

The NetCDF Fortran 77 Interface Guide

NetCDF Version 3.6.2
Last Updated 29 October 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	3
1.4	Adding New Dimensions, Variables, Attributes	4
1.5	Error Handling	5
1.6	Compiling and Linking with the NetCDF Library	5
2	Datasets	7
2.1	Datasets Introduction	7
2.2	NetCDF Library Interface Descriptions	7
2.3	NF_STRERROR	8
2.4	Get netCDF library version: NF_INQ_LIBVERS	8
2.5	NF_CREATE	9
2.6	NF__CREATE	10
2.7	NF_OPEN	12
2.8	NF__OPEN	13
2.9	NF_REDEF	14
2.10	NF_ENDDEF	15
2.11	NF__ENDDEF	15
2.12	NF_CLOSE	17
2.13	NF_INQ Family	18
2.14	NF_SYNC	19
2.15	NF_ABORT	21
2.16	NF_SET_FILL	22
2.17	NF_SET_DEFAULT_FORMAT	23
3	Dimensions	25
3.1	Dimensions Introduction	25
3.2	NF_DEF_DIM	25
3.3	NF_INQ_DIMID	26
3.4	NF_INQ_DIM Family	27
3.5	NF_RENAME_DIM	28
4	Variables	31
4.1	Variables Introduction	31
4.2	Language Types Corresponding to netCDF external data types	31
4.3	Create a Variable: NF_DEF_VAR	32
4.4	Get a Variable ID from Its Name: NF_INQ_VARID	34
4.5	Get Information about a Variable from Its ID: NF_INQ_VAR family	34

4.6	Write a Single Data Value: <code>NF_PUT_VAR1_</code> type	36
4.7	Write an Entire Variable: <code>NF_PUT_VAR_</code> type	38
4.8	Write an Array of Values: <code>NF_PUT_VARA_</code> type	39
4.9	<code>NF_PUT_VARS_</code> type	42
4.10	<code>NF_PUT_VARM_</code> type	44
4.11	<code>NF_GET_VAR1_</code> type	47
4.12	<code>NF_GET_VAR_</code> type	49
4.13	<code>NF_GET_VARA_</code> type	50
4.14	<code>NF_GET_VARS_</code> type	52
4.15	<code>NF_GET_VARM_</code> type	54
4.16	Reading and Writing Character String Values	58
4.17	Fill Values	60
4.18	<code>NF_RENAME_VAR_</code>	60
5	Attributes	63
5.1	Attributes Introduction	63
5.2	Attribute Conventions	63
5.3	<code>NF_PUT_ATT_</code> type	66
5.4	<code>NF_INQ_ATT</code> Family	68
5.5	<code>NF_GET_ATT_</code> type	70
5.6	<code>NF_COPY_ATT</code>	72
5.7	<code>NF_RENAME_ATT</code>	73
5.8	<code>NF_DEL_ATT</code>	75
Appendix A Summary of FORTRAN 77		
	Interface	77
Appendix B Overview of C interface changes		
	from NetCDF 2 to NetCDF 3	83
B.1	The NetCDF-3 C Interface	83
B.2	Function Naming Conventions	84
B.3	Type Conversion	85
B.4	Error handling	86
B.5	<code>NF_LONG</code> and <code>NF_INT</code>	86
B.6	What's Missing?	86
B.7	Other Changes	86
	Index	91

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 *The 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:

```

NF_CREATE          ! create netCDF dataset: enter define mode
...
NF_DEF_DIM         ! define dimensions: from name and length
...
NF_DEF_VAR        ! define variables: from name, type, dims
...
NF_PUT_ATT        ! assign attribute values
...
NF_ENDDEF         ! end definitions: leave define mode
...
NF_PUT_VAR        ! provide values for variable
...
NF_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 `NF_DEF_DIM` is needed for each dimension created. Similarly, one call to `NF_DEF_VAR` is needed for each variable creation, and one call to a member of the `NF_PUT_ATT` family is needed for each attribute defined