n) result(m)
12     integer :: n, m
13
14     ! if n>1, call factorial recursively
15     ! otherwise 1 factorial is 1

```

```

16     if (n>1) then
17         m = n*factorial(n-1)
18     else
19         m = 1
20     end if
21
22 end function factorial
23
24 end program main

```

120

7 Object-oriented programming

7.1 Derived types

Data types can be defined by the programmer. Variables and procedures that belong to a defined data type are declared between a **type / end type** pair. Type-bound procedures, i.e. functions and subroutines, are defined by the **procedure** keyword followed by **::** and the name of the procedure within the **type / end type** pair after the **contains** keyword. A variable with defined type is declared with the **type** keyword and the name of the type. The variables and procedures of a defined type variable can be accessed by appending a **%** symbol to the name of the variable.

```

1  ! define a 'matrix' type
2  ! type-bound variables: shape, data
3  ! type-bound procedures: construct, destruct
4  type matrix
5      integer :: shape(2)
6      real, allocatable :: data(:, :)
7  contains
8      procedure :: construct
9      procedure :: destruct
10 end type matrix
11
12 ! declare a matrix variable
13 type(matrix) :: mat
14
15 ! assign value to type-bound variable
16 mat%shape = [3,3]

```

7.2 Modules

A type-bound procedure can be defined after the `contains` keyword in the same program construct, i.e. a `module`, as the type definition. The first argument in the definition of a type-bound procedure is of the defined type and is declared within the procedure body with the `class` keyword and the name of the type.

```
1  module matrix_module
2    implicit none
3
4    type matrix
5      integer :: shape(2)
6      real, allocatable :: data(:, :)
7    contains
8      procedure :: construct
9      procedure :: destruct
10   end type matrix
11
12  contains
13
14   ! construct: populate shape and allocate memory for matrix
15   ! m,n: number of rows,cols of matrix
16   subroutine construct(this, m, n)
17     class(matrix) :: this
18     integer :: m, n
19     this%shape = [m,n]
20     allocate(this%data(m,n))
21   end subroutine construct
22
23   ! destruct: deallocate memory that matrix occupies
24   subroutine destruct(this)
25     class(matrix) :: this
26     deallocate(this%data)
27   end subroutine destruct
28
29  end module matrix_module
```

To define variables of the `matrix` type in the main program, tell it to `use` the module defined above with `use matrix_module` immediately after the `program main` line. The procedures bound to a defined type can be access

through variables of that type by appending the % symbol to the name of the variable.

```
1 program main
2   use matrix_module
3   implicit none
4
5   type(matrix) :: mat
6   mat%shape = [3,3]
7
8   ! create matrix
9   call mat%construct(3,3)
10
11  ! treat matrix variable 'data' like an array
12  mat%data(1,1) = 1.0
13  ! etc...
14
15  ! destruct matrix
16  call matrix%destruct()
17 end program main
```

7.3 Example: determinant of random matrix

The following module defines a `matrix` type with two variables: an `integer` array `shape` that stores the number of rows and columns of the matrix and a `real` array `data` that stores the elements of the matrix. The type has four procedures: a subroutine `construct` that sets the shape and allocates memory for the data, a subroutine `destruct` that deallocates memory, a subroutine `print` that prints a matrix, and a function `det` that computes the determinant of a matrix. Note `det` is based on the definition of determinant using cofactors, and is very inefficient. A function `random_matrix` defined within the module generates a matrix with uniform random entries in $[-1, 1]$.

```
1 module matrix_module
2   implicit none
3
4   type matrix
5     integer :: shape(2)
6     real, allocatable :: data(:, :)
7   contains
```

```

8      procedure :: construct
9      procedure :: destruct
10     procedure :: print
11     procedure :: det
12 end type matrix
13
14 contains
15
16     subroutine construct(this, m, n)
17         class(matrix) :: this
18         integer :: m,n
19         this%shape = [m,n]
20         allocate(this%data(m,n))
21     end subroutine construct
22
23     subroutine destruct(this)
24         class(matrix) :: this
25         deallocate(this%data)
26     end subroutine destruct
27
28     ! print: formatted print of matrix
29     subroutine print(this)
30         class(matrix) :: this
31         ! row_fmt: format character string for row printing
32         ! fmt: temporary format string
33         character(32) :: row_fmt, fmt = '(a,i0,a,i0,a,i0,a)'
34         ! w: width of each entry printed
35         ! d: number of decimal digits printed
36         integer :: w, d = 2, row
37         ! find largest width of element in matrix
38         w = ceiling(log10(maxval(abs(this%data)))) + d + 2
39         ! write row formatting to 'row_fmt' variable
40         write(row_fmt,fmt) '(, ,this%shape(2), (f',w,',',d,',1x))'
41         ! print matrix row by row
42         do row = 1,this%shape(1)
43             print row_fmt, this%data(row,:)
44         end do
45     end subroutine print
46
47     ! det: compute determinant of matrix

```

