

The NetCDF Fortran 90 Interface Guide

NetCDF Version 3.6.1
31 January 2006

Russ Rew, Glenn Davis, Steve Emmerson, and Harvey Davies
Unidata Program Center

Copyright © 2005-2006 University Corporation for Atmospheric Research

Permission is granted to make and distribute verbatim copies of this manual provided that the copyright notice and these paragraphs are preserved on all copies. The software and any accompanying written materials are provided “as is” without warranty of any kind. UCAR expressly disclaims all warranties of any kind, either expressed or implied, including but not limited to the implied warranties of merchantability and fitness for a particular purpose.

The Unidata Program Center is managed by the University Corporation for Atmospheric Research and sponsored by the National Science Foundation. Any opinions, findings, conclusions, or recommendations expressed in this publication are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

Mention of any commercial company or product in this document does not constitute an endorsement by the Unidata Program Center. Unidata does not authorize any use of information from this publication for advertising or publicity purposes.

Table of Contents

1	Use of the NetCDF Library	1
1.1	Creating a NetCDF Dataset	1
1.2	Reading a NetCDF Dataset with Known Names	2
1.3	Reading a netCDF Dataset with Unknown Names	2
1.4	Writing Data in an Existing NetCDF Dataset	3
1.5	Adding New Dimensions, Variables, Attributes	4
1.6	Error Handling	5
1.7	Compiling and Linking with the NetCDF Library	6
2	Datasets	7
2.1	Datasets Introduction	7
2.2	NetCDF Library Interface Descriptions	7
2.3	NF90_STRERROR	8
2.4	Get netCDF library version: NF90_INQ_LIBVERS	8
2.5	NF90_CREATE	9
2.6	NF90_OPEN	10
2.7	NF90_REDEF	12
2.8	NF90_ENDDEF	13
2.9	NF90_CLOSE	14
2.10	NF90_INQUIRE Family	15
2.11	NF90_SYNC	16
2.12	NF90_ABORT	18
2.13	NF90_SET_FILL	19
3	Dimensions	21
3.1	Dimensions Introduction	21
3.2	NF90_DEF_DIM	21
3.3	NF90_INQ_DIMID	22
3.4	NF90_INQUIRE_DIMENSION	23
3.5	NF90_RENAME_DIM	24
4	Variables	27
4.1	Variables Introduction	27
4.2	Language Types Corresponding to netCDF external data types	27
4.3	Create a Variable: NF90_DEF_VAR	28
4.4	Get a Variable ID from Its Name: NF90_INQ_VARID	29
4.5	Get Information about a Variable from Its ID: NF90_INQUIRE_VARIABLE	30
4.6	Writing Data Values: NF90_PUT_VAR	32
4.7	Reading Data Values: NF90_GET_VAR	36
4.8	Reading and Writing Character String Values	41

4.9	Fill Values	42
4.10	NF90_RENAME_VAR	43
5	Attributes	45
5.1	Attributes Introduction	45
5.2	Attribute Conventions	45
5.3	Create an Attribute: NF90_PUT_ATT	48
5.4	Get Information about an Attribute: NF90_INQUIRE_ATTRIBUTE and NF90_INQ_ATTNAME ...	50
5.5	Get Attribute's Values: NF90_GET_ATT	51
5.6	Copy Attribute from One NetCDF to Another: NF90_COPY_ATT	53
5.7	Rename an Attribute: NF90_RENAME_ATT	55
5.8	NF90_DEL_ATT	56
Appendix A	Appendix A - Summary of Fortran 90 Interface	59
Appendix B	Appendix B - FORTRAN 77 to Fortran 90 Transition Guide	63
	The new Fortran 90 interface	63
	Changes to Inquiry functions	63
	Changes to put and get function	63
Index	65

1 Use of the NetCDF Library

You can use the netCDF library without knowing about all of the netCDF interface. If you are creating a netCDF dataset, only a handful of routines are required to define the necessary dimensions, variables, and attributes, and to write the data to the netCDF dataset. (Even less are needed if you use the `ncgen` utility to create the dataset before running a program using netCDF library calls to write data. See [section “ncgen” in *NetCDF Users Guide*](#).) Similarly, if you are writing software to access data stored in a particular netCDF object, only a small subset of the netCDF library is required to open the netCDF dataset and access the data. Authors of generic applications that access arbitrary netCDF datasets need to be familiar with more of the netCDF library.

In this chapter we provide templates of common sequences of netCDF calls needed for common uses. For clarity we present only the names of routines; omit declarations and error checking; omit the type-specific suffixes of routine names for variables and attributes; indent statements that are typically invoked multiple times; and use `...` to represent arbitrary sequences of other statements. Full parameter lists are described in later chapters.

1.1 Creating a NetCDF Dataset

Here is a typical sequence of netCDF calls used to create a new netCDF dataset:

```

NF90_CREATE          ! create netCDF dataset: enter define mode
...
NF90_DEF_DIM        ! define dimensions: from name and length
...
NF90_DEF_VAR        ! define variables: from name, type, dims
...
NF90_PUT_ATT        ! assign attribute values
...
NF90_ENDDEF         ! end definitions: leave define mode
...
NF90_PUT_VAR        ! provide values for variable
...
NF90_CLOSE          ! close: save new netCDF dataset

```

Only one call is needed to create a netCDF dataset, at which point you will be in the first of two netCDF modes. When accessing an open netCDF dataset, it is either in define mode or data mode. In define mode, you can create dimensions, variables, and new attributes, but you cannot read or write variable data. In data mode, you can access data and change existing attributes, but you are not permitted to create new dimensions, variables, or attributes.

One call to `NF90_DEF_DIM` is needed for each dimension created. Similarly, one call to `NF90_DEF_VAR` is needed for each variable creation, and one call to a member of the `NF90_PUT_ATT` family is needed for each attribute defined and assigned a value. To leave define mode and enter data mode, call `NF90_ENDDEF`.

Once in data mode, you can add new data to variables, change old values, and change values of existing attributes (so long as the attribute changes do not require more storage space). Data of all types is written to a netCDF variable using the `NF90_PUT_VAR`

subroutine. Single values, arrays, or array sections may be supplied to `NF90_PUT_VAR`; optional arguments allow the writing of subsampled or mapped portions of the variable. (Subsampled and mapped access are general forms of data access that are explained later.)

Finally, you should explicitly close all netCDF datasets that have been opened for writing by calling `NF90_CLOSE`. By default, access to the file system is buffered by the netCDF library. If a program terminates abnormally with netCDF datasets open for writing, your most recent modifications may be lost. This default buffering of data is disabled by setting the `NF90_SHARE` flag when opening the dataset. But even if this flag is set, changes to attribute values or changes made in define mode are not written out until `NF90_SYNC` or `NF90_CLOSE` is called.

1.2 Reading a NetCDF Dataset with Known Names

Here we consider the case where you know the names of not only the netCDF datasets, but also the names of their dimensions, variables, and attributes. (Otherwise you would have to do "inquire" calls.) The order of typical C calls to read data from those variables in a netCDF dataset is:

```

NF90_OPEN           ! open existing netCDF dataset
...
NF90_INQ_DIMID      ! get dimension IDs
...
NF90_INQ_VARID      ! get variable IDs
...
NF90_GET_ATT        ! get attribute values
...
NF90_GET_VAR        ! get values of variables
...
NF90_CLOSE          ! close netCDF dataset

```

First, a single call opens the netCDF dataset, given the dataset name, and returns a netCDF ID that is used to refer to the open netCDF dataset in all subsequent calls.

Next, a call to `NF90_INQ_DIMID` for each dimension of interest gets the dimension ID from the dimension name. Similarly, each required variable ID is determined from its name by a call to `NF90_INQ_VARID`. Once variable IDs are known, variable attribute values can be retrieved using the netCDF ID, the variable ID, and the desired attribute name as input to `NF90_GET_ATT` for each desired attribute. Variable data values can be directly accessed from the netCDF dataset with calls to `NF90_GET_VAR`.

Finally, the netCDF dataset is closed with `NF90_CLOSE`. There is no need to close a dataset open only for reading.

1.3 Reading a netCDF Dataset with Unknown Names

It is possible to write programs (e.g., generic software) which do such things as processing every variable, without needing to know in advance the names of these variables. Similarly, the names of dimensions and attributes may be unknown.

Names and other information about netCDF objects may be obtained from netCDF datasets by calling inquire functions. These return information about a whole netCDF

dataset, a dimension, a variable, or an attribute. The following template illustrates how they are used:

```

NF90_OPEN                ! open existing netCDF dataset
...
NF90_INQUIRE            ! find out what is in it
...
NF90_INQUIRE_DIMENSION ! get dimension names, lengths
...
NF90_INQUIRE_VARIABLE  ! get variable names, types, shapes
...
NF90_INQ_ATTNAME        ! get attribute names
...
NF90_INQUIRE_ATTRIBUTE ! get other attribute information
...
NF90_GET_ATT            ! get attribute values
...
NF90_GET_VAR            ! get values of variables
...
NF90_CLOSE              ! close netCDF dataset

```

As in the previous example, a single call opens the existing netCDF dataset, returning a netCDF ID. This netCDF ID is given to the NF90_INQUIRE routine, which returns the number of dimensions, the number of variables, the number of global attributes, and the ID of the unlimited dimension, if there is one.

All the inquire functions are inexpensive to use and require no I/O, since the information they provide is stored in memory when a netCDF dataset is first opened.

Dimension IDs use consecutive integers, beginning at 1. Also dimensions, once created, cannot be deleted. Therefore, knowing the number of dimension IDs in a netCDF dataset means knowing all the dimension IDs: they are the integers 1, 2, 3, ...up to the number of dimensions. For each dimension ID, a call to the inquire function NF90_INQUIRE_DIMENSION returns the dimension name and length.

Variable IDs are also assigned from consecutive integers 1, 2, 3, ... up to the number of variables. These can be used in NF90_INQUIRE_VARIABLE calls to find out the names, types, shapes, and the number of attributes assigned to each variable.

Once the number of attributes for a variable is known, successive calls to NF90_INQ_ATTNAME return the name for each attribute given the netCDF ID, variable ID, and attribute number. Armed with the attribute name, a call to NF90_INQUIRE_ATTRIBUTE returns its type and length. Given the type and length, you can allocate enough space to hold the attribute values. Then a call to NF90_GET_ATT returns the attribute values.

Once the IDs and shapes of netCDF variables are known, data values can be accessed by calling NF90_GET_VAR.

1.4 Writing Data in an Existing NetCDF Dataset

With write access to an existing netCDF dataset, you can overwrite data values in existing variables or append more data to record variables along the unlimited (record) dimension.

To append more data to non-record variables requires changing the shape of such variables, which means creating a new netCDF dataset, defining new variables with the desired shape, and copying data. The netCDF data model was not designed to make such "schema changes" efficient or easy, so it is best to specify the shapes of variables correctly when you create a netCDF dataset, and to anticipate which variables will later grow by using the unlimited dimension in their definition.

The following code template lists a typical sequence of calls to overwrite some existing values and add some new records to record variables in an existing netCDF dataset with known variable names:

```

NF90_OPEN           ! open existing netCDF dataset
...
NF90_INQ_VARID      ! get variable IDs
...
NF90_PUT_VAR        ! provide new values for variables, if any
...
NF90_PUT_ATT        ! provide new values for attributes, if any
...
NF90_CLOSE          ! close netCDF dataset

```

A netCDF dataset is first opened by the `NF90_OPEN` call. This call puts the open dataset in data mode, which means existing data values can be accessed and changed, existing attributes can be changed, but no new dimensions, variables, or attributes can be added.

Next, calls to `NF90_INQ_VARID` get the variable ID from the name, for each variable you want to write. Then each call to `NF90_PUT_VAR` writes data into a specified variable, either a single value at a time, or a whole set of values at a time, depending on which variant of the interface is used. The calls used to overwrite values of non-record variables are the same as are used to overwrite values of record variables or append new data to record variables. The difference is that, with record variables, the record dimension is extended by writing values that don't yet exist in the dataset. This extends all record variables at once, writing "fill values" for record variables for which the data has not yet been written (but see [Section 4.9 \[Fill Values\]](#), page 42 to specify different behavior).

Calls to `NF90_PUT_ATT` may be used to change the values of existing attributes, although data that changes after a file is created is typically stored in variables rather than attributes.

Finally, you should explicitly close any netCDF datasets into which data has been written by calling `NF90_CLOSE` before program termination. Otherwise, modifications to the dataset may be lost.

1.5 Adding New Dimensions, Variables, Attributes

An existing netCDF dataset can be extensively altered. New dimensions, variables, and attributes can be added or existing ones renamed, and existing attributes can be deleted. Existing dimensions, variables, and attributes can be renamed. The following code template lists a typical sequence of calls to add new netCDF components to an existing dataset:

```

NF90_OPEN           ! open existing netCDF dataset
...

```

```

NF90_REDEF          ! put it into define mode
...
NF90_DEF_DIM       ! define additional dimensions (if any)
...
NF90_DEF_VAR       ! define additional variables (if any)
...
NF90_PUT_ATT       ! define other attributes (if any)
...
NF90_ENDDEF        ! check definitions, leave define mode
...
NF90_PUT_VAR       ! provide new variable values
...
NF90_CLOSE         ! close netCDF dataset

```

A netCDF dataset is first opened by the `NF90_OPEN` call. This call puts the open dataset in data mode, which means existing data values can be accessed and changed, existing attributes can be changed (so long as they do not grow), but nothing can be added. To add new netCDF dimensions, variables, or attributes you must enter define mode, by calling `NF90_REDEF`. In define mode, call `NF90_DEF_DIM` to define new dimensions, `NF90_DEF_VAR` to define new variables, and `NF90_PUT_ATT` to assign new attributes to variables or enlarge old attributes.

You can leave define mode and reenter data mode, checking all the new definitions for consistency and committing the changes to disk, by calling `NF90_ENDDEF`. If you do not wish to reenter data mode, just call `NF90_CLOSE`, which will have the effect of first calling `NF90_ENDDEF`.

Until the `NF90_ENDDEF` call, you may back out of all the redefinitions made in define mode and restore the previous state of the netCDF dataset by calling `NF90_ABORT`. You may also use the `NF90_ABORT` call to restore the netCDF dataset to a consistent state if the call to `NF90_ENDDEF` fails. If you have called `NF90_CLOSE` from definition mode and the implied call to `NF90_ENDDEF` fails, `NF90_ABORT` will automatically be called to close the netCDF dataset and leave it in its previous consistent state (before you entered define mode).

At most one process should have a netCDF dataset open for writing at one time. The library is designed to provide limited support for multiple concurrent readers with one writer, via disciplined use of the `NF90_SYNC` function and the `NF90_SHARE` flag. If a writer makes changes in define mode, such as the addition of new variables, dimensions, or attributes, some means external to the library is necessary to prevent readers from making concurrent accesses and to inform readers to call `NF90_SYNC` before the next access.

1.6 Error Handling

The netCDF library provides the facilities needed to handle errors in a flexible way. Each netCDF function returns an integer status value. If the returned status value indicates an error, you may handle it in any way desired, from printing an associated error message and exiting to ignoring the error indication and proceeding (not recommended!). For simplicity, the examples in this guide check the error status and call a separate function to handle any errors.

The `NF90_STRERROR` function is available to convert a returned integer error status into an error message string.

Occasionally, low-level I/O errors may occur in a layer below the netCDF library. For example, if a write operation causes you to exceed disk quotas or to attempt to write to a device that is no longer available, you may get an error from a layer below the netCDF library, but the resulting write error will still be reflected in the returned status value.

