| AMath 483/583 - Lecture 8 |
| :---: |
| This lecture: <br> - Fortran subroutines and functions <br> - Arrays <br> - Dynamic memory |
| Reading: <br> - class notes: Fortran Arrays <br> - class notes: Fortran Subroutines and Functions <br> - class notes: gfortran flags |
| R.J. Leveque, Univesity of Wastington AMait 4835633, Lecture 8 |

## Fortran functions and subroutines

For now, assume we have a single file filename. $f 90$ that contains the main program and also any functions or subroutines needed.

Later we will see how to split into separate files.
Will also discuss use of modules.

Functions take some input arguments and return a single value.
Usage: $\quad y=f(x)$ or $\quad z=g(x, y)$

Should be declared as external with the type of value returned:
real(kind=8), external :: f
R.J. LeVeque, University of Washington AMath 483/583, Lecture 8

## Fortran functions

```
! $UWHPSC/codes/fortran/fcn1.f90
program fcn1
    implicit none
    real(kind=8) :: y,z
    real(kind=8), external :: f
    y = 2.
    z=f(y)
    print *, "z = ",z
end program fcn1
real(kind=8) function f(x)
    implicit none
    real(kind=8), intent(in) :: x
    f= x**2
end function f
```

Prints out: z

## Notes:

## Notes:

R.J. LeVeque, University of Washington AMath 483/583, Lecture 8

## Notes:

## Fortran subroutines

Subroutines have arguments, each of which might be for input or output or both.

Usage: call sub1 ( $x, y, z, a, b)$

Can specify the intent of each argument, e.g.

```
real(kind=8), intent(in) :: x,y
real(kind=8), intent(out) :: z
real(kind=8), intent(inout) :: a,b
```

specifies that $\mathrm{x}, \mathrm{y}$ are passed in and not modified, z may not have a value coming in but will be set by sub1, $\mathrm{a}, \mathrm{b}$ are passed in and may be modified.

After this call, $\mathrm{z}, \mathrm{a}, \mathrm{b}$ may all have changed.
R.J. LeVeque, University of Washington AMath 483/583, Lecture 8

| Fortran subro | utines |
| :---: | :---: |
| 1 | ! \$UWHPSC/codes/fortran/sub1.f9e |
| 2 |  |
| 3 | program sub1 |
| 4 | implicit none |
| 5 6 | real (kind=8) :: y,z |
| 7 | $y=2$. |
| 8 | call fsub(y,z) |
| 9 | print *, "z = ", z |
| 18 | end program sub1 |
| 11 | subroutine fsub(x,f) |
| 13 | implicit none |
| 14 | real (kind=8), ${ }^{\text {intent }}$ (in) $: ~: ~$ ¢ |
| 15 16 | $\underset{\substack{\text { real (kind=8), intent(out) } \\ f=x * * 2}}{\text { a }}$ f |
| 17 | end subroutine fsub |

R.J. LeVeque, University of Washington AMath 483/583, Lecture 8

## Fortran subroutines

A version that takes an array as input and squares each value:

```
```

! \$UWHPSC/codes/fortran/sub2.f90

```
```

! \$UWHPSC/codes/fortran/sub2.f90
program sub2
program sub2
implicit none
implicit none
real(kind=8), dimension(3) :: y,z
real(kind=8), dimension(3) :: y,z
integer n
integer n
y = (/2., 3., 4./)
y = (/2., 3., 4./)
n = size(y)
n = size(y)
call fsub(y,n,z)
call fsub(y,n,z)
print *, "z = ",z
print *, "z = ",z
end program sub2
end program sub2
subroutine fsub(x,n,f)
subroutine fsub(x,n,f)
! compute f(x)= x**2 for all elements of the array }
! compute f(x)= x**2 for all elements of the array }
! of length n.
! of length n.
implicit none
implicit none
integer, intent(in) :: n
integer, intent(in) :: n
real(kind=8), dimension(n), intent(in) :: x
real(kind=8), dimension(n), intent(in) :: x
real(kind=8), dimension(n), intent(out) :: f
real(kind=8), dimension(n), intent(out) :: f
f = x**2
f = x**2
end subroutine fsub

```
```

end subroutine fsub

```
```


R.J. LeVeque, University of Washington AMath 483/583, Lecture 8

## Array operations in Fortran

! \$UWHPSC/codes/fortran/vectorops.f90
! continued...
print *, "x*y = " print *, $x * y \quad!=(x(1) y(1), x(2) y(2), \ldots)$
print *, "sqrt(y) = "
print *, sqrt(y) ! componentwise
print *, "dot_product $(x, y)="$ print *, dot_product $(x, y)$ ! scalar product
end program vectorops
R.J. LeVeque, University of Washington AMath 483/583, Lecture 8