```

48   ! using recursive definition based on cofactors
49   recursive function det(this) result(d)
50     class(matrix) :: this
51     type(matrix) :: submatrix
52     real :: d, sgn, element, minor
53     integer :: m, n, row, col, i, j
54
55     m = this%shape(1)
56     n = this%shape(2)
57     d = 0.0
58
59     ! compute cofactor
60     ! if 1x1 matrix, return value
61     if (m==1.and.n==1) then
62       d = this%data(1,1)
63     ! if square and not 1x1
64     else if (m==n) then
65       ! cofactor sum down the first column
66       do row = 1,m
67         ! sign of term
68         sgn = (-1.0)**(row+1)
69         ! matrix element
70         element = this%data(row,1)
71         ! construct the cofactor submatrix and compute its determinant
72         call submatrix%construct(m-1,n-1)
73         if (row==1) then
74           submatrix%data = this%data(2:,2:)
75         else if (row==m) then
76           submatrix%data = this%data(:m-1,2:)
77         else
78           submatrix%data(:row-1,:) = this%data(:row-1,2:)
79           submatrix%data(row,:,:) = this%data(row+1:,2:)
80         end if
81         minor = submatrix%det()
82         call submatrix%destruct()
83
84         ! determinant accumulator
85         d = d + sgn*element*minor
86       end do
87     end if

```

```

88     end function det
89
90     ! random_matrix: generate matrix with random entries in [-1,1]
91     ! m,n: number of rows,cols
92     function random_matrix(m,n) result(mat)
93         integer :: m,n,i,j
94         type(matrix) :: mat
95         ! allocate memory for matrix
96         call mat%construct(m,n)
97         ! seed random number generator
98         call srand(time())
99         ! populate matrix
100        do i = 1,m
101            do j = 1,n
102                mat%data(i,j) = 2.0*rand() - 1.0
103            end do
104        end do
105    end function random_matrix
106
107 end module matrix_module

```

The main program uses the `matrix_module` defined above to find the determinants of a number of random matrices of increasing size.

```

1  program main
2      use matrix_module
3      implicit none
4
5      type(matrix) :: mat
6      integer :: n
7
8      ! compute determinants of random matrices
9      do n = 1,5
10         ! generate random matrix
11         mat = random_matrix(n,n)
12
13         ! print determinant of matrix
14         print "('n: ', i0, ', det: ', f0.5)", n, det(mat)
15
16         ! destruct matrix

```



```
17     call mat%destruct()
18   end do
19
20 end program main
```

```
./main
```

```
n: 1, det: -.68676
n: 2, det: .45054
n: 3, det: .37319
n: 4, det: -.27328
n: 5, det: .26695
```

7.4 Example: matrix module

```
1  module matrix_module
2    implicit none
3
4    public :: zeros
5    public :: identity
6    public :: random
7
8    type matrix
9      integer :: shape(2)
10     real, allocatable :: data(:, :)
11     contains
12     procedure :: construct => matrix_construct
13     procedure :: destruct => matrix_destruct
14     procedure :: norm => matrix_norm
15   end type matrix
16
17   type vector
18     integer :: length
19     real, allocatable :: data(:)
20     contains
21     procedure :: construct => vector_construct
22     procedure :: destruct => vector_destruct
23     procedure :: norm => vector_norm
24   end type vector
25
```

```

26   ! assignments
27   interface assignment(=)
28       procedure vec_num_assign, vec_vec_assign, mat_num_assign, mat_mat_assign
29   end interface assignment(=)
30
31   ! operations
32   interface operator(+)
33       procedure vec_vec_sum, mat_mat_sum
34   end interface operator(+)
35
36   interface operator(-)
37       procedure vec_vec_diff, mat_mat_diff
38   end interface operator(-)
39
40   interface operator(*)
41       procedure num_vec_prod, num_mat_prod, mat_vec_prod, mat_mat_prod
42   end interface operator(*)
43
44   interface operator(/)
45       procedure vec_num_quot, mat_num_quot
46   end interface operator(/)
47
48   interface operator(**)
49       procedure mat_pow
50   end interface operator(**)
51
52   ! functions
53   interface norm
54       procedure vector_norm, matrix_norm
55   end interface norm
56
57   ! structured vectors/matrices
58   interface zeros
59       procedure zeros_vector, zeros_matrix
60   end interface zeros
61
62   interface random
63       procedure random_vector, random_matrix
64   end interface random
65

```

```

66 contains
67
68   subroutine matrix_construct(this, m, n)
69     class(matrix) :: this
70     integer :: m,n
71     this%shape = [m,n]
72     allocate(this%data(m,n))
73   end subroutine matrix_construct
74
75   subroutine vector_construct(this, n)
76     class(vector) :: this
77     integer :: n
78     this%length = n
79     allocate(this%data(n))
80   end subroutine vector_construct
81
82   subroutine matrix_destruct(this)
83     class(matrix) :: this
84     deallocate(this%data)
85   end subroutine matrix_destruct
86
87   subroutine vector_destruct(this)
88     class(vector) :: this
89     deallocate(this%data)
90   end subroutine vector_destruct
91
92   ! assignment
93   subroutine vec_num_assign(vec,num)
94     type(vector), intent(inout) :: vec
95     real, intent(in) :: num
96     vec%data = num
97   end subroutine vec_num_assign
98
99   subroutine vec_vec_assign(vec1,vec2)
100    type(vector), intent(inout) :: vec1
101    type(vector), intent(in) :: vec2
102    vec1%data = vec2%data
103  end subroutine vec_vec_assign
104
105  subroutine mat_num_assign(mat,num)

```

```

106     type(matrix), intent(inout) :: mat
107     real, intent(in) :: num
108     mat%data = num
109 end subroutine mat_num_assign
110
111 subroutine mat_mat_assign(mat1,mat2)
112     type(matrix), intent(inout) :: mat1
113     type(matrix), intent(in) :: mat2
114     mat1%data = mat2%data
115 end subroutine mat_mat_assign
116
117 ! operations
118 function vec_vec_sum(vec1,vec2) result(s)
119     type(vector), intent(in) :: vec1, vec2
120     type(vector) :: s
121     call s%construct(vec1%length)
122     s%data = vec1%data + vec2%data
123 end function vec_vec_sum
124
125 function mat_mat_sum(mat1,mat2) result(s)
126     type(matrix), intent(in) :: mat1, mat2
127     type(matrix) :: s
128     call s%construct(mat1%shape(1),mat1%shape(2))
129     s%data = mat1%data+mat2%data
130 end function mat_mat_sum
131
132 function vec_vec_diff(vec1,vec2) result(diff)
133     type(vector), intent(in) :: vec1, vec2
134     type(vector) :: diff
135     call diff%construct(vec1%length)
136     diff%data = vec1%data-vec2%data
137 end function vec_vec_diff
138
139 function mat_mat_diff(mat1,mat2) result(diff)
140     type(matrix), intent(in) :: mat1, mat2
141     type(matrix) :: diff
142     call diff%construct(mat1%shape(1),mat1%shape(2))
143     diff%data = mat1%data-mat2%data
144 end function mat_mat_diff
145

```

```

146 function num_vec_prod(num,vec) result(prod)
147     real, intent(in) :: num
148     type(vector), intent(in) :: vec
149     type(vector) :: prod
150     call prod%construct(vec%length)
151     prod%data = num*vec%data
152 end function num_vec_prod
153
154 function num_mat_prod(num,mat) result(prod)
155     real, intent(in) :: num
156     type(matrix), intent(in) :: mat
157     type(matrix) :: prod
158     call prod%construct(mat%shape(1),mat%shape(2))
159     prod%data = num*mat%data
160 end function num_mat_prod
161
162 function mat_vec_prod(mat,vec) result(prod)
163     type(matrix), intent(in) :: mat
164     type(vector), intent(in) :: vec
165     type(vector) :: prod
166     call prod%construct(mat%shape(1))
167     prod%data = matmul(mat%data,vec%data)
168 end function mat_vec_prod
169
170 function mat_mat_prod(mat1,mat2) result(prod)
171     type(matrix), intent(in) :: mat1, mat2
172     type(matrix) :: prod
173     call prod%construct(mat1%shape(1),mat2%shape(2))
174     prod%data = matmul(mat1%data,mat2%data)
175 end function mat_mat_prod
176
177 function vec_num_quot(vec,num) result(quot)
178     type(vector), intent(in) :: vec
179     real, intent(in) :: num
180     type(vector) :: quot
181     call quot%construct(vec%length)
182     quot%data = vec%data/num
183 end function vec_num_quot
184
185 function mat_num_quot(mat,num) result(quot)

```