1.7 Compiling and Linking with the NetCDF Library

Details of how to compile and link a program that uses the netCDF C or FORTRAN interfaces differ, depending on the operating system, the available compilers, and where the netCDF library and include files are installed. Nevertheless, we provide here examples of how to compile and link a program that uses the netCDF library on a Unix platform, so that you can adjust these examples to fit your installation.

Every Fortran 90 procedure or module which references netCDF constants or procedures must have access to the module information created when the netCDF module was compiled. The suffix for this file depends on the compiler, but is often `.MOD`. Most Fortran 90 compilers do not allow you to specify an alternative location for this file as you might the location of external libraries. The simplest solution, therefore, is to create a symbolic link from the directory in which your code resides to the location of the pre-compiled netCDF module. For example:

```
ln -s /usr/local/netcdf/src/f90/netcdf.mod .
```

You may then compile source files which reference netCDF constants or procedures.

```
f90 -c mymodule.f90
```

Unless the netCDF library is installed in a standard directory where the linker always looks, you must use the `-L` and `-l` options to link an object file that uses the netCDF library. For example:

```
f90 -o myprogram myprogram.o -L/usr/local/netcdf/lib -lnetcdf
```

2 Datasets

2.1 Datasets Introduction

This chapter presents the interfaces of the netCDF functions that deal with a netCDF dataset or the whole netCDF library.

A netCDF dataset that has not yet been opened can only be referred to by its dataset name. Once a netCDF dataset is opened, it is referred to by a netCDF ID, which is a small nonnegative integer returned when you create or open the dataset. A netCDF ID is much like a file descriptor in C or a logical unit number in FORTRAN. In any single program, the netCDF IDs of distinct open netCDF datasets are distinct. A single netCDF dataset may be opened multiple times and will then have multiple distinct netCDF IDs; however at most one of the open instances of a single netCDF dataset should permit writing. When an open netCDF dataset is closed, the ID is no longer associated with a netCDF dataset.

Functions that deal with the netCDF library include:

- Get version of library.
- Get error message corresponding to a returned error code.

The operations supported on a netCDF dataset as a single object are:

- Create, given dataset name and whether to overwrite or not.
- Open for access, given dataset name and read or write intent.
- Put into define mode, to add dimensions, variables, or attributes.
- Take out of define mode, checking consistency of additions.
- Close, writing to disk if required.
- Inquire about the number of dimensions, number of variables, number of global attributes, and ID of the unlimited dimension, if any.
- Synchronize to disk to make sure it is current.
- Set and unset nofill mode for optimized sequential writes.
- After a summary of conventions used in describing the netCDF interfaces, the rest of this chapter presents a detailed description of the interfaces for these operations.

2.2 NetCDF Library Interface Descriptions

Each interface description for a particular netCDF function in this and later chapters contains:

- a description of the purpose of the function;
- a Fortran 90 interface block that presents the type and order of the formal parameters to the function;
- a description of each formal parameter in the C interface;
- a list of possible error conditions; and
- an example of a Fortran 90 program fragment calling the netCDF function (and perhaps other netCDF functions).

The examples follow a simple convention for error handling, always checking the error status returned from each netCDF function call and calling a `handle_error` function in case an error was detected. For an example of such a function, see Section 5.2 "Get error message corresponding to error status: `nc_strerror`".

2.3 NF90_STRERROR

The function `NF90_STRERROR` returns a static reference to an error message string corresponding to an integer netCDF error status or to a system error number, presumably returned by a previous call to some other netCDF function. The list of netCDF error status codes is available in the appropriate include file for each language binding.

Usage

```
function nf90_strerror(ncerr)
  integer, intent( in) :: ncerr
  character(len = 80)  :: nf90_strerror
```

NCERR An error status that might have been returned from a previous call to some netCDF function.

Errors

If you provide an invalid integer error status that does not correspond to any netCDF error message or to any system error message (as understood by the system `strerror` function), `NF90_STRERROR` returns a string indicating that there is no such error status.

Example

Here is an example of a simple error handling function that uses `NF90_STRERROR` to print the error message corresponding to the netCDF error status returned from any netCDF function call and then exit:

```
subroutine handle_err(status)
  integer, intent ( in) :: status

  if(status /= nf90_noerr) then
    print *, trim(nf90_strerror(status))
    stop "Stopped"
  end if
end subroutine handle_err
```

2.4 Get netCDF library version: NF90_INQ_LIBVERS

The function `NF90_INQ_LIBVERS` returns a string identifying the version of the netCDF library, and when it was built.

Usage

```
function nf90_inq_libvers()
  character(len = 80) :: nf90_inq_libvers
```

Errors

This function takes no arguments, and returns no error status.

Example

Here is an example using `nc_inq_libvers` to print the version of the netCDF library with which the program is linked:

```
print *, trim(nf90_inq_libvers())
```

2.5 NF90_CREATE

This function creates a new netCDF dataset, returning a netCDF ID that can subsequently be used to refer to the netCDF dataset in other netCDF function calls. The new netCDF dataset opened for write access and placed in define mode, ready for you to add dimensions, variables, and attributes.

A creation mode flag specifies whether to overwrite any existing dataset with the same name and whether access to the dataset is shared.

Usage

```
function nf90_create(path, cmode, ncid)
  character (len = *), intent(in  ) :: path
  integer,          intent(in  ) :: cmode
  integer, optional, intent(in  ) :: initialsize
  integer, optional, intent(inout) :: chunksize
  integer,          intent( out) :: ncid
  integer          :: nf90_create
```

path The file name of the new netCDF dataset.

cmode The creation mode flag. The following flags are available: `NF90_NO_CLOBBER`, `NF90_SHARE`, and `NF90_64BIT_OFFSET`.

Setting `NF90_NO_CLOBBER` means you do not want to clobber (overwrite) an existing dataset; an error (`NF90_EEXIST`) is returned if the specified dataset already exists.

The `NF90_SHARE` flag is appropriate when one process may be writing the dataset and one or more other processes reading the dataset concurrently; it means that dataset accesses are not buffered and caching is limited. Since the buffering scheme is optimized for sequential access, programs that do not access data sequentially may see some performance improvement by setting the `NF90_SHARE` flag.

Setting `NF90_64BIT_OFFSET` causes netCDF to create a 64-bit offset format file, instead of a netCDF classic format file. The 64-bit offset format imposes far fewer restrictions on very large (i.e. over 2 GB) data files. See [section “Large File Support” in *NetCDF Users’ Guide*](#).

A zero value (defined for convenience as `NF90_CLOBBER`) specifies the default behavior: overwrite any existing dataset with the same file name and buffer and

cache accesses for efficiency. The dataset will be in netCDF classic format. See section “NetCDF Classic Format Limitations” in *NetCDF Users’ Guide*.

ncid Returned netCDF ID.

The following optional arguments allow additional performance tuning.

initialsize

The initial size of the file (in bytes) at creation time. A value of 0 causes the file size to be computed when `nf90_enddef` is called.

chunksize

Controls a space versus time trade-off, memory allocated in the netcdf library versus number of system calls. Because of internal requirements, the value may not be set to exactly the value requested. The actual value chosen is returned.

The library chooses a system-dependent default value if `NF90_SIZEHINT_DEFAULT` is supplied as input. If the "preferred I/O block size" is available from the `stat()` system call as member `st_blksize` this value is used. Lacking that, twice the system pagesize is used. Lacking a call to discover the system pagesize, the default chunksize is set to 8192 bytes.

The chunksize is a property of a given open netcdf descriptor `ncid`, it is not a persistent property of the netcdf dataset.

Errors

`NF90_CREATE` returns the value `NF90_NOERR` if no errors occurred. Possible causes of errors include:

- Passing a dataset name that includes a directory that does not exist.
- Specifying a dataset name of a file that exists and also specifying `NF90_NOCLOBBER`.
- Specifying a meaningless value for the creation mode.
- Attempting to create a netCDF dataset in a directory where you don’t have permission to create files.

Example

In this example we create a netCDF dataset named `foo.nc`; we want the dataset to be created in the current directory only if a dataset with that name does not already exist:

```
use netcdf
implicit none
integer :: ncid, status
...
status = nf90_create(path = "foo.nc", cmode = nf90_noclobber, ncid = ncid)
if (status /= nf90_noerr) call handle_err(status)
```

2.6 NF90_OPEN

The function `NF90_OPEN` opens an existing netCDF dataset for access.

Usage

```
function nf90_open(path, mode, ncid, chunksize)
  character (len = *), intent(in  ) :: path
  integer,          intent(in  ) :: mode
  integer,          intent( out) :: ncid
  integer, optional, intent(inout) :: chunksize
  integer                               :: nf90_open
```

path File name for netCDF dataset to be opened.

omode A zero value (or `NF90_NOWRITE`) specifies the default behavior: open the dataset with read-only access, buffering and caching accesses for efficiency

Otherwise, the creation mode is `NF90_WRITE`, `NF90_SHARE`, or `NF90_WRITE|NF90_SHARE`. Setting the `NF90_WRITE` flag opens the dataset with read-write access. ("Writing" means any kind of change to the dataset, including appending or changing data, adding or renaming dimensions, variables, and attributes, or deleting attributes.) The `NF90_SHARE` flag is appropriate when one process may be writing the dataset and one or more other processes reading the dataset concurrently; it means that dataset accesses are not buffered and caching is limited. Since the buffering scheme is optimized for sequential access, programs that do not access data sequentially may see some performance improvement by setting the `NF90_SHARE` flag.

ncid Returned netCDF ID.

The following optional argument allows additional performance tuning.

chunksize

Controls a space versus time trade-off, memory allocated in the netcdf library versus number of system calls. Because of internal requirements, the value may not be set to exactly the value requested. The actual value chosen is returned.

The library chooses a system-dependent default value if `NF90_SIZEHINT_DEFAULT` is supplied as input. If the "preferred I/O block size" is available from the `stat()` system call as member `st_blksize` this value is used. Lacking that, twice the system pagesize is used. Lacking a call to discover the system pagesize, the default chunksize is set to 8192 bytes.

The chunksize is a property of a given open netcdf descriptor `ncid`, it is not a persistent property of the netcdf dataset.

Errors

`NF90_OPEN` returns the value `NF90_NOERR` if no errors occurred. Otherwise, the returned status indicates an error. Possible causes of errors include:

- The specified netCDF dataset does not exist.
- A meaningless mode was specified.

Example

Here is an example using `NF90_OPEN` to open an existing netCDF dataset named `foo.nc` for read-only, non-shared access:

```

use netcdf
implicit none
integer :: ncid, status
...
status = nf90_open(path = "foo.nc", cmode = nf90_nowrite, ncid = ncid)
if (status /= nf90_noerr) call handle_err(status)

```

2.7 NF90_REDEF

The function `NF90_REDEF` puts an open netCDF dataset into define mode, so dimensions, variables, and attributes can be added or renamed and attributes can be deleted.

Usage

```

function nf90_redef(ncid)
  integer, intent( in) :: ncid
  integer                :: nf90_redef

```

`ncid` netCDF ID, from a previous call to `NF90_OPEN` or `NF90_CREATE`.

Errors

`NF90_REDEF` returns the value `NF90_NOERR` if no errors occurred. Otherwise, the returned status indicates an error. Possible causes of errors include:

- The specified netCDF dataset is already in define mode.
- The specified netCDF dataset was opened for read-only.
- The specified netCDF ID does not refer to an open netCDF dataset.

Example

Here is an example using `NF90_REDEF` to open an existing netCDF dataset named `foo.nc` and put it into define mode:

```

use netcdf
implicit none
integer :: ncid, status
...
status = nf90_open("foo.nc", nf90_write, ncid) ! Open dataset
if (status /= nf90_noerr) call handle_err(status)
...
status = nf90_redef(ncid)                       ! Put the file in define mode
if (status /= nf90_noerr) call handle_err(status)

```

2.8 NF90_ENDDEF

The function `NF90_ENDDEF` takes an open netCDF dataset out of define mode. The changes made to the netCDF dataset while it was in define mode are checked and committed to disk if no problems occurred. Non-record variables may be initialized to a "fill value" as well (see [Section 2.13 \[NF90_SET_FILL\]](#), page 19). The netCDF dataset is then placed in data mode, so variable data can be read or written.

This call may involve copying data under some circumstances. For a more extensive discussion See [section "File Structure and Performance" in *NetCDF Users Guide*](#).

Usage

```
function nf90_enddef(ncid, h_minfree, v_align, v_minfree, r_align)
    integer,          intent( in) :: ncid
    integer, optional, intent( in) :: h_minfree, v_align, v_minfree, r_align
    integer          :: nf90_enddef
```

ncid NetCDF ID, from a previous call to `NF90_OPEN` or `NF90_CREATE`.

The following arguments allow additional performance tuning. Note: these arguments expose internals of the netcdf version 1 file format, and may not be available in future netcdf implementations.

The current netcdf file format has three sections: the "header" section, the data section for fixed size variables, and the data section for variables which have an unlimited dimension (record variables). The header begins at the beginning of the file. The index (offset) of the beginning of the other two sections is contained in the header. Typically, there is no space between the sections. This causes copying overhead to accrue if one wishes to change the size of the sections, as may happen when changing the names of things, text attribute values, adding attributes or adding variables. Also, for buffered i/o, there may be advantages to aligning sections in certain ways.

The following parameters allow one to control costs of future calls to `nf90_redef` or `nf90_enddef` by requesting that some space be available at the end of the section. The default value for both arguments is 0.

h_minfree

Size of the pad (in bytes) at the end of the "header" section.

v_minfree

Size of the pad (in bytes) at the end of the data section for fixed size variables.

The align parameters allow one to set the alignment of the beginning of the corresponding sections. The beginning of the section is rounded up to an index which is a multiple of the align parameter. The flag value `NF90_ALIGN_CHUNK` tells the library to use the chunksize (see above) as the align parameter. The default value for both arguments is 4 bytes.

v_align

The alignment of the beginning of the data section for fixed size variables.

r_align

The alignment of the beginning of the data section for variables which have an unlimited dimension (record variables).

Errors

NF90_ENDDEF returns the value NF90_NOERR if no errors occurred. Otherwise, the returned status indicates an error. Possible causes of errors include:

- The specified netCDF dataset is not in define mode.
- The specified netCDF ID does not refer to an open netCDF dataset.
- The size of one or more variables exceed the size constraints for whichever variant of the file format is in use). See section “Large File Support” in *The NetCDF Users Guide*.

Example

Here is an example using NF90_ENDDEF to finish the definitions of a new netCDF dataset named foo.nc and put it into data mode:

```

use netcdf
implicit none
integer :: ncid, status
...
status = nf90_create("foo.nc", nf90_noclobber, ncid)
if (status /= nf90_noerr) call handle_err(status)
... ! create dimensions, variables, attributes
status = nf90_enddef(ncid)
if (status /= nf90_noerr) call handle_err(status)

```

2.9 NF90_CLOSE

The function NF90_CLOSE closes an open netCDF dataset. If the dataset is in define mode, NF90_ENDDEF will be called before closing. (In this case, if NF90_ENDDEF returns an error, NF90_ABORT will automatically be called to restore the dataset to the consistent state before define mode was last entered.) After an open netCDF dataset is closed, its netCDF ID may be reassigned to the next netCDF dataset that is opened or created.

Usage

```

function nf90_close(ncid)
  integer, intent( in) :: ncid
  integer                :: nf90_close

```

ncid NetCDF ID, from a previous call to NF90_OPEN or NF90_CREATE.

Errors

NF90_CLOSE returns the value NF90_NOERR if no errors occurred. Otherwise, the returned status indicates an error. Possible causes of errors include:

- Define mode was entered and the automatic call made to NF90_ENDDEF failed.
- The specified netCDF ID does not refer to an open netCDF dataset.