## Array operations in Fortran - Matrices

! \$UWHPSC/codes/fortran/arrayops.f90
program arrayops
implicit none
real(kind=8), dimension(3,2) :: a
! create a as $3 x 2$ array:
$A=\operatorname{reshape}((/ 1,2,3,4,5,6 /),(/ 3,2 /))$

## Note:

- Fortran is case insensitive: $\mathrm{A}=\mathrm{a}$
- Reshape fills array by columns, so

$$
A=\left[\begin{array}{ll}
1 & 4 \\
2 & 5 \\
3 & 6
\end{array}\right]
$$



| Array | perations in Fortran - Matrices |
| :---: | :---: |
|  | ```$UWHPSC/codes/fortran/arrayops.f90 (continued) real(kind=8), dimension(3,2) :: a real(kind=8), dimension(2) :: x real(kind=8), dimension(3) :: y x = (/5,6/) y = matmul(a,x) ! matrix-vector product print *, "x = ",x print *, "y = ",y``` |

## Array operations in Fortran - Matrices

! \$UWHPSC/codes/fortran/arrayops.f90 (continued)
real(kind=8), dimension(3,2) : : a
real(kind=8), dimension(2) : : x
real(kind=8), dimension(3) : : y
$x=(/ 5,6 /)$
print *, "x = ", x
print *, "Y = ", Y
R.J. LeVeque, University of Washington AMath 483/583, Lecture 8

## Linear systems in Fortran

There is no equivalent of the Matlab backslash operator for solving a linear system $A x=b \quad(b=A \backslash b)$

Must call a library subroutine to solve a system.
Later we will see how to use LAPACK for this.

Note: Under the hood, Matlab calls LAPACK too!
So does NumPy.

## Notes:

R.J. LeVeque, University of Washington AMath 483/583, Lecture 8

## Notes:

R.J. LeVeque, University of Washington AMath 483/583, Lecture 8

## Notes:

## Array storage

Rank 1 arrays have a single index, for example:

```
real(kind=8) :: x(3)
real(kind=8), dimension(3) :: x
```

are equivalent ways to define x with elements $x(1), x(2), x(3)$.

You can also specify a different starting index:

```
real(kind=8) :: x(0:2), y(4:6), z(-2:0)
```

These are all arrays of length 3 and this would be a valid assignment:

$$
y(5)=z(-2)
$$

R.J. LeVeque, University of Washington

AMath $483 / 583$, Lecture 8

## Multi-dimensional array storage

Memory can be thought of linear, indexed by a single address.
A one-dimensional array of length N will generally occupy N consecutive memory locations: 8 N bytes for floats.

A two-dimensional array (e.g. matrix) of size $m \times n$ will require $m n$ memory locations.

Might be stored by rows, e.g. first row, followed by second row, etc.

This what's done in Python or C , as suggested by notation:
$A=[[10,20,30],[40,50,60]]$
Or, could be stored by columns, as done in Fortran!
R.J. LeVeque, University of Washington AMath 483/583, Lecture 8

## Multi-dimensional array storage

$$
A=\left[\begin{array}{lll}
10 & 20 & 30 \\
40 & 50 & 60
\end{array}\right]
$$

Apy $=$ reshape (array $([10,20,30,40,50,60]),(3,2))$
Afort = reshape((/10,20,30,40,50,60/), (/3,2/))
Suppose the array storage starts at memory location 3401.
In Python or Fortran, the elements will be stored in the order:

| loc 3401 | Apy $[0,0]=10$ | Afort $(1,1)=10$ |
| :--- | :--- | :--- |
| loc 3402 | Apy $[0,1]=20$ | Afort $(2,1)=40$ |
| loc 3403 | Apy $[0,2]=30$ | Afort $(1,2)=20$ |
| loc 3404 | Apy $[1,0]=40$ | Afort $(2,2)=50$ |
| loc 3405 | Apy $[1,1]=50$ | Afort $(1,3)=30$ |
| loc 3406 | Apy $[1,2]=60$ | Afort $(2,3)=60$ |