```

186     type(matrix), intent(in) :: mat
187     real, intent(in) :: num
188     type(matrix) :: quot
189     call quot%construct(mat%shape(1),mat%shape(2))
190     quot%data = mat%data/num
191 end function mat_num_quot
192
193 function mat_pow(mat1,pow) result(mat2)
194     type(matrix), intent(in) :: mat1
195     integer, intent(in) :: pow
196     type(matrix) :: mat2
197     integer :: i
198     mat2 = mat1
199     do i = 2,pow
200         mat2 = mat1*mat2
201     end do
202 end function mat_pow
203
204 ! functions
205 function vector_norm(this,p) result(mag)
206     class(vector), intent(in) :: this
207     integer, intent(in) :: p
208     real :: mag
209     integer :: i
210     ! inf-norm
211     if (p==0) then
212         mag = 0.0
213         do i = 1,this%length
214             mag = max(mag,abs(this%data(i)))
215         end do
216     ! p-norm
217     else if (p>0) then
218         mag = (sum(abs(this%data)**p))**(1./p)
219     end if
220 end function vector_norm
221
222 function matrix_norm(this, p) result(mag)
223     class(matrix), intent(in) :: this
224     integer, intent(in) :: p
225     real :: mag, tol = 1e-6

```

```

226     integer :: m, n, row, col, iter, max_iters = 1000
227     type(vector) :: vec, last_vec
228     m = size(this%data(:,1)); n = size(this%data(1,:))
229
230     ! entry-wise norms
231     if (p<0) then
232         mag = (sum(abs(this%data)**(-p)))**(-1./p)
233     ! inf-norm
234     else if (p==0) then
235         mag = 0.0
236         do row = 1,m
237             mag = max(mag,sum(abs(this%data(row,:)))
238         end do
239     ! 1-norm
240     else if (p==1) then
241         mag = 0.0
242         do col = 1,n
243             mag = max(mag,sum(abs(this%data(:,col))))
244         end do
245     ! p-norm
246     else if (p>0) then
247         vec = random(n)
248         vec = vec/vec%norm(p)
249         last_vec = zeros(n)
250         mag = 0.0
251         do iter = 1,max_iters
252             last_vec = vec
253             vec = this*last_vec
254             vec = vec/vec%norm(p)
255             if (vector_norm(vec-last_vec,p)<tol) exit
256         end do
257         mag = vector_norm(this*vec,p)
258     end if
259 end function matrix_norm
260
261 ! structured vectors/matrices
262 function random_matrix(m,n) result(mat)
263     integer :: m,n
264     type(matrix) :: mat
265     call mat%construct(m,n)

```

```

266     call random_seed()
267     call random_number(mat%data)
268 end function random_matrix
269
270 function random_vector(n) result(vec)
271     integer :: n
272     type(vector) :: vec
273     call vec%construct(n)
274     call random_seed()
275     call random_number(vec%data)
276 end function random_vector
277
278 function zeros_vector(n) result(vec)
279     integer :: n
280     type(vector) :: vec
281     call vec%construct(n)
282     vec = 0.0
283 end function zeros_vector
284
285 function zeros_matrix(m,n) result(mat)
286     integer :: m,n
287     type(matrix) :: mat
288     call mat%construct(m,n)
289     mat = 0.0
290 end function zeros_matrix
291
292 function identity(m,n) result(mat)
293     integer :: m,n,i
294     type(matrix) :: mat
295     call mat%construct(m,n)
296     do i = 1,min(m,n)
297         mat%data(i,i) = 1.0
298     end do
299 end function identity
300
301 end module matrix_module

1 program main
2     use matrix_module
3     implicit none

```



```

4
5  type(vector) :: vec1, vec2
6  type(matrix) :: mat1, mat2
7  real :: x
8  integer :: i
9
10 ! 0s, id, random
11 mat1 = zeros(3,3)
12 call mat1%destruct()
13 mat1 = identity(3,3)
14 mat2 = random(3,3)
15 mat1 = mat1*mat1
16 vec1 = zeros(3)
17 call vec1%destruct()
18 vec1 = random(3)
19 vec2 = random(3)
20 ! +,-,*,/,**
21 mat1 = mat1+mat2
22 vec1 = vec1+vec2
23 mat1 = mat1-mat2
24 vec1 = vec1-vec2
25 vec1 = mat1*vec2
26 mat1 = mat2*mat1
27 mat1 = 2.0*mat1
28 vec1 = 2.0*vec1
29 mat1 = mat1/2.0
30 vec1 = vec1/2.0
31 mat2 = mat1**3
32 ! norm
33 x = norm(vec1,0)
34 x = norm(vec1,1)
35 x = norm(mat1,-1)
36 x = norm(mat1,0)
37 x = norm(mat1,1)
38 x = norm(mat1,2)
39 call vec1%destruct
40 call vec2%destruct
41 call mat1%destruct
42 call mat2%destruct
43 end program main

```

./main