Example

Here is an example using `NF90_CLOSE` to finish the definitions of a new netCDF dataset named `foo.nc` and release its netCDF ID:

```

use netcdf
implicit none
integer :: ncid, status
...
status = nf90_create("foo.nc", nf90_noclobber, ncid)
if (status /= nf90_noerr) call handle_err(status)
... ! create dimensions, variables, attributes
status = nf90_close(ncid)
if (status /= nf90_noerr) call handle_err(status)

```

2.10 NF90_INQUIRE Family

The `NF90_INQUIRE` subroutine returns information about an open netCDF dataset, given its netCDF ID. The subroutine can be called from either define mode or data mode, and returns values for any or all of the following: the number of dimensions, the number of variables, the number of global attributes, and the dimension ID of the dimension defined with unlimited length, if any. An additional function, `NF90_INQ_FORMAT`, returns the (rarely needed) format version.

No I/O is performed when `NF90_INQUIRE` is called, since the required information is available in memory for each open netCDF dataset.

Usage

```

function nf90_inquire(ncid, nDimensions, nVariables, nAttributes, &
                    unlimitedDimId, formatNum)
integer,          intent( in) :: ncid
integer, optional, intent(out) :: nDimensions, nVariables, &
                                nAttributes, unlimitedDimId, &
                                formatNum
integer          :: nf90_inquire

```

ncid NetCDF ID, from a previous call to `NF90_OPEN` or `NF90_CREATE`.

nDimensions

Returned number of dimensions defined for this netCDF dataset.

nVariables

Returned number of variables defined for this netCDF dataset.

nAttributes

Returned number of global attributes defined for this netCDF dataset.

unlimitedDimID

Returned ID of the unlimited dimension, if there is one for this netCDF dataset. If no unlimited length dimension has been defined, -1 is returned.

format

Returned integer indicating format version for this dataset, one of `nf90_format_classic`, `nf90_format_64bit`, `nf90_format_netcdf4`, or

`nf90_format_netcdf4_classic`. These are rarely needed by users or applications, since the library recognizes the format of a file it is accessing and handles it accordingly.

Errors

Function `NF90_INQUIRE` returns the value `NF90_NOERR` if no errors occurred. Otherwise, the returned status indicates an error. Possible causes of errors include:

- The specified netCDF ID does not refer to an open netCDF dataset.

Example

Here is an example using `NF90_INQUIRE` to find out about a netCDF dataset named `foo.nc`:

```

use netcdf
implicit none
integer :: ncid, status, nDims, nVars, nGlobalAtts, unlimDimID
...
status = nf90_open("foo.nc", nf90_nowrite, ncid)
if (status /= nf90_noerr) call handle_err(status)
...
status = nf90_inquire(ncid, nDims, nVars, nGlobalAtts, unlimdimid)
if (status /= nf90_noerr) call handle_err(status)
status = nf90_inquire(ncid, nDimensions = nDims, &
                      unlimitedDimID = unlimdimid)
if (status /= nf90_noerr) call handle_err(status)

```

2.11 NF90_SYNC

The function `NF90_SYNC` offers a way to synchronize the disk copy of a netCDF dataset with in-memory buffers. There are two reasons you might want to synchronize after writes:

- To minimize data loss in case of abnormal termination, or
- To make data available to other processes for reading immediately after it is written. But note that a process that already had the dataset open for reading would not see the number of records increase when the writing process calls `NF90_SYNC`; to accomplish this, the reading process must call `NF90_SYNC`.

This function is backward-compatible with previous versions of the netCDF library. The intent was to allow sharing of a netCDF dataset among multiple readers and one writer, by having the writer call `NF90_SYNC` after writing and the readers call `NF90_SYNC` before each read. For a writer, this flushes buffers to disk. For a reader, it makes sure that the next read will be from disk rather than from previously cached buffers, so that the reader will see changes made by the writing process (e.g., the number of records written) without having to close and reopen the dataset. If you are only accessing a small amount of data, it can be expensive in computer resources to always synchronize to disk after every write, since you are giving up the benefits of buffering.

An easier way to accomplish sharing (and what is now recommended) is to have the writer and readers open the dataset with the `NF90_SHARE` flag, and then it will not be

necessary to call `NF90_SYNC` at all. However, the `NF90_SYNC` function still provides finer granularity than the `NF90_SHARE` flag, if only a few netCDF accesses need to be synchronized among processes.

It is important to note that changes to the ancillary data, such as attribute values, are not propagated automatically by use of the `NF90_SHARE` flag. Use of the `NF90_SYNC` function is still required for this purpose.

Sharing datasets when the writer enters define mode to change the data schema requires extra care. In previous releases, after the writer left define mode, the readers were left looking at an old copy of the dataset, since the changes were made to a new copy. The only way readers could see the changes was by closing and reopening the dataset. Now the changes are made in place, but readers have no knowledge that their internal tables are now inconsistent with the new dataset schema. If netCDF datasets are shared across redefinition, some mechanism external to the netCDF library must be provided that prevents access by readers during redefinition and causes the readers to call `NF90_SYNC` before any subsequent access.

When calling `NF90_SYNC`, the netCDF dataset must be in data mode. A netCDF dataset in define mode is synchronized to disk only when `NF90_ENDDEF` is called. A process that is reading a netCDF dataset that another process is writing may call `NF90_SYNC` to get updated with the changes made to the data by the writing process (e.g., the number of records written), without having to close and reopen the dataset.

Data is automatically synchronized to disk when a netCDF dataset is closed, or whenever you leave define mode.

Usage

```
function nf90_sync(ncid)
  integer, intent( in) :: ncid
  integer                :: nf90_sync
```

`ncid` NetCDF ID, from a previous call to `NF90_OPEN` or `NF90_CREATE`.

Errors

`NF90_SYNC` returns the value `NF90_NOERR` if no errors occurred. Otherwise, the returned status indicates an error. Possible causes of errors include:

- The netCDF dataset is in define mode.
- The specified netCDF ID does not refer to an open netCDF dataset.

Example

Here is an example using `NF90_SYNC` to synchronize the disk writes of a netCDF dataset named `foo.nc`:

```
use netcdf
implicit none
integer :: ncid, status
...
status = nf90_open("foo.nc", nf90_write, ncid)
```

```

if (status /= nf90_noerr) call handle_err(status)
...
! write data or change attributes
...
status = NF90_SYNC(ncid)
if (status /= nf90_noerr) call handle_err(status)

```

2.12 NF90_ABORT

You no longer need to call this function, since it is called automatically by NF90_CLOSE in case the dataset is in define mode and something goes wrong with committing the changes. The function NF90_ABORT just closes the netCDF dataset, if not in define mode. If the dataset is being created and is still in define mode, the dataset is deleted. If define mode was entered by a call to NF90_REDEF, the netCDF dataset is restored to its state before definition mode was entered and the dataset is closed.

Usage

```

function nf90_abort(ncid)
  integer, intent( in) :: ncid
  integer                :: nf90_abort

```

ncid NetCDF ID, from a previous call to NF90_OPEN or NF90_CREATE.

Errors

NF90_ABORT returns the value NF90_NOERR if no errors occurred. Otherwise, the returned status indicates an error. Possible causes of errors include:

- When called from define mode while creating a netCDF dataset, deletion of the dataset failed.
- The specified netCDF ID does not refer to an open netCDF dataset.

Example

Here is an example using NF90_ABORT to back out of redefinitions of a dataset named foo.nc:

```

use netcdf
implicit none
integer :: ncid, status, LatDimID
...
status = nf90_open("foo.nc", nf90_write, ncid)
if (status /= nf90_noerr) call handle_err(status)
...
status = nf90_redef(ncid)
if (status /= nf90_noerr) call handle_err(status)
...
status = nf90_def_dim(ncid, "Lat", 18, LatDimID)
if (status /= nf90_noerr) then ! Dimension definition failed

```

```

    call handle_err(status)
    status = nf90_abort(ncid) ! Abort redefinitions
    if (status /= nf90_noerr) call handle_err(status)
end if
...

```

2.13 NF90_SET_FILL

This function is intended for advanced usage, to optimize writes under some circumstances described below. The function `NF90_SET_FILL` sets the fill mode for a netCDF dataset open for writing and returns the current fill mode in a return parameter. The fill mode can be specified as either `NF90_FILL` or `NF90_NOFILL`. The default behavior corresponding to `NF90_FILL` is that data is pre-filled with fill values, that is fill values are written when you create non-record variables or when you write a value beyond data that has not yet been written. This makes it possible to detect attempts to read data before it was written. See [Section 4.9 \[Fill Values\], page 42](#), for more information on the use of fill values. See [Section 5.2 \[Attribute Conventions\], page 45](#), for information about how to define your own fill values.

The behavior corresponding to `NF90_NOFILL` overrides the default behavior of pre-filling data with fill values. This can be used to enhance performance, because it avoids the duplicate writes that occur when the netCDF library writes fill values that are later overwritten with data.

A value indicating which mode the netCDF dataset was already in is returned. You can use this value to temporarily change the fill mode of an open netCDF dataset and then restore it to the previous mode.

After you turn on `NF90_NOFILL` mode for an open netCDF dataset, you must be certain to write valid data in all the positions that will later be read. Note that nofill mode is only a transient property of a netCDF dataset open for writing: if you close and reopen the dataset, it will revert to the default behavior. You can also revert to the default behavior by calling `NF90_SET_FILL` again to explicitly set the fill mode to `NF90_FILL`.

There are three situations where it is advantageous to set nofill mode:

1. Creating and initializing a netCDF dataset. In this case, you should set nofill mode before calling `NF90_ENDDEF` and then write completely all non-record variables and the initial records of all the record variables you want to initialize.
2. Extending an existing record-oriented netCDF dataset. Set nofill mode after opening the dataset for writing, then append the additional records to the dataset completely, leaving no intervening unwritten records.
3. Adding new variables that you are going to initialize to an existing netCDF dataset. Set nofill mode before calling `NF90_ENDDEF` then write all the new variables completely.

If the netCDF dataset has an unlimited dimension and the last record was written while in nofill mode, then the dataset may be shorter than if nofill mode was not set, but this will be completely transparent if you access the data only through the netCDF interfaces.

The use of this feature may not be available (or even needed) in future releases. Programmers are cautioned against heavy reliance upon this feature.

Usage

```
function nf90_set_fill(ncid, fillmode, old_mode)
  integer, intent( in) :: ncid, fillmode
  integer, intent(out) :: old_mode
  integer                :: nf90_set_fill
```

ncid NetCDF ID, from a previous call to `NF90_OPEN` or `NF90_CREATE`.

fillmode Desired fill mode for the dataset, either `NF90_NOFILL` or `NF90_FILL`.

old_mode Returned current fill mode of the dataset before this call, either `NF90_NOFILL` or `NF90_FILL`.

Errors

`NF90_SET_FILL` returns the value `NF90_NOERR` if no errors occurred. Otherwise, the returned status indicates an error. Possible causes of errors include:

- The specified netCDF ID does not refer to an open netCDF dataset.
- The specified netCDF ID refers to a dataset open for read-only access.
- The fill mode argument is neither `NF90_NOFILL` nor `NF90_FILL`.

Example

Here is an example using `NF90_SET_FILL` to set nofill mode for subsequent writes of a netCDF dataset named `foo.nc`:

```
use netcdf
implicit none
integer :: ncid, status, oldMode
...
status = nf90_open("foo.nc", nf90_write, ncid)
if (status /= nf90_noerr) call handle_err(status)
...
! Write data with prefilling behavior
...
status = nf90_set_fill(ncid, nf90_nofill, oldMode)
if (status /= nf90_noerr) call handle_err(status)
...
! Write data with no prefilling
...
```

3 Dimensions

3.1 Dimensions Introduction

Dimensions for a netCDF dataset are defined when it is created, while the netCDF dataset is in define mode. Additional dimensions may be added later by reentering define mode. A netCDF dimension has a name and a length. At most one dimension in a netCDF dataset can have the unlimited length, which means variables using this dimension can grow along this dimension.

There is a suggested limit (512) to the number of dimensions that can be defined in a single netCDF dataset. The limit is the value of the constant `NF90_MAX_DIMS`. The purpose of the limit is to make writing generic applications simpler. They need only provide an array of `NF90_MAX_DIMS` dimensions to handle any netCDF dataset. The implementation of the netCDF library does not enforce this advisory maximum, so it is possible to use more dimensions, if necessary, but netCDF utilities that assume the advisory maximums may not be able to handle the resulting netCDF datasets.

Ordinarily, the name and length of a dimension are fixed when the dimension is first defined. The name may be changed later, but the length of a dimension (other than the unlimited dimension) cannot be changed without copying all the data to a new netCDF dataset with a redefined dimension length.

A netCDF dimension in an open netCDF dataset is referred to by a small integer called a dimension ID. In the Fortran 90 interface, dimension IDs are 1, 2, 3, ..., in the order in which the dimensions were defined.

Operations supported on dimensions are:

- Create a dimension, given its name and length.
- Get a dimension ID from its name.
- Get a dimension's name and length from its ID.
- Rename a dimension.

3.2 NF90_DEF_DIM

The function `NF90_DEF_DIM` adds a new dimension to an open netCDF dataset in define mode. It returns (as an argument) a dimension ID, given the netCDF ID, the dimension name, and the dimension length. At most one unlimited length dimension, called the record dimension, may be defined for each netCDF dataset.

Usage

```
function nf90_def_dim(ncid, name, len, dimid)
  integer,          intent( in) :: ncid
  character (len = *), intent( in) :: name
  integer,          intent( in) :: len
  integer,          intent(out) :: dimid
  integer          :: nf90_def_dim
```

`ncid` NetCDF ID, from a previous call to `NF90_OPEN` or `NF90_CREATE`.

name	Dimension name. Must begin with an alphabetic character, followed by zero or more alphanumeric characters including the underscore ('_'). Case is significant.
len	Length of dimension; that is, number of values for this dimension as an index to variables that use it. This should be either a positive integer or the predefined constant <code>NF90_UNLIMITED</code> .
dimid	Returned dimension ID.

Errors

`NF90_DEF_DIM` returns the value `NF90_NOERR` if no errors occurred. Otherwise, the returned status indicates an error. Possible causes of errors include:

- The netCDF dataset is not in definition mode.
- The specified dimension name is the name of another existing dimension.
- The specified length is not greater than zero.
- The specified length is unlimited, but there is already an unlimited length dimension defined for this netCDF dataset.
- The specified netCDF ID does not refer to an open netCDF dataset.

Example

Here is an example using `NF90_DEF_DIM` to create a dimension named `lat` of length 18 and a unlimited dimension named `rec` in a new netCDF dataset named `foo.nc`:

```

use netcdf
implicit none
integer :: ncid, status, LatDimID, RecordDimID
...
status = nf90_create("foo.nc", nf90_noclobber, ncid)
if (status /= nf90_noerr) call handle_err(status)
...
status = nf90_def_dim(ncid, "Lat", 18, LatDimID)
if (status /= nf90_noerr) call handle_err(status)
status = nf90_def_dim(ncid, "Record", nf90_unlimited, RecordDimID)
if (status /= nf90_noerr) call handle_err(status)

```

3.3 NF90_INQ_DIMID

The function `NF90_INQ_DIMID` returns (as an argument) the ID of a netCDF dimension, given the name of the dimension. If `ndims` is the number of dimensions defined for a netCDF dataset, each dimension has an ID between 1 and `ndims`.

Usage

```

function nf90_inq_dimid(ncid, name, dimid)
integer,          intent( in) :: ncid
character (len = *), intent( in) :: name
integer,          intent(out) :: dimid
integer           :: nf90_inq_dimid

```

<code>ncid</code>	NetCDF ID, from a previous call to <code>NF90_OPEN</code> or <code>NF90_CREATE</code> .
<code>name</code>	Dimension name, a character string beginning with a letter and followed by any sequence of letters, digits, or underscore ('_') characters. Case is significant in dimension names.
<code>dimid</code>	Returned dimension ID.

Errors

`NF90_INQ_DIMID` returns the value `NF90_NOERR` if no errors occurred. Otherwise, the returned status indicates an error. Possible causes of errors include:

- The name that was specified is not the name of a dimension in the netCDF dataset.
- The specified netCDF ID does not refer to an open netCDF dataset.

Example

Here is an example using `NF90_INQ_DIMID` to determine the dimension ID of a dimension named `lat`, assumed to have been defined previously in an existing netCDF dataset named `foo.nc`:

```
use netcdf
implicit none
integer :: ncid, status, LatDimID
...
status = nf90_open("foo.nc", nf90_nowrite, ncid)
if (status /= nf90_noerr) call handle_err(status)
...
status = nf90_inq_dimid(ncid, "Lat", LatDimID)
if (status /= nf90_noerr) call handle_err(status)
```

3.4 NF90_INQUIRE_DIMENSION

This function information about a netCDF dimension. Information about a dimension includes its name and its length. The length for the unlimited dimension, if any, is the number of records written so far.

Usage

```
function nf90_inquire_dimension(ncid, dimid, name, len)
integer, intent(in) :: ncid, dimid
character(len = *), optional, intent(out) :: name
integer, optional, intent(out) :: len
integer :: nf90_inquire_dimension
```

<code>ncid</code>	NetCDF ID, from a previous call to <code>NF90_OPEN</code> or <code>NF90_CREATE</code> .
<code>dimid</code>	Dimension ID, from a previous call to <code>NF90_INQ_DIMID</code> or <code>NF90_DEF_DIM</code> .
<code>name</code>	Returned dimension name. The caller must allocate space for the returned name. The maximum possible length, in characters, of a dimension name is given by the predefined constant <code>NF90_MAX_NAME</code> .

len Returned length of dimension. For the unlimited dimension, this is the current maximum value used for writing any variables with this dimension, that is the maximum record number.

Errors

These functions return the value `NF90_NOERR` if no errors occurred. Otherwise, the returned status indicates an error. Possible causes of errors include:

- The dimension ID is invalid for the specified netCDF dataset.
- The specified netCDF ID does not refer to an open netCDF dataset.

Example

Here is an example using `NF90_INQ_DIM` to determine the length of a dimension named `lat`, and the name and current maximum length of the unlimited dimension for an existing netCDF dataset named `foo.nc`:

```

use netcdf
implicit none
integer :: ncid, status, LatDimID, RecordDimID
integer :: nLats, nRecords
character(len = nf90_max_name) :: RecordDimName
...
status = nf90_open("foo.nc", nf90_nowrite, ncid)
if (status /= nf90_noerr) call handle_err(status)
! Get ID of unlimited dimension
status = nf90_inquire(ncid, unlimitedDimId = RecordDimID)
if (status /= nf90_noerr) call handle_err(status)
...
status = nf90_inq_dimid(ncid, "Lat", LatDimID)
if (status /= nf90_noerr) call handle_err(status)
! How many values of "lat" are there?
status = nf90_inquire_dimension(ncid, LatDimID, len = nLats)
if (status /= nf90_noerr) call handle_err(status)
! What is the name of the unlimited dimension, how many records are there?
status = nf90_inquire_dimension(ncid, RecordDimID, &
                               name = RecordDimName, len = Records)
if (status /= nf90_noerr) call handle_err(status)

```

3.5 NF90_RENAME_DIM

The function `NF90_RENAME_DIM` renames an existing dimension in a netCDF dataset open for writing. If the new name is longer than the old name, the netCDF dataset must be in define mode. You cannot rename a dimension to have the same name as another dimension.

Usage

```
function nf90_rename_dim(ncid, dimid, name)
```

```

integer,          intent( in) :: ncid
character (len = *), intent( in) :: name
integer,          intent( in) :: dimid
integer          :: nf90_rename_dim

```

`ncid` NetCDF ID, from a previous call to `NF90_OPEN` or `NF90_CREATE`.

`dimid` Dimension ID, from a previous call to `NF90_INQ_DIMID` or `NF90_DEF_DIM`.

`name` New dimension name.

Errors

`NF90_RENAME_DIM` returns the value `NF90_NOERR` if no errors occurred. Otherwise, the returned status indicates an error. Possible causes of errors include:

- The new name is the name of another dimension.
- The dimension ID is invalid for the specified netCDF dataset.
- The specified netCDF ID does not refer to an open netCDF dataset.
- The new name is longer than the old name and the netCDF dataset is not in define mode.

Example

Here is an example using `NF90_RENAME_DIM` to rename the dimension `lat` to `latitude` in an existing netCDF dataset named `foo.nc`:

```

use netcdf
implicit none
integer :: ncid, status, LatDimID
...
status = nf90_open("foo.nc", nf90_write, ncid)
if (status /= nf90_noerr) call handle_err(status)
...
! Put in define mode so we can rename the dimension
status = nf90_redef(ncid)
if (status /= nf90_noerr) call handle_err(status)
! Get the dimension ID for "Lat"...
status = nf90_inq_dimid(ncid, "Lat", LatDimID)
if (status /= nf90_noerr) call handle_err(status)
! ... and change the name to "Latitude".
status = nf90_rename_dim(ncid, LatDimID, "Latitude")
if (status /= nf90_noerr) call handle_err(status)
! Leave define mode
status = nf90_enddef(ncid)
if (status /= nf90_noerr) call handle_err(status)

```


4 Variables

4.1 Variables Introduction

Variables for a netCDF dataset are defined when the dataset is created, while the netCDF dataset is in define mode. Other variables may be added later by reentering define mode. A netCDF variable has a name, a type, and a shape, which are specified when it is defined. A variable may also have values, which are established later in data mode.

Ordinarily, the name, type, and shape are fixed when the variable is first defined. The name may be changed, but the type and shape of a variable cannot be changed. However, a variable defined in terms of the unlimited dimension can grow without bound in that dimension.

A netCDF variable in an open netCDF dataset is referred to by a small integer called a variable ID.

Variable IDs reflect the order in which variables were defined within a netCDF dataset. Variable IDs are 1, 2, 3,..., in the order in which the variables were defined. A function is available for getting the variable ID from the variable name and vice-versa.

Attributes [Chapter 5 \[Attributes\], page 45](#) may be associated with a variable to specify such properties as units.

Operations supported on variables are:

- Create a variable, given its name, data type, and shape.
- Get a variable ID from its name.
- Get a variable's name, data type, shape, and number of attributes from its ID.
- Put a data value into a variable, given variable ID, indices, and value.
- Put an array of values into a variable, given variable ID, corner indices, edge lengths, and a block of values.
- Put a subsampled or mapped array-section of values into a variable, given variable ID, corner indices, edge lengths, stride vector, index mapping vector, and a block of values.
- Get a data value from a variable, given variable ID and indices.
- Get an array of values from a variable, given variable ID, corner indices, and edge lengths.
- Get a subsampled or mapped array-section of values from a variable, given variable ID, corner indices, edge lengths, stride vector, and index mapping vector.
- Rename a variable.

4.2 Language Types Corresponding to netCDF external data types

The following table gives the netCDF external data types and the corresponding type constants for defining variables in the FORTRAN interface:

Type	FORTRAN API Mnemonic	Bits
byte	NF90_BYTE	8

char	NF90_CHAR	8
short	NF90_SHORT	16
int	NF90_INT	32
float	NF90_FLOAT	32
double	NF90_DOUBLE	64

The first column gives the netCDF external data type, which is the same as the CDL data type. The next column gives the corresponding Fortran 90 parameter for use in netCDF functions (the parameters are defined in the netCDF Fortran 90 module `netcdf.f90`). The last column gives the number of bits used in the external representation of values of the corresponding type.

Note that there are no netCDF types corresponding to 64-bit integers or to characters wider than 8 bits in the current version of the netCDF library.

4.3 Create a Variable: NF90_DEF_VAR

The function `NF90_DEF_VAR` adds a new variable to an open netCDF dataset in define mode. It returns (as an argument) a variable ID, given the netCDF ID, the variable name, the variable type, the number of dimensions, and a list of the dimension IDs.

Usage

```
function nf90_def_var(ncid, name, xtype, dimids, varid)
  integer,          intent( in) :: ncid
  character (len = *), intent( in) :: name
  integer,          intent( in) :: xtype
  integer, dimension(:), intent( in) :: dimids
  integer,          intent(out) :: varid
  integer          :: nf90_def_var
```

<code>ncid</code>	NetCDF ID, from a previous call to <code>NF90_OPEN</code> or <code>NF90_CREATE</code> .
<code>name</code>	Variable name. Must begin with an alphabetic character, followed by zero or more alphanumeric characters including the underscore ('_'). Case is significant.
<code>xtype</code>	One of the set of predefined netCDF external data types. The type of this parameter, <code>NF90_TYPE</code> , is defined in the netCDF header file. The valid netCDF external data types are <code>NF90_BYTE</code> , <code>NF90_CHAR</code> , <code>NF90_SHORT</code> , <code>NF90_INT</code> , <code>NF90_FLOAT</code> , and <code>NF90_DOUBLE</code> .
<code>dimids</code>	Vector of dimension IDs corresponding to the variable dimensions. For example, a vector of 2 dimension IDs specifies a matrix, 1 specifies a vector, and 0 means the variable is a scalar with no dimensions. If the ID of the unlimited dimension is included, it must be last. This argument is optional, and if absent specifies a scalar with no dimensions.
<code>varid</code>	Returned variable ID.

Errors

NF90_DEF_VAR returns the value NF90_NOERR if no errors occurred. Otherwise, the returned status indicates an error. Possible causes of errors include:

- The netCDF dataset is not in define mode.
- The specified variable name is the name of another existing variable.
- The specified type is not a valid netCDF type.
- The specified number of dimensions is negative or more than the constant NF90_MAX_VAR_DIMS, the maximum number of dimensions permitted for a netCDF variable.
- One or more of the dimension IDs in the list of dimensions is not a valid dimension ID for the netCDF dataset.
- The number of variables would exceed the constant NF90_MAX_VARS, the maximum number of variables permitted in a netCDF dataset.
- The specified netCDF ID does not refer to an open netCDF dataset.

Example

Here is an example using NF90_DEF_VAR to create a variable named rh of type double with three dimensions, time, lat, and lon in a new netCDF dataset named foo.nc:

```

use netcdf
implicit none
integer :: status, ncid
integer :: LonDimId, LatDimId, TimeDimId
integer :: RhVarId
...
status = nf90_create("foo.nc", nf90_NoClobber, ncid)
if(status /= nf90_NoErr) call handle_error(status)
...
! Define the dimensions
status = nf90_def_dim(ncid, "lat", 5, LatDimId)
if(status /= nf90_NoErr) call handle_error(status)
status = nf90_def_dim(ncid, "lon", 10, LonDimId)
if(status /= nf90_NoErr) call handle_error(status)
status = nf90_def_dim(ncid, "time", nf90_unlimited, TimeDimId)
if(status /= nf90_NoErr) call handle_error(status)
...
! Define the variable
status = nf90_def_var(ncid, "rh", nf90_double, &
                    (/ LonDimId, LatDimID, TimeDimID /), RhVarId)
if(status /= nf90_NoErr) call handle_error(status)

```

4.4 Get a Variable ID from Its Name: NF90_INQ_VARID

The function NF90_INQ_VARID returns the ID of a netCDF variable, given its name.

Usage

```
function nf90_inq_varid(ncid, name, varid)
  integer,          intent( in) :: ncid
  character (len = *), intent( in) :: name
  integer,          intent(out) :: varid
  integer           :: nf90_inq_varid
```

ncid NetCDF ID, from a previous call to NF90_OPEN or NF90_CREATE.

name Variable name for which ID is desired.

varid Returned variable ID.

Errors

NF90_INQ_VARID returns the value NF90_NOERR if no errors occurred. Otherwise, the returned status indicates an error. Possible causes of errors include:

- The specified variable name is not a valid name for a variable in the specified netCDF dataset.
- The specified netCDF ID does not refer to an open netCDF dataset.

Example

Here is an example using NF90_INQ_VARID to find out the ID of a variable named rh in an existing netCDF dataset named foo.nc:

```
      use netcdf
      implicit none
      integer :: status, ncid, RhVarId
      ...
      status = nf90_open("foo.nc", nf90_NoWrite, ncid)
      if(status /= nf90_NoErr) call handle_err(status)
      ...
      status = nf90_inq_varid(ncid, "rh", RhVarId)
      if(status /= nf90_NoErr) call handle_err(status)
```

4.5 Get Information about a Variable from Its ID: NF90_INQUIRE_VARIABLE

NF90_INQUIRE_VARIABLE returns information about a netCDF variable given its ID. Information about a variable includes its name, type, number of dimensions, a list of dimension IDs describing the shape of the variable, and the number of variable attributes that have been assigned to the variable.

Usage

```
function nf90_inquire_variable(ncid, varid, name, xtype, ndims, dimids, nAtts)
  integer,          intent( in) :: ncid, varid
  character (len = *), optional, intent(out) :: name
  integer,          optional, intent(out) :: xtype, ndims
```

	<code>integer, dimension(*), optional, intent(out) :: dimids</code>
	<code>integer, optional, intent(out) :: nAtts</code>
	<code>integer :: nf90_inquire_variable</code>
<code>ncid</code>	NetCDF ID, from a previous call to <code>NF90_OPEN</code> or <code>NF90_CREATE</code> .
<code>varid</code>	Variable ID.
<code>name</code>	Returned variable name. The caller must allocate space for the returned name. The maximum possible length, in characters, of a variable name is given by the predefined constant <code>NF90_MAX_NAME</code> .
<code>xtype</code>	Returned variable type, one of the set of predefined netCDF external data types. The type of this parameter, <code>NF90_TYPE</code> , is defined in the netCDF header file. The valid netCDF external data types are <code>NF90_BYTE</code> , <code>NF90_CHAR</code> , <code>NF90_SHORT</code> , <code>NF90_INT</code> , <code>NF90_FLOAT</code> , AND <code>NF90_DOUBLE</code> .
<code>ndims</code>	Returned number of dimensions the variable was defined as using. For example, 2 indicates a matrix, 1 indicates a vector, and 0 means the variable is a scalar with no dimensions.
<code>dimids</code>	Returned vector of <code>*ndimsp</code> dimension IDs corresponding to the variable dimensions. The caller must allocate enough space for a vector of at least <code>*ndimsp</code> integers to be returned. The maximum possible number of dimensions for a variable is given by the predefined constant <code>NF90_MAX_VAR_DIMS</code> .
<code>natts</code>	Returned number of variable attributes assigned to this variable.

These functions return the value `NF90_NOERR` if no errors occurred. Otherwise, the returned status indicates an error. Possible causes of errors include:

- The variable ID is invalid for the specified netCDF dataset.
- The specified netCDF ID does not refer to an open netCDF dataset.

Example

Here is an example using `NF90_INQ_VAR` to find out about a variable named `rh` in an existing netCDF dataset named `foo.nc`:

```

use netcdf
implicit none
integer :: status, ncid, &
         RhVarId &
         numDims, numAtts
integer, dimension(nf90_max_var_dims) :: rhDimIds
...
status = nf90_open("foo.nc", nf90_NoWrite, ncid)
if(status /= nf90_NoErr) call handle_error(status)
...
status = nf90_inq_varid(ncid, "rh", RhVarId)
if(status /= nf90_NoErr) call handle_err(status)
status = nf90_inquire_variable(ncid, RhVarId, ndims = numDims, natts = numAtts)
if(status /= nf90_NoErr) call handle_err(status)

```

```
status = nf90_inquire_variable(ncid, RhVarId, dimids = rhDimIds(:numDims))
if(status /= nf90_NoErr) call handle_err(status)
```

4.6 Writing Data Values: NF90_PUT_VAR

The function `NF90_PUT_VAR` puts one or more data values into the variable of an open netCDF dataset that is in data mode. Required inputs are the netCDF ID, the variable ID, and one or more data values. Optional inputs may indicate the starting position of the data values in the netCDF variable (argument `start`), the sampling frequency with which data values are written into the netCDF variable (argument `stride`), and a mapping between the dimensions of the data array and the netCDF variable (argument `map`). The values to be written are associated with the netCDF variable by assuming that the first dimension of the netCDF variable varies fastest in the Fortran 90 interface. Data values converted to the external type of the variable, if necessary.

Take care when using the simplest forms of this interface with record variables when you don't specify how many records are to be written. If you try to write all the values of a record variable into a netCDF file that has no record data yet (hence has 0 records), nothing will be written. Similarly, if you try to write all of a record variable but there are more records in the file than you assume, more data may be written to the file than you supply, which may result in a segmentation violation.

Usage

```
function nf90_put_var(ncid, varid, values, start, count, stride, map)
  integer,                                intent( in) :: ncid, varid
  any valid type, scalar or array of any rank, &
                                          intent( in) :: values
  integer, dimension(:), optional, intent( in) :: start, count, stride, map
  integer                                  :: nf90_put_var
```

ncid NetCDF ID, from a previous call to `NF90_OPEN` or `NF90_CREATE`.

varid Variable ID.

values The data value(s) to be written. The data may be of any type, and may be a scalar or an array of any rank. You cannot put `CHARACTER` data into a numeric variable or numeric data into a text variable. For numeric data, if the type of data differs from the netCDF variable type, type conversion will occur. See section “Type Conversion” in *NetCDF Users Guide*.

start A vector of integers specifying the index in the variable where the first (or only) of the data values will be written. The indices are relative to 1, so for example, the first data value of a variable would have index (1, 1, ..., 1). The elements of `start` correspond, in order, to the variable's dimensions. Hence, if the variable is a record variable, the last index would correspond to the starting record number for writing the data values.

By default, `start(:) = 1`.

count A vector of integers specifying the number of indices selected along each dimension. To write a single value, for example, specify `count` as (1, 1, ..., 1). The

elements of count correspond, in order, to the variable's dimensions. Hence, if the variable is a record variable, the last element of count corresponds to a count of the number of records to write.

By default, `count(:numDims) = shape(values)` and `count(numDims + 1:) = 1`, where `numDims = size(shape(values))`.

stride A vector of integers that specifies the sampling interval along each dimension of the netCDF variable. The elements of the stride vector correspond, in order, to the netCDF variable's dimensions (`stride(1)` gives the sampling interval along the most rapidly varying dimension of the netCDF variable). Sampling intervals are specified in type-independent units of elements (a value of 1 selects consecutive elements of the netCDF variable along the corresponding dimension, a value of 2 selects every other element, etc.).

By default, `stride(:) = 1`.

imap A vector of integers that specifies the mapping between the dimensions of a netCDF variable and the in-memory structure of the internal data array. The elements of the index mapping vector correspond, in order, to the netCDF variable's dimensions (`map(1)` gives the distance between elements of the internal array corresponding to the most rapidly varying dimension of the netCDF variable). Distances between elements are specified in units of elements.

By default, `edgeLengths = shape(values)`, and `map = (/ 1, (product(edgeLengths(:i)), i = 1, size(edgeLengths) - 1) /)`, that is, there is no mapping.

Use of Fortran 90 intrinsic functions (including `reshape`, `transpose`, and `spread`) may let you avoid using this argument.

Errors

`NF90_PUT_VAR1_` type returns the value `NF90_NOERR` if no errors occurred. Otherwise, the returned status indicates an error. Possible causes of errors include:

- The variable ID is invalid for the specified netCDF dataset.
- The specified indices were out of range for the rank of the specified variable. For example, a negative index or an index that is larger than the corresponding dimension length will cause an error.
- The specified value is out of the range of values representable by the external data type of the variable.
- The specified netCDF is in define mode rather than data mode.
- The specified netCDF ID does not refer to an open netCDF dataset.

Example

Here is an example using `NF90_PUT_VAR` to set the (4,3,2) element of the variable named `rh` to 0.5 in an existing netCDF dataset named `foo.nc`. For simplicity in this example, we assume that we know that `rh` is dimensioned with `lon`, `lat`, and `time`, so we want to set the value of `rh` that corresponds to the fourth `lon` value, the third `lat` value, and the second `time` value:

```

use netcdf
implicit none
integer :: ncId, rhVarId, status
...
status = nf90_open("foo.nc", nf90_Write, ncId)
if(status /= nf90_NoErr) call handle_err(status)
...
status = nf90_inq_varid(ncId, "rh", rhVarId)
if(status /= nf90_NoErr) call handle_err(status)
status = nf90_put_var(ncId, rhVarId, 0.5, start = (/ 4, 3, 2 /) )
if(status /= nf90_NoErr) call handle_err(status)

```

In this example we use NF90_PUT_VAR to add or change all the values of the variable named rh to 0.5 in an existing netCDF dataset named foo.nc. We assume that we know that rh is dimensioned with lon, lat, and time. In this example we query the netCDF file to discover the lengths of the dimensions, then use the Fortran 90 intrinsic function reshape to create a temporary array of data values which is the same shape as the netCDF variable.

```

use netcdf
implicit none
integer :: ncId, rhVarId, status, &
lonDimID, latDimID, timeDimID, &
numLons, numLats, numTimes, &
i
integer, dimension(nf90_max_var_dims) :: dimIDs
...
status = nf90_open("foo.nc", nf90_Write, ncId)
if(status /= nf90_NoErr) call handle_err(status)
...
status = nf90_inq_varid(ncId, "rh", rhVarId)
if(status /= nf90_NoErr) call handle_err(status)
! How big is the netCDF variable, that is, what are the lengths of
! its constituent dimensions?
status = nf90_inquire_variable(ncId, rhVarId, dimids = dimIDs)
if(status /= nf90_NoErr) call handle_err(status)
status = nf90_inquire_dimension(ncId, dimIDs(1), len = numLons)
if(status /= nf90_NoErr) call handle_err(status)
status = nf90_inquire_dimension(ncId, dimIDs(2), len = numLats)
if(status /= nf90_NoErr) call handle_err(status)
status = nf90_inquire_dimension(ncId, dimIDs(3), len = numTimes)
if(status /= nf90_NoErr) call handle_err(status)
...
! Make a temporary array the same shape as the netCDF variable.
status = nf90_put_var(ncId, rhVarId, &
reshape( &
(/ (0.5, i = 1, numLons * numLats * numTimes) /) , &
shape = (/ numLons, numLats, numTimes /) )
if(status /= nf90_NoErr) call handle_err(status)

```

Here is an example using `NF90_PUT_VAR` to add or change a section of the variable named `rh` to 0.5 in an existing netCDF dataset named `foo.nc`. For simplicity in this example, we assume that we know that `rh` is dimensioned with `lon`, `lat`, and `time`, that there are ten `lon` values, five `lat` values, and three `time` values, and that we want to replace all the values at the last time.

```

use netcdf
implicit none
integer          :: ncId, rhVarId, status
integer, parameter :: numLons = 10, numLats = 5, numTimes = 3
real, dimension(numLons, numLats) &
               :: rhValues
...
status = nf90_open("foo.nc", nf90_Write, ncId)
if(status /= nf90_NoErr) call handle_err(status)
...
status = nf90_inq_varid(ncId, "rh", rhVarId)
if(status /= nf90_NoErr) call handle_err(status)
! Fill in all values at the last time
rhValues(:, :) = 0.5
status = nf90_put_var(ncId, rhVarId, rhValues,          &
                    start = (/ 1, 1, numTimes /), &
                    count = (/ numLats, numLons, 1 /))
if(status /= nf90_NoErr) call handle_err(status)

```

Here is an example of using `NF_PUT_VAR` to write every other point of a netCDF variable named `rh` having dimensions (6, 4).

```

use netcdf
implicit none
integer          :: ncId, rhVarId, status
integer, parameter :: numLons = 6, numLats = 4
real, dimension(numLons, numLats) &
               :: rhValues = 0.5
...
status = nf90_open("foo.nc", nf90_Write, ncId)
if(status /= nf90_NoErr) call handle_err(status)
...
status = nf90_inq_varid(ncId, "rh", rhVarId)
if(status /= nf90_NoErr) call handle_err(status)
...
! Fill in every other value using an array section
status = nf90_put_var(ncId, rhVarId, rhValues(::2, ::2), &
                    stride = (/ 2, 2 /))
if(status /= nf90_NoErr) call handle_err(status)

```

The following map vector shows the default mapping between a 2x3x4 netCDF variable and an internal array of the same shape:

```

real,    dimension(2, 3, 4):: a ! same shape as netCDF variable
integer, dimension(3)      :: map = (/ 1, 2, 6 /)

```

```

! netCDF dimension inter-element distance
! -----
! most rapidly varying      1
! intermediate              2 (= map(1)*2)
! most slowly varying      6 (= map(2)*3)

```

Using the map vector above obtains the same result as simply not passing a map vector at all.

Here is an example of using `nf90_put_var` to write a netCDF variable named `rh` whose dimensions are the transpose of the Fortran 90 array:

```

use netcdf
implicit none
integer                :: ncId, rhVarId, status
integer, parameter    :: numLons = 6, numLats = 4
real, dimension(numLons, numLats) :: rhValues
! netCDF variable has dimensions (numLats, numLons)
...
status = nf90_open("foo.nc", nf90_Write, ncId)
if(status /= nf90_NoErr) call handle_err(status)
...
status = nf90_inq_varid(ncId, "rh", rhVarId)
if(status /= nf90_NoErr) call handle_err(status)
...
!Write transposed values: map vector would be (/ 1, numLats /) for
! no transposition
status = nf90_put_var(ncId, rhVarId, rhValues, map = (/ numLons, 1 /))
if(status /= nf90_NoErr) call handle_err(status)

```

The same effect can be obtained more simply using Fortran 90 intrinsic functions:

```

use netcdf
implicit none
integer                :: ncId, rhVarId, status
integer, parameter    :: numLons = 6, numLats = 4
real, dimension(numLons, numLats) :: rhValues
! netCDF variable has dimensions (numLats, numLons)
...
status = nf90_open("foo.nc", nf90_Write, ncId)
if(status /= nf90_NoErr) call handle_err(status)
...
status = nf90_inq_varid(ncId, "rh", rhVarId)
if(status /= nf90_NoErr) call handle_err(status)
...
status = nf90_put_var(ncId, rhVarId, transpose(rhValues))
if(status /= nf90_NoErr) call handle_err(status)

```

4.7 Reading Data Values: `NF90_GET_VAR`

The function `NF90_GET_VAR` gets one or more data values from a netCDF variable of an open netCDF dataset that is in data mode. Required inputs are the netCDF ID, the variable ID, and a specification for the data values into which the data will be read. Optional inputs may indicate the starting position of the data values in the netCDF variable (argument `start`), the sampling frequency with which data values are read from the netCDF variable (argument `stride`), and a mapping between the dimensions of the data array and the netCDF variable (argument `map`). The values to be read are associated with the netCDF variable by assuming that the first dimension of the netCDF variable varies fastest in the Fortran 90 interface. Data values are converted from the external type of the variable, if necessary.

Take care when using the simplest forms of this interface with record variables when you don't specify how many records are to be read. If you try to read all the values of a record variable into an array but there are more records in the file than you assume, more data will be read than you expect, which may cause a segmentation violation.

Usage

```
function nf90_get_var(ncid, varid, values, start, count, stride, map)
  integer,                                intent( in) :: ncid, varid
  any valid type, scalar or array of any rank, &
                                          intent(out) :: values
  integer, dimension(:), optional, intent( in) :: start, count, stride, map
  integer                                  :: nf90_get_var
```

ncid NetCDF ID, from a previous call to `NF90_OPEN` or `NF90_CREATE`.

varid Variable ID.

values The data value(s) to be read. The data may be of any type, and may be a scalar or an array of any rank. You cannot read `CHARACTER` data from a numeric variable or numeric data from a text variable. For numeric data, if the type of data differs from the netCDF variable type, type conversion will occur. See section “Type Conversion” in *NetCDF Users Guide*.

start A vector of integers specifying the index in the variable from which the first (or only) of the data values will be read. The indices are relative to 1, so for example, the first data value of a variable would have index (1, 1, ..., 1). The elements of `start` correspond, in order, to the variable's dimensions. Hence, if the variable is a record variable, the last index would correspond to the starting record number for writing the data values.
By default, `start(:) = 1`.

count A vector of integers specifying the number of indices selected along each dimension. To read a single value, for example, specify `count` as (1, 1, ..., 1). The elements of `count` correspond, in order, to the variable's dimensions. Hence, if the variable is a record variable, the last element of `count` corresponds to a count of the number of records to read.
By default, `count(:numDims) = shape(values)` and `count(numDims + 1:) = 1`, where `numDims = size(shape(values))`.

- stride** A vector of integers that specifies the sampling interval along each dimension of the netCDF variable. The elements of the stride vector correspond, in order, to the netCDF variable's dimensions (stride(1) gives the sampling interval along the most rapidly varying dimension of the netCDF variable). Sampling intervals are specified in type-independent units of elements (a value of 1 selects consecutive elements of the netCDF variable along the corresponding dimension, a value of 2 selects every other element, etc.).
By default, stride(:) = 1.
- map** A vector of integers that specifies the mapping between the dimensions of a netCDF variable and the in-memory structure of the internal data array. The elements of the index mapping vector correspond, in order, to the netCDF variable's dimensions (map(1) gives the distance between elements of the internal array corresponding to the most rapidly varying dimension of the netCDF variable). Distances between elements are specified in units of elements.
By default, edgeLengths = shape(values), and map = (/ 1, (product(edgeLengths(:i)), i = 1, size(edgeLengths) - 1) /), that is, there is no mapping.
Use of Fortran 90 intrinsic functions (including reshape, transpose, and spread) may let you avoid using this argument.

Errors

NF90_GET_VAR returns the value NF90_NOERR if no errors occurred. Otherwise, the returned status indicates an error. Possible causes of errors include:

- The variable ID is invalid for the specified netCDF dataset.
- The assumed or specified start, count, and stride generate an index which is out of range. Note that no error checking is possible on the map vector.
- One or more of the specified values are out of the range of values representable by the desired type.
- The specified netCDF is in define mode rather than data mode.
- The specified netCDF ID does not refer to an open netCDF dataset.

(As noted above, another possible source of error is using this interface to read all the values of a record variable without specifying the number of records. If there are more records in the file than you assume, more data will be read than you expect!)

Example

Here is an example using NF90_GET_VAR to read the (4,3,2) element of the variable named rh from an existing netCDF dataset named foo.nc. For simplicity in this example, we assume that we know that rh is dimensioned with lon, lat, and time, so we want to read the value of rh that corresponds to the fourth lon value, the third lat value, and the second time value:

```
use netcdf
implicit none
integer :: ncId, rhVarId, status
```

```

real    :: rhValue
...
status = nf90_open("foo.nc", nf90_NoWrite, ncid)
if(status /= nf90_NoErr) call handle_err(status)
-
status = nf90_inq_varid(ncid, "rh", rhVarId)
if(status /= nf90_NoErr) call handle_err(status)
status = nf90_get_var(ncid, rhVarId, rhValue, start = (/ 4, 3, 2 /) )
if(status /= nf90_NoErr) call handle_err(status)

```

In this example we use `NF90_GET_VAR` to read all the values of the variable named `rh` from an existing netCDF dataset named `foo.nc`. We assume that we know that `rh` is dimensioned with `lon`, `lat`, and `time`. In this example we query the netCDF file to discover the lengths of the dimensions, then allocate a Fortran 90 array the same shape as the netCDF variable.

```

use netcdf
implicit none
integer                                :: ncId, rhVarId, &
                                       lonDimID, latDimId, timeDimId, &
                                       numLons, numLats, numTimes, &
                                       status
integer, dimension(nf90_max_var_dims) :: dimIDs
real, dimension(:, :, :), allocatable :: rhValues
...
status = nf90_open("foo.nc", nf90_NoWrite, ncid)
if(status /= nf90_NoErr) call handle_err(status)
...
status = nf90_inq_varid(ncid, "rh", rhVarId)
if(status /= nf90_NoErr) call handle_err(status)
! How big is the netCDF variable, that is, what are the lengths of
!   its constituent dimensions?
status = nf90_inquire_variable(ncid, rhVarId, dimids = dimIDs)
if(status /= nf90_NoErr) call handle_err(status)
status = nf90_inquire_dimension(ncid, dimIDs(1), len = numLons)
if(status /= nf90_NoErr) call handle_err(status)
status = nf90_inquire_dimension(ncid, dimIDs(2), len = numLats)
if(status /= nf90_NoErr) call handle_err(status)
status = nf90_inquire_dimension(ncid, dimIDs(3), len = numTimes)
if(status /= nf90_NoErr) call handle_err(status)
allocate(rhValues(numLons, numLats, numTimes))
...
status = nf90_get_var(ncid, rhVarId, rhValues)
if(status /= nf90_NoErr) call handle_err(status)

```

Here is an example using `NF90_GET_VAR` to read a section of the variable named `rh` from an existing netCDF dataset named `foo.nc`. For simplicity in this example, we assume that we know that `rh` is dimensioned with `lon`, `lat`, and `time`, that there are ten `lon` values,

five lat values, and three time values, and that we want to replace all the values at the last time.

```

use netcdf
implicit none
integer          :: ncId, rhVarId, status
integer, parameter :: numLons = 10, numLats = 5, numTimes = 3
real, dimension(numLons, numLats, numTimes) &
               :: rhValues
...
status = nf90_open("foo.nc", nf90_NoWrite, ncId)
if(status /= nf90_NoErr) call handle_err(status)
...
status = nf90_inq_varid(ncId, "rh", rhVarId)
if(status /= nf90_NoErr) call handle_err(status)
!Read the values at the last time by passing an array section
status = nf90_get_var(ncId, rhVarId, rhValues(:, :, 3), &
                    start = (/ 1, 1, numTimes /), &
                    count = (/ numLats, numLons, 1 /))
if(status /= nf90_NoErr) call handle_err(status)

```

Here is an example of using `NF_GET_VAR` to read every other point of a netCDF variable named `rh` having dimensions (6, 4).

```

use netcdf
implicit none
integer          :: ncId, rhVarId, status
integer, parameter :: numLons = 6, numLats = 4
real, dimension(numLons, numLats) &
               :: rhValues
...
status = nf90_open("foo.nc", nf90_NoWrite, ncId)
if(status /= nf90_NoErr) call handle_err(status)
...
status = nf90_inq_varid(ncId, "rh", rhVarId)
if(status /= nf90_NoErr) call handle_err(status)
...
! Read every other value into an array section
status = nf90_get_var(ncId, rhVarId, rhValues(::2, ::2) &
                    stride = (/ 2, 2 /))
if(status /= nf90_NoErr) call handle_err(status)

```

The following map vector shows the default mapping between a 2x3x4 netCDF variable and an internal array of the same shape:

```

real,    dimension(2, 3, 4):: a ! same shape as netCDF variable
integer, dimension(3)      :: map = (/ 1, 2, 6 /)
                                ! netCDF dimension inter-element distance
                                ! -----
                                ! most rapidly varying           1
                                ! intermediate                     2 (= map(1)*2)

```

```
! most slowly varying      6 (= map(2)*3)
```

Using the map vector above obtains the same result as simply not passing a map vector at all.

Here is an example of using `nf90_get_var` to read a netCDF variable named `rh` whose dimensions are the transpose of the Fortran 90 array:

```
use netcdf
implicit none
integer          :: ncId, rhVarId, status
integer, parameter :: numLons = 6, numLats = 4
real, dimension(numLons, numLats) :: rhValues
! netCDF variable has dimensions (numLats, numLons)
...
status = nf90_open("foo.nc", nf90_NoWrite, ncId)
if(status /= nf90_NoErr) call handle_err(status)
...
status = nf90_inq_varid(ncId, "rh", rhVarId)
if(status /= nf90_NoErr) call handle_err(status)
...
! Read transposed values: map vector would be (/ 1, numLats /) for
!   no transposition
status = nf90_get_var(ncId, rhVarId, rhValues, map = (/ numLons, 1 /))
if(status /= nf90_NoErr) call handle_err(status)
```

The same effect can be obtained more simply, though using more memory, using Fortran 90 intrinsic functions:

```
use netcdf
implicit none
integer          :: ncId, rhVarId, status
integer, parameter :: numLons = 6, numLats = 4
real, dimension(numLons, numLats) :: rhValues
! netCDF variable has dimensions (numLats, numLons)
real, dimension(numLons, numLats) :: tempValues
...
status = nf90_open("foo.nc", nf90_NoWrite, ncId)
if(status /= nf90_NoErr) call handle_err(status)
...
status = nf90_inq_varid(ncId, "rh", rhVarId)
if(status /= nf90_NoErr) call handle_err(status)
...
status = nf90_get_var(ncId, rhVarId, tempValues)
if(status /= nf90_NoErr) call handle_err(status)
rhValues(:, :) = transpose(tempValues)
```

4.8 Reading and Writing Character String Values

Character strings are not a primitive netCDF external data type, in part because FORTRAN does not support the abstraction of variable-length character strings (the FORTRAN

LEN function returns the static length of a character string, not its dynamic length). As a result, a character string cannot be written or read as a single object in the netCDF interface. Instead, a character string must be treated as an array of characters, and array access must be used to read and write character strings as variable data in netCDF datasets. Furthermore, variable-length strings are not supported by the netCDF interface except by convention; for example, you may treat a zero byte as terminating a character string, but you must explicitly specify the length of strings to be read from and written to netCDF variables.

Character strings as attribute values are easier to use, since the strings are treated as a single unit for access. However, the value of a character-string attribute is still an array of characters with an explicit length that must be specified when the attribute is defined.

When you define a variable that will have character-string values, use a character-position dimension as the most quickly varying dimension for the variable (the first dimension for the variable in Fortran 90). The length of the character-position dimension will be the maximum string length of any value to be stored in the character-string variable. Space for maximum-length strings will be allocated in the disk representation of character-string variables whether you use the space or not. If two or more variables have the same maximum length, the same character-position dimension may be used in defining the variable shapes.

To write a character-string value into a character-string variable, use either entire variable access or array access. The latter requires that you specify both a corner and a vector of edge lengths. The character-position dimension at the corner should be one for Fortran 90. If the length of the string to be written is n , then the vector of edge lengths will specify n in the character-position dimension, and one for all the other dimensions: $(n, 1, 1, \dots, 1)$.

In Fortran 90, fixed-length strings may be written to a netCDF dataset without a terminating character, to save space. Variable-length strings should follow the C convention of writing strings with a terminating zero byte so that the intended length of the string can be determined when it is later read by either C or Fortran 90 programs.

4.9 Fill Values

What happens when you try to read a value that was never written in an open netCDF dataset? You might expect that this should always be an error, and that you should get an error message or an error status returned. You do get an error if you try to read data from a netCDF dataset that is not open for reading, if the variable ID is invalid for the specified netCDF dataset, or if the specified indices are not properly within the range defined by the dimension lengths of the specified variable. Otherwise, reading a value that was not written returns a special fill value used to fill in any undefined values when a netCDF variable is first written.

You may ignore fill values and use the entire range of a netCDF external data type, but in this case you should make sure you write all data values before reading them. If you know you will be writing all the data before reading it, you can specify that no prefilling of variables with fill values will occur by calling `writing`. This may provide a significant performance gain for netCDF writes.

The variable attribute `_FillValue` may be used to specify the fill value for a variable. There are default fill values for each type, defined in module `netcdf`: `NF90_FILL_CHAR`, `NF90_FILL_INT1` (same as `NF90_FILL_BYTE`), `NF90_FILL_INT2`

(same as `NF90_FILL_SHORT`), `NF90_FILL_INT`, `NF90_FILL_REAL` (same as `NF90_FILL_FLOAT`), and `NF90_FILL_DOUBLE`

The netCDF byte and character types have different default fill values. The default fill value for characters is the zero byte, a useful value for detecting the end of variable-length C character strings. If you need a fill value for a byte variable, it is recommended that you explicitly define an appropriate `_FillValue` attribute, as generic utilities such as `ncdump` will not assume a default fill value for byte variables.

Type conversion for fill values is identical to type conversion for other values: attempting to convert a value from one type to another type that can't represent the value results in a range error. Such errors may occur on writing or reading values from a larger type (such as double) to a smaller type (such as float), if the fill value for the larger type cannot be represented in the smaller type.

4.10 NF90_RENAME_VAR

The function `NF90_RENAME_VAR` changes the name of a netCDF variable in an open netCDF dataset. If the new name is longer than the old name, the netCDF dataset must be in define mode. You cannot rename a variable to have the name of any existing variable.

Usage

```
function nf90_rename_var(ncid, varid, newname)
    integer,          intent( in) :: ncid, varid
    character (len = *), intent( in) :: newname
    integer
                                :: nf90_rename_var
```

`ncid` NetCDF ID, from a previous call to `NF90_OPEN` or `NF90_CREATE`.

`varid` Variable ID.

`newname` New name for the specified variable.

Errors

`NF90_RENAME_VAR` returns the value `NF90_NOERR` if no errors occurred. Otherwise, the returned status indicates an error. Possible causes of errors include:

- The new name is in use as the name of another variable.
- The variable ID is invalid for the specified netCDF dataset.
- The specified netCDF ID does not refer to an open netCDF dataset.

Example

Here is an example using `NF90_RENAME_VAR` to rename the variable `rh` to `rel_hum` in an existing netCDF dataset named `foo.nc`:

```
use netcdf
implicit none
integer :: ncId, rhVarId, status
...
status = nf90_open("foo.nc", nf90_Write, ncid)
```

```
if(status /= nf90_NoErr) call handle_err(status)
...
status = nf90_inq_varid(ncid, "rh", rhVarId)
if(status /= nf90_NoErr) call handle_err(status)
status = nf90_redef(ncid) ! Enter define mode to change variable name
if(status /= nf90_NoErr) call handle_err(status)
status = nf90_rename_var(ncid, rhVarId, "rel_hum")
if(status /= nf90_NoErr) call handle_err(status)
status = nf90_enddef(ncid) ! Leave define mode
if(status /= nf90_NoErr) call handle_err(status)
```

5 Attributes

5.1 Attributes Introduction

Attributes may be associated with each netCDF variable to specify such properties as units, special values, maximum and minimum valid values, scaling factors, and offsets. Attributes for a netCDF dataset are defined when the dataset is first created, while the netCDF dataset is in define mode. Additional attributes may be added later by reentering define mode. A netCDF attribute has a netCDF variable to which it is assigned, a name, a type, a length, and a sequence of one or more values. An attribute is designated by its variable ID and name. When an attribute name is not known, it may be designated by its variable ID and number in order to determine its name, using the function `NF90_INQ_ATTNAME`.

The attributes associated with a variable are typically defined immediately after the variable is created, while still in define mode. The data type, length, and value of an attribute may be changed even when in data mode, as long as the changed attribute requires no more space than the attribute as originally defined.

It is also possible to have attributes that are not associated with any variable. These are called global attributes and are identified by using `NF90_GLOBAL` as a variable pseudo-ID. Global attributes are usually related to the netCDF dataset as a whole and may be used for purposes such as providing a title or processing history for a netCDF dataset.

Operations supported on attributes are:

- Create an attribute, given its variable ID, name, data type, length, and value.
- Get attribute's data type and length from its variable ID and name.
- Get attribute's value from its variable ID and name.
- Copy attribute from one netCDF variable to another.
- Get name of attribute from its number.
- Rename an attribute.
- Delete an attribute.

5.2 Attribute Conventions

Names commencing with underscore ('_') are reserved for use by the netCDF library. Most generic applications that process netCDF datasets assume standard attribute conventions and it is strongly recommended that these be followed unless there are good reasons for not doing so. Below we list the names and meanings of recommended standard attributes that have proven useful. Note that some of these (e.g. `units`, `valid_range`, `scale_factor`) assume numeric data and should not be used with character data.

`units` A character string that specifies the units used for the variable's data. Unidata has developed a freely-available library of routines to convert between character string and binary forms of unit specifications and to perform various useful operations on the binary forms. This library is used in some netCDF applications. Using the recommended units syntax permits data represented in conformable units to be automatically converted to common units for arithmetic operations. See [section "Appendix A - Units" in *The NetCDF Users' Guide*](#).

long_name

A long descriptive name. This could be used for labeling plots, for example. If a variable has no long_name attribute assigned, the variable name should be used as a default.

valid_min

A scalar specifying the minimum valid value for this variable.

valid_max

A scalar specifying the maximum valid value for this variable.

valid_range

A vector of two numbers specifying the minimum and maximum valid values for this variable, equivalent to specifying values for both valid_min and valid_max attributes. Any of these attributes define the valid range. The attribute valid_range must not be defined if either valid_min or valid_max is defined.

Generic applications should treat values outside the valid range as missing. The type of each valid_range, valid_min and valid_max attribute should match the type of its variable (except that for byte data, these can be of a signed integral type to specify the intended range).

If neither valid_min, valid_max nor valid_range is defined then generic applications should define a valid range as follows. If the data type is byte and _FillValue is not explicitly defined, then the valid range should include all possible values. Otherwise, the valid range should exclude the _FillValue (whether defined explicitly or by default) as follows. If the _FillValue is positive then it defines a valid maximum, otherwise it defines a valid minimum. For integer types, there should be a difference of 1 between the _FillValue and this valid minimum or maximum. For floating point types, the difference should be twice the minimum possible (1 in the least significant bit) to allow for rounding error.

scale_factor

If present for a variable, the data are to be multiplied by this factor after the data are read by the application that accesses the data.

add_offset

If present for a variable, this number is to be added to the data after it is read by the application that accesses the data. If both scale_factor and add_offset attributes are present, the data are first scaled before the offset is added. The attributes scale_factor and add_offset can be used together to provide simple data compression to store low-resolution floating-point data as small integers in a netCDF dataset. When scaled data are written, the application should first subtract the offset and then divide by the scale factor.

When scale_factor and add_offset are used for packing, the associated variable (containing the packed data) is typically of type byte or short, whereas the unpacked values are intended to be of type float or double. The attributes scale_factor and add_offset should both be of the type intended for the unpacked data, e.g. float or double.

_FillValue

The `_FillValue` attribute specifies the fill value used to pre-fill disk space allocated to the variable. Such pre-fill occurs unless `nofill` mode is set using `NF90_SET_FILL`. See [Section 2.13 \[NF90_SET_FILL\]](#), page 19. The fill value is returned when reading values that were never written. If `_FillValue` is defined then it should be scalar and of the same type as the variable. It is not necessary to define your own `_FillValue` attribute for a variable if the default fill value for the type of the variable is adequate. However, use of the default fill value for data type byte is not recommended. Note that if you change the value of this attribute, the changed value applies only to subsequent writes; previously written data are not changed.

Generic applications often need to write a value to represent undefined or missing values. The fill value provides an appropriate value for this purpose because it is normally outside the valid range and therefore treated as missing when read by generic applications. It is legal (but not recommended) for the fill value to be within the valid range.

See [Section 4.9 \[Fill Values\]](#), page 42.

missing_value

This attribute is not treated in any special way by the library or conforming generic applications, but is often useful documentation and may be used by specific applications. The `missing_value` attribute can be a scalar or vector containing values indicating missing data. These values should all be outside the valid range so that generic applications will treat them as missing.

signedness

Deprecated attribute, originally designed to indicate whether byte values should be treated as signed or unsigned. The attributes `valid_min` and `valid_max` may be used for this purpose. For example, if you intend that a byte variable store only nonnegative values, you can use `valid_min = 0` and `valid_max = 255`. This attribute is ignored by the `netCDF` library.

C_format

A character array providing the format that should be used by C applications to print values for this variable. For example, if you know a variable is only accurate to three significant digits, it would be appropriate to define the `C_format` attribute as `"%.3g"`. The `ncdump` utility program uses this attribute for variables for which it is defined. The format applies to the scaled (internal) type and value, regardless of the presence of the scaling attributes `scale_factor` and `add_offset`.

FORTRAN_format

A character array providing the format that should be used by FORTRAN applications to print values for this variable. For example, if you know a variable is only accurate to three significant digits, it would be appropriate to define the `FORTRAN_format` attribute as `"(G10.3)"`.

title

A global attribute that is a character array providing a succinct description of what is in the dataset.

history A global attribute for an audit trail. This is a character array with a line for each invocation of a program that has modified the dataset. Well-behaved generic netCDF applications should append a line containing: date, time of day, user name, program name and command arguments.

Conventions

If present, 'Conventions' is a global attribute that is a character array for the name of the conventions followed by the dataset, in the form of a string that is interpreted as a directory name relative to a directory that is a repository of documents describing sets of discipline-specific conventions. This permits a hierarchical structure for conventions and provides a place where descriptions and examples of the conventions may be maintained by the defining institutions and groups. The conventions directory name is currently interpreted relative to the directory `pub/netcdf/Conventions/` on the host machine `ftp.unidata.ucar.edu`. Alternatively, a full URL specification may be used to name a WWW site where documents that describe the conventions are maintained.

For example, if a group named NUWG agrees upon a set of conventions for dimension names, variable names, required attributes, and netCDF representations for certain discipline-specific data structures, they may store a document describing the agreed-upon conventions in a dataset in the NUWG/ subdirectory of the Conventions directory. Datasets that followed these conventions would contain a global Conventions attribute with value "NUWG".

Later, if the group agrees upon some additional conventions for a specific subset of NUWG data, for example time series data, the description of the additional conventions might be stored in the NUWG/Time_series/ subdirectory, and datasets that adhered to these additional conventions would use the global Conventions attribute with value "NUWG/Time_series", implying that this dataset adheres to the NUWG conventions and also to the additional NUWG time-series conventions.

5.3 Create an Attribute: NF90_PUT_ATT

The function `NF90_PUT_ATT` adds or changes a variable attribute or global attribute of an open netCDF dataset. If this attribute is new, or if the space required to store the attribute is greater than before, the netCDF dataset must be in define mode.

Usage

Although it's possible to create attributes of all types, text and double attributes are adequate for most purposes.

```
function nf90_put_att(ncid, varid, name, values)
  integer,          intent( in) :: ncid, varid
  character(len = *), intent( in) :: name
  any valid type, scalar or array of rank 1, &
  intent( in) :: values
  integer          :: nf90_put_att
```

ncid NetCDF ID, from a previous call to `NF90_OPEN` or `NF90_CREATE`.

varid	Variable ID of the variable to which the attribute will be assigned or NF90_GLOBAL for a global attribute.
name	Attribute name. Must begin with an alphabetic character, followed by zero or more alphanumeric characters including the underscore ('_'). Case is significant. Attribute name conventions are assumed by some netCDF generic applications, e.g., units as the name for a string attribute that gives the units for a netCDF variable. For examples of attribute conventions Section 5.2 [Attribute Conventions] , page 45.
values	An array of attribute values. Values may be supplied as scalars or as arrays of rank one (one dimensional vectors). The external data type of the attribute is set to match the internal representation of the argument, that is if values is a two byte integer array, the attribute will be of type NF90_INT2. Fortran 90 intrinsic functions can be used to convert attributes to the desired type.

Errors

NF90_PUT_ATT returns the value NF90_NOERR if no errors occurred. Otherwise, the returned status indicates an error. Possible causes of errors include:

- The variable ID is invalid for the specified netCDF dataset.
- The specified netCDF type is invalid.
- The specified length is negative.
- The specified open netCDF dataset is in data mode and the specified attribute would expand.
- The specified open netCDF dataset is in data mode and the specified attribute does not already exist.
- The specified netCDF ID does not refer to an open netCDF dataset.
- The number of attributes for this variable exceeds NF90_MAX_ATTRS.

Example

Here is an example using NF90_PUT_ATT to add a variable attribute named valid_range for a netCDF variable named rh and a global attribute named title to an existing netCDF dataset named foo.nc:

```

use netcdf
implicit none
integer :: ncid, status, RHVarID
...
status = nf90_open("foo.nc", nf90_write, ncid)
if (status /= nf90_noerr) call handle_err(status)
...
! Enter define mode so we can add the attribute
status = nf90_redef(ncid)
if (status /= nf90_noerr) call handle_err(status)
! Get the variable ID for "rh"...
status = nf90_inq_varid(ncid, "rh", RHVarID)

```

```

if (status /= nf90_noerr) call handle_err(status)
! ... put the range attribute, setting it to eight byte reals...
status = nf90_put_att(ncid, RHVarID, "valid_range", real((/ 0, 100 /))
! ... and the title attribute.
if (status /= nf90_noerr) call handle_err(status)
status = nf90_put_att(ncid, RHVarID, "title", "example netCDF dataset" )
if (status /= nf90_noerr) call handle_err(status)
! Leave define mode
status = nf90_enddef(ncid)
if (status /= nf90_noerr) call handle_err(status)

```

5.4 Get Information about an Attribute: NF90_INQUIRE_ATTRIBUTE and NF90_INQ_ATTNAME

The function `NF90_INQUIRE_ATTRIBUTE` returns information about a netCDF attribute given the variable ID and attribute name. Information about an attribute includes its type, length, name, and number. See `NF90_GET_ATT` for getting attribute values.

The function `NF90_INQ_ATTNAME` gets the name of an attribute, given its variable ID and number. This function is useful in generic applications that need to get the names of all the attributes associated with a variable, since attributes are accessed by name rather than number in all other attribute functions. The number of an attribute is more volatile than the name, since it can change when other attributes of the same variable are deleted. This is why an attribute number is not called an attribute ID.

Usage

```

function nf90_inquire_attribute(ncid, varid, name, xtype, len, attnum)
  integer,          intent( in)          :: ncid, varid
  character (len = *), intent( in)      :: name
  integer,          intent(out), optional :: xtype, len, attnum
  integer          :: nf90_inquire_attribute
function nf90_inq_attname(ncid, varid, attnum, name)
  integer,          intent( in) :: ncid, varid, attnum
  character (len = *), intent(out) :: name
  integer          :: nf90_inq_attname

```

<code>ncid</code>	NetCDF ID, from a previous call to <code>NF90_OPEN</code> or <code>NF90_CREATE</code> .
<code>varid</code>	Variable ID of the attribute's variable, or <code>NF90_GLOBAL</code> for a global attribute.
<code>name</code>	Attribute name. For <code>NF90_INQ_ATTNAME</code> , this is a pointer to the location for the returned attribute name.
<code>xtype</code>	Returned attribute type, one of the set of predefined netCDF external data types. The valid netCDF external data types are <code>NF90_BYTE</code> , <code>NF90_CHAR</code> , <code>NF90_SHORT</code> , <code>NF90_INT</code> , <code>NF90_FLOAT</code> , and <code>NF90_DOUBLE</code> .
<code>len</code>	Returned number of values currently stored in the attribute. For a string-valued attribute, this is the number of characters in the string.

attnum For `NF90_INQ_ATTNAME`, the input attribute number; for `NF90_INQ_ATTID`, the returned attribute number. The attributes for each variable are numbered from 1 (the first attribute) to `NATTS`, where `NATTS` is the number of attributes for the variable, as returned from a call to `NF90_INQ_VARNATTS`.

(If you already know an attribute name, knowing its number is not very useful, because accessing information about an attribute requires its name.)

Errors

Each function returns the value `NF90_NOERR` if no errors occurred. Otherwise, the returned status indicates an error. Possible causes of errors include:

- The variable ID is invalid for the specified netCDF dataset.
- The specified attribute does not exist.
- The specified netCDF ID does not refer to an open netCDF dataset.
- For `NF90_INQ_ATTNAME`, the specified attribute number is negative or more than the number of attributes defined for the specified variable.

Example

Here is an example using `NF90_INQUIRE_ATTRIBUTE` to inquire about the lengths of an attribute named `valid_range` for a netCDF variable named `rh` and a global attribute named `title` in an existing netCDF dataset named `foo.nc`:

```

use netcdf
implicit none
integer :: ncid, status
integer :: RHVarID                      ! Variable ID
integer :: validRangeLength, titleLength ! Attribute lengths
...
status = nf90_open("foo.nc", nf90_nowrite, ncid)
if (status /= nf90_noerr) call handle_err(status)
...
! Get the variable ID for "rh"...
status = nf90_inq_varid(ncid, "rh", RHVarID)
if (status /= nf90_noerr) call handle_err(status)
! ... get the length of the "valid_range" attribute...
status = nf90_inquire_attribute(ncid, RHVarID, "valid_range", &
                               len = validRangeLength)
if (status /= nf90_noerr) call handle_err(status)
! ... and the global title attribute.
status = nf90_inquire_attribute(ncid, nf90_global, "title", len = titleLength)
if (status /= nf90_noerr) call handle_err(status)

```

5.5 Get Attribute's Values: `NF90_GET_ATT`

Function `nf90_get_att` gets the value(s) of a netCDF attribute, given its variable ID and name.

Usage

```
function nf90_get_att(ncid, varid, name, values)
  integer,          intent( in) :: ncid, varid
  character(len = *), intent( in) :: name
  any valid type, scalar or array of rank 1, &
                      intent(out) :: values
  integer          :: nf90_get_att
```

ncid NetCDF ID, from a previous call to `NF90_OPEN` or `NF90_CREATE`.

varid Variable ID of the attribute's variable, or `NF90_GLOBAL` for a global attribute.

name Attribute name.

values Returned attribute values. All elements of the vector of attribute values are returned, so you must provide enough space to hold them. If you don't know how much space to reserve, call `NF90_INQUIRE_ATTRIBUTE` first to find out the length of the attribute. If there is only a single attribute values may be a scalar. If the attribute is of type character values should be a variable of type character with the `len` Fortran 90 attribute set to an appropriate value (i.e. `character (len = 80) :: values`). You cannot read character data from a numeric variable or numeric data from a text variable. For numeric data, if the type of data differs from the netCDF variable type, type conversion will occur. See section "Type Conversion" in *NetCDF Users Guide*.

Errors

`NF90_GET_ATT_` type returns the value `NF90_NOERR` if no errors occurred. Otherwise, the returned status indicates an error. Possible causes of errors include:

- The variable ID is invalid for the specified netCDF dataset.
- The specified attribute does not exist.
- The specified netCDF ID does not refer to an open netCDF dataset.
- One or more of the attribute values are out of the range of values representable by the desired type.

Example

Here is an example using `NF90_GET_ATT` to determine the values of an attribute named `valid_range` for a netCDF variable named `rh` and a global attribute named `title` in an existing netCDF dataset named `foo.nc`. In this example, it is assumed that we don't know how many values will be returned, so we first inquire about the length of the attributes to make sure we have enough space to store them:

```
use netcdf
implicit none
integer          :: ncid, status
integer          :: RHVarID                ! Variable ID
integer          :: validRangeLength, titleLength ! Attribute lengths
real, dimension(:), allocatable, &
```

```

                                :: validRange
character (len = 80) :: title
...
status = nf90_open("foo.nc", nf90_nowrite, ncid)
if (status /= nf90_noerr) call handle_err(status)
...
! Find the lengths of the attributes
status = nf90_inq_varid(ncid, "rh", RHVarID)
if (status /= nf90_noerr) call handle_err(status)
status = nf90_inquire_attribute(ncid, RHVarID, "valid_range", &
                                len = validRangeLength)
if (status /= nf90_noerr) call handle_err(status)
status = nf90_inquire_attribute(ncid, nf90_global, "title", len = titleLength)
if (status /= nf90_noerr) call handle_err(status)
...
!Allocate space to hold attribute values, check string lengths
allocate(validRange(validRangeLength), stat = status)
if(status /= 0 .or. len(title) < titleLength)
  print *, "Not enough space to put attribute values."
  exit
end if
! Read the attributes.
status = nf90_get_att(ncid, RHVarID, "valid_range", validRange)
if (status /= nf90_noerr) call handle_err(status)
status = nf90_get_att(ncid, nf90_global, "title", title)
if (status /= nf90_noerr) call handle_err(status)

```

5.6 Copy Attribute from One NetCDF to Another: NF90_COPY_ATT

The function `NF90_COPY_ATT` copies an attribute from one open netCDF dataset to another. It can also be used to copy an attribute from one variable to another within the same netCDF dataset.

Usage

```

function nf90_copy_att(ncid_in, varid_in, name, ncid_out, varid_out)
  integer,          intent( in) :: ncid_in,  varid_in
  character (len = *), intent( in) :: name
  integer,          intent( in) :: ncid_out,  varid_out
  integer           :: nf90_copy_att

```

- ncid_in** The netCDF ID of an input netCDF dataset from which the attribute will be copied, from a previous call to `NF90_OPEN` or `NF90_CREATE`.
- varid_in** ID of the variable in the input netCDF dataset from which the attribute will be copied, or `NF90_GLOBAL` for a global attribute.
- name** Name of the attribute in the input netCDF dataset to be copied.

ncid_out The netCDF ID of the output netCDF dataset to which the attribute will be copied, from a previous call to NF90_OPEN or NF90_CREATE. It is permissible for the input and output netCDF IDs to be the same. The output netCDF dataset should be in define mode if the attribute to be copied does not already exist for the target variable, or if it would cause an existing target attribute to grow.

varid_out ID of the variable in the output netCDF dataset to which the attribute will be copied, or NF90_GLOBAL to copy to a global attribute.

Errors

NF90_COPY_ATT returns the value NF90_NOERR if no errors occurred. Otherwise, the returned status indicates an error. Possible causes of errors include:

- The input or output variable ID is invalid for the specified netCDF dataset.
- The specified attribute does not exist.
- The output netCDF is not in define mode and the attribute is new for the output dataset is larger than the existing attribute.
- The input or output netCDF ID does not refer to an open netCDF dataset.

Example

Here is an example using NF90_COPY_ATT to copy the variable attribute units from the variable rh in an existing netCDF dataset named foo.nc to the variable avgrh in another existing netCDF dataset named bar.nc, assuming that the variable avgrh already exists, but does not yet have a units attribute:

```

use netcdf
implicit none
integer :: ncid1, ncid2, status
integer :: RHVarID, avgRHVarID    ! Variable ID
...
status = nf90_open("foo.nc", nf90_nowrite, ncid1)
if (status /= nf90_noerr) call handle_err(status)
status = nf90_open("bar.nc", nf90_write, ncid2)
if (status /= nf90_noerr) call handle_err(status)
...
! Find the IDs of the variables
status = nf90_inq_varid(ncid1, "rh", RHVarID)
if (status /= nf90_noerr) call handle_err(status)
status = nf90_inq_varid(ncid1, "avgrh", avgRHVarID)
if (status /= nf90_noerr) call handle_err(status)
...
status = nf90_redef(ncid2)    ! Enter define mode
if (status /= nf90_noerr) call handle_err(status)
! Copy variable attribute from "rh" in file 1 to "avgrh" in file 1
status = nf90_copy_att(ncid1, RHVarID, "units", ncid2, avgRHVarID)

```

```

if (status /= nf90_noerr) call handle_err(status)
status = nf90_enddef(ncid2)
if (status /= nf90_noerr) call handle_err(status)

```

5.7 Rename an Attribute: NF90_RENAME_ATT

The function `NF90_RENAME_ATT` changes the name of an attribute. If the new name is longer than the original name, the netCDF dataset must be in define mode. You cannot rename an attribute to have the same name as another attribute of the same variable.

Usage

```

function nf90_rename_att(ncid, varid, curname, newname)
  integer,          intent( in) :: ncid,  varid
  character (len = *), intent( in) :: curname, newname
  integer          :: nf90_rename_att

```

ncid NetCDF ID, from a previous call to `NF90_OPEN` or `NF90_CREATE`

varid ID of the attribute's variable, or `NF90_GLOBAL` for a global attribute

curname The current attribute name.

newname The new name to be assigned to the specified attribute. If the new name is longer than the current name, the netCDF dataset must be in define mode.

Errors

`NF90_RENAME_ATT` returns the value `NF90_NOERR` if no errors occurred. Otherwise, the returned status indicates an error. Possible causes of errors include:

- The specified variable ID is not valid.
- The new attribute name is already in use for another attribute of the specified variable.
- The specified netCDF dataset is in data mode and the new name is longer than the old name.
- The specified attribute does not exist.
- The specified netCDF ID does not refer to an open netCDF dataset.

Example

Here is an example using `NF90_RENAME_ATT` to rename the variable attribute units to Units for a variable `rh` in an existing netCDF dataset named `foo.nc`:

```

use netcdf
implicit none
integer :: ncid1, status
integer :: RHVarID      ! Variable ID
...
status = nf90_open("foo.nc", nf90_nowrite, ncid)
if (status /= nf90_noerr) call handle_err(status)
...

```

```

! Find the IDs of the variables
status = nf90_inq_varid(ncid, "rh", RHVarID)
if (status /= nf90_noerr) call handle_err(status)
...
status = nf90_rename_att(ncid, RHVarID, "units", "Units")
if (status /= nf90_noerr) call handle_err(status)

```

5.8 NF90_DEL_ATT

The function NF90_DEL_ATT deletes a netCDF attribute from an open netCDF dataset. The netCDF dataset must be in define mode.

Usage

```

function nf90_del_att(ncid, varid, name)
  integer,          intent( in) :: ncid, varid
  character (len = *), intent( in) :: name
  integer          :: nf90_del_att

```

ncid NetCDF ID, from a previous call to NF90_OPEN or NF90_CREATE.

varid ID of the attribute's variable, or NF90_GLOBAL for a global attribute.

name The name of the attribute to be deleted.

Errors

NF90_DEL_ATT returns the value NF90_NOERR if no errors occurred. Otherwise, the returned status indicates an error. Possible causes of errors include:

- The specified variable ID is not valid.
- The specified netCDF dataset is in data mode.
- The specified attribute does not exist.
- The specified netCDF ID does not refer to an open netCDF dataset.

Example

Here is an example using NF90_DEL_ATT to delete the variable attribute Units for a variable rh in an existing netCDF dataset named foo.nc:

```

use netcdf
implicit none
integer :: ncid1, status
integer :: RHVarID      ! Variable ID
...
status = nf90_open("foo.nc", nf90_nowrite, ncid)
if (status /= nf90_noerr) call handle_err(status)
...
! Find the IDs of the variables
status = nf90_inq_varid(ncid, "rh", RHVarID)
if (status /= nf90_noerr) call handle_err(status)

```

```
...
status = nf90_redef(ncid)    ! Enter define mode
if (status /= nf90_noerr) call handle_err(status)
status = nf90_del_att(ncid, RHVarID, "Units")
if (status /= nf90_noerr) call handle_err(status)
status = nf90_enddef(ncid)
if (status /= nf90_noerr) call handle_err(status)
```


Appendix A Appendix A - Summary of Fortran 90 Interface

Dataset Functions

```

function nf90_inq_libvers()
  character(len = 80) :: nf90_inq_libvers
function nf90_strerror(ncerr)
  integer, intent( in) :: ncerr
  character(len = 80)  :: nf90_strerror
function nf90_create(path, cmode, ncid)
  character (len = *), intent(in  ) :: path
  integer,          intent(in  ) :: cmode
  integer, optional, intent(in  ) :: initialsize
  integer, optional, intent(inout) :: chunksize
  integer,          intent( out) :: ncid
  integer          :: nf90_create
function nf90_open(path, mode, ncid, chunksize)
  character (len = *), intent(in  ) :: path
  integer,          intent(in  ) :: mode
  integer,          intent( out) :: ncid
  integer, optional, intent(inout) :: chunksize
  integer          :: nf90_open
function nf90_set_fill(ncid, fillmode, old_mode)
  integer, intent( in) :: ncid, fillmode
  integer, intent(out) :: old_mode
  integer          :: nf90_set_fill
function nf90_redef(ncid)
  integer, intent( in) :: ncid
  integer          :: nf90_redef
function nf90_enddef(ncid, h_minfree, v_align, v_minfree, r_align)
  integer,          intent( in) :: ncid
  integer, optional, intent( in) :: h_minfree, v_align, v_minfree, r_align
  integer          :: nf90_enddef
function nf90_sync(ncid)
  integer, intent( in) :: ncid
  integer          :: nf90_sync
function nf90_abort(ncid)
  integer, intent( in) :: ncid
  integer          :: nf90_abort
function nf90_close(ncid)
  integer, intent( in) :: ncid
  integer          :: nf90_close
function nf90_Inquire(ncid, nDimensions, nVariables, nAttributes, &
                     unlimitedDimId)
  integer,          intent( in) :: ncid
  integer, optional, intent(out) :: nDimensions, nVariables, nAttributes, &
                                   unlimitedDimId
  integer          :: nf90_Inquire

```

Dimension functions

```

function nf90_def_dim(ncid, name, len, dimid)
  integer,          intent( in) :: ncid
  character (len = *), intent( in) :: name
  integer,          intent( in) :: len
  integer,          intent(out) :: dimid
  integer           :: nf90_def_dim
function nf90_inq_dimid(ncid, name, dimid)
  integer,          intent( in) :: ncid
  character (len = *), intent( in) :: name
  integer,          intent(out) :: dimid
  integer           :: nf90_inq_dimid
function nf90_inquire_dimension(ncid, dimid, name, len)
  integer,          intent( in) :: ncid, dimid
  character (len = *), optional, intent(out) :: name
  integer,          optional, intent(out) :: len
  integer           :: nf90_inquire_dimension
function nf90_rename_dim(ncid, dimid, name)
  integer,          intent( in) :: ncid
  character (len = *), intent( in) :: name
  integer,          intent( in) :: dimid
  integer           :: nf90_rename_dim

```

Variable functions

```

function nf90_def_var(ncid, name, xtype, dimids, varid)
  integer,          intent( in) :: ncid
  character (len = *), intent( in) :: name
  integer,          intent( in) :: xtype
  integer, dimension(:), intent( in) :: dimids ! May be omitted, scalar,
                                              ! vector
  integer           :: nf90_def_var
function nf90_inq_varid(ncid, name, varid)
  integer,          intent( in) :: ncid
  character (len = *), intent( in) :: name
  integer,          intent(out) :: varid
  integer           :: nf90_inq_varid
function nf90_inquire_variable(ncid, varid, name, xtype, ndims, &
                               dimids, nAtts)
  integer,          intent( in) :: ncid, varid
  character (len = *), optional, intent(out) :: name
  integer,          optional, intent(out) :: xtype, ndims
  integer, dimension(*), optional, intent(out) :: dimids
  integer,          optional, intent(out) :: nAtts
  integer           :: nf90_inquire_variable
function nf90_put_var(ncid, varid, values, start, stride, map)
  integer,          intent( in) :: ncid, varid
  any valid type, scalar or array of any rank, &

```



```
character (len = *), intent( in) :: name  
integer                :: nf90_del_att
```

Appendix B Appendix B - FORTRAN 77 to Fortran 90 Transition Guide

The new Fortran 90 interface

The Fortran 90 interface to the netCDF library closely follows the FORTRAN 77 interface. In most cases, function and constant names and argument lists are the same, except that `nf90_` replaces `nf_` in names. The Fortran 90 interface is much smaller than the FORTRAN 77 interface, however. This has been accomplished by using optional arguments and overloaded functions wherever possible.

Because FORTRAN 77 is a subset of Fortran 90, there is no reason to modify working FORTRAN code to use the Fortran 90 interface. New code, however, can easily be patterned after existing FORTRAN while taking advantage of the simpler interface. Some compilers may provide additional support when using Fortran 90. For example, compilers may issue warnings if arguments with `intent(in)` are not set before they are passed to a procedure.

The Fortran 90 interface is currently implemented as a set of wrappers around the base FORTRAN subroutines in the netCDF distribution. Future versions may be implemented entirely in Fortran 90, adding additional error checking possibilities.

Changes to Inquiry functions

In the Fortran 90 interface there are two inquiry functions each for dimensions, variables, and attributes, and a single inquiry function for datasets. These functions take optional arguments, allowing users to request only the information they need. These functions replace the many-argument and single-argument inquiry functions in the FORTRAN interface.

As an example, compare the attribute inquiry functions in the Fortran 90 interface

```
function nf90_inquire_attribute(ncid, varid, name, xtype, len, attnum)
  integer,          intent( in)          :: ncid, varid
  character (len = *), intent( in)       :: name
  integer,          intent(out), optional :: xtype, len, attnum
  integer          :: nf90_inquire_attribute
function nf90_inq_attname(ncid, varid, attnum, name)
  integer,          intent( in) :: ncid, varid, attnum
  character (len = *), intent(out) :: name
  integer          :: nf90_inq_attname
```

with those in the FORTRAN interface

```
INTEGER FUNCTION  NF_INQ_ATT      (NCID, VARID, NAME, xtype, len)
INTEGER FUNCTION  NF_INQ_ATTID   (NCID, VARID, NAME, attnum)
INTEGER FUNCTION  NF_INQ_ATTTYPE (NCID, VARID, NAME, xtype)
INTEGER FUNCTION  NF_INQ_ATTLEN  (NCID, VARID, NAME, len)
INTEGER FUNCTION  NF_INQ_ATTNAME (NCID, VARID, ATTNUM, name)
```

Changes to put and get function

The biggest simplification in the Fortran 90 is in the `nf90_put_var` and `nf90_get_var` functions. Both functions are overloaded: the values argument can be a scalar or an array any rank (7 is the maximum rank allowed by Fortran 90), and may be of any numeric type or the default character type. The netCDF library provides transparent conversion between the external representation of the data and the desired internal representation.

The start, count, stride, and map arguments to `nf90_put_var` and `nf90_get_var` are optional. By default, data is read from or written to consecutive values of starting at the origin of the netCDF variable; the shape of the argument determines how many values are read from or written to each dimension. Any or all of these arguments may be supplied to override the default behavior.

Note also that Fortran 90 allows arbitrary array sections to be passed to any procedure, which may greatly simplify programming. For examples [Section 4.6 \[NF90_PUT_VAR\]](#), page 32 and [Section 4.7 \[NF90_GET_VAR\]](#), page 36.

Index

A

attributes, adding 4

C

common netcdf commands 1
 compiling with netCDF library 6

D

dataset, creating 1
 datasets, overview 7
 dimensions, adding 4

E

error handling 5

I

interface descriptions 7

L

linking to netCDF library 6

N

NF90_ABORT 18
 NF90_ABORT , example 18
 NF90_CLOSE 14
 NF90_CLOSE , example 14
 NF90_CLOSE, typical use 1
 NF90_COPY_ATT 53
 NF90_COPY_ATT, example 53
 NF90_CREATE 9
 NF90_CREATE , example 9
 NF90_CREATE, typical use 1
 NF90_DEF_DIM 21
 NF90_DEF_DIM, example 21
 NF90_DEF_DIM, typical use 1
 NF90_DEF_VAR 28
 NF90_DEF_VAR, example 28
 NF90_DEF_VAR, typical use 1
 NF90_DEL_ATT 56
 NF90_DEL_ATT , example 56
 NF90_ENDDEF 13
 NF90_ENDDEF , example 13
 NF90_ENDDEF, typical use 1
 NF90_GET_ATT 51
 NF90_GET_ATT, example 51
 NF90_GET_ATT, typical use 2
 NF90_GET_VAR 37
 NF90_GET_VAR, example 37
 NF90_GET_VAR, typical use 2
 NF90_INQ_ATTNAME, typical use 2
 NF90_INQ_DIMID 22
 NF90_INQ_DIMID , example 22

NF90_INQ_DIMID, typical use 2
 NF90_INQ_LIBVERS 8
 NF90_INQ_LIBVERS, example 8
 NF90_INQ_VARID 29
 NF90_INQ_VARID , example 29
 NF90_INQ_VARID, typical use 2, 3
 NF90_INQUIRE, typical use 2
 NF90_INQUIRE_ATTRIBUTE and
 NF90_INQ_ATTNAME 50
 NF90_INQUIRE_ATTRIBUTE and
 NF90_INQ_ATTNAME, example 50
 NF90_INQUIRE_ATTRIBUTE, typical use 2
 NF90_INQUIRE_DIMENSION 23
 NF90_INQUIRE_DIMENSION , example 23
 NF90_INQUIRE_DIMENSION, typical use 2
 NF90_INQUIRE_VARIABLE 30
 NF90_INQUIRE_VARIABLE , example 30
 NF90_INQUIRE_VARIABLE, typical use 2
 NF90_OPEN 10
 NF90_OPEN , example 10
 NF90_OPEN, typical use 2
 NF90_PUT_ATT 48
 NF90_PUT_ATT, example 48
 NF90_PUT_ATT, typical use 1, 3
 NF90_PUT_VAR 32
 NF90_PUT_VAR, example 32
 NF90_PUT_VAR, typical use 1, 3
 NF90_REDEF 12
 NF90_REDEF , example 12
 NF90_REDEF, typical use 4
 NF90_RENAME_ATT 55
 NF90_RENAME_ATT, example 55
 NF90_RENAME_DIM 24
 NF90_RENAME_DIM , example 24
 NF90_RENAME_VAR 43
 NF90_RENAME_VAR , example 43
 NF90_SET_FILL 19
 NF90_SET_FILL , example 19
 NF90_STRERROR 8
 NF90_STRERROR, example 8
 NF90_STRERROR, introduction 5
 NF90_SYNC 16
 NF90_SYNC , example 16

R

reading dataset with unknown names 2

U

users' guide, netcdf 1

V

variables, adding 4

W

writing to existing dataset 3