## Notes:

R.J. LeVeque, University of Washington AMath 483/583, Lecture 8

## Notes:

R.J. LeVeque, University of Washington AMath 483/583, Lecture 8

## Notes:

## Aside on np.reshape

The np.reshape method can go through data in either order:

```
>>> v = linspace (10,60,6)
>>> v
array([ 10., 20., 30., 40., 50., 60.])
>>> reshape(v, (2,3)) # order='C' by default
array([[ 10., 20., 30.],
    [ 40., 50., 60.]])
>>> reshape(v, (2,3), order=' F')
array([[ 10., 30., 50.],
    [ 20., 40., 60.]])
```


## Aside on np.reshape

The np.reshape method can go through data in either order:

```
>>> A
array([[ 10., 20., 30.],
    [ 40., 50., 60.]])
>>> A.reshape(3,2) # order='C' by default
array([[ 10., 20.],
    [ 30., 40.],
    [ 50., 60.]])
>>> A.reshape((3,2),order=' F')
array([[ 10., 50.],
    [ 40., 30.],
    [ 20., 60.]])
```

Note: reshape can be called as function or method of A...

```
>>> reshape(A, (3,2), order=' F')
R.J. LeVeque, University of Washington AMath 483/583, Lecture 8
```


## Aside on np.flatten

The np.flatten method converts an N -dim array to a 1-dimensional one:

```
>>> A = np.array([[10.,20,30],[40,50,60]])
>>> A
array([[ 10., 20., 30.],
    [ 40., 50., 60.]])
>>> A.flatten() # Default is 'C'
array([ 10., 20., 30., 40., 50., 60.])
>>> A.flatten('F') # Fortran ordering
array([ 10., 40., 20., 50., 30., 60.])
```


## Memory management for arrays

Often a program needs to be written to handle arrays whose size is not known until the program is running.

Fortran 77 approaches:

- Allocate arrays large enough for any application,
- Use "work arrays" that are partitioned into pieces.

We will look at some examples from LAPACK since you will probably see this in other software!

## The good news:

Fortran 90 allows dynamic memory allocation.
R.J. LeVeque, University of Washington AMath 483/583, Lecture 8

```
Memory allocation
real(kind=8) dimension(:), allocatable :: x
real(kind=8) dimension(:,:), allocatable :: a
allocate(x(10))
allocate (a(30,10))
! use arrays...
! then clean up:
deallocate(x)
deallocate(a)
```


## Memory allocation

If you might run out of memory, use optional argument stat to return the status...

```
real(kind=8), dimension(:,:), allocatable :: a
allocate(a(30000,10000), stat=alloc_error)
if (alloc_error /= 0) then
    print *, "Insufficient memory"
    stop
    endif
```


## Notes:

## Notes:

## R.J. LeVeque, University of Washington AMath 483/583, Lecture 8

## Notes:



Note: x is a scalar, setvals dummy argument a is an array.

## Passing arrays to subroutines

The call setvals(x) statement passes the address where $x$ is stored.

In the subroutine, the array a of length 3 is assumed to start at this address. So next $3 * 8=24$ bytes are assumed to be elements of $a(1: 3)$.

In fact these 24 bytes are occupied by
x (8 bytes),
y (8 bytes),
i (4 bytes),
j (4 bytes).
So setting a (1:3) changes all these variables!

## R.J. LeVeque, University of Washington AMath 483/583, Lecture 8

## Passing arrays to subroutines

This produces:

```
x = 5.00000000000000
y = 5.00000000000000
i = 1075052544
j = 0
```

Nasty!!

- The storage location of $x$ and the next 2 storage locations were all set to the floating point value 5.0 e 0
- This messed up the values originally stored in $y, i, j$.
- Integers are stored differently than floats. Two integers take up 8 bytes, the same as one float, so the assignment a $(3)=5$. overwrites both $i$ and $j$.
- The first half of the float 5 ., when interpreted as an integer, is huge.


Note: We now try to set 1000 elements in memory!
R.J. LeVeque, University of Washington AMath 483/583, Lecture 8

## Passing arrays to subroutines

This compiles fine, but running it gives:
Segmentation fault

This means that the program tried to change a value of memory it was not allowed to.

Only a small amount of memory is devoted to the variables declared.

The memory we tried to access might be where the program itself is stored, or something related to another program that's running.

## Segmentation faults

Debugging segmentation faults can be difficult.
Tip: Compile using -fbounds-check option of gfortran.
This catches some cases when you try to access an array out of bounds.

But not the case just shown! The variable was passed to a subroutine that doesn't know how long the array should be.

For a case where this helps, see
\$UWHPSC/codes/fortran/segfault1.f90

## Notes:

R.J. LeVeque, University of Washington AMath 483/583, Lecture 8

## Notes:

R.J. LeVeque, University of Washington AMath 483/583, Lecture 8

## Notes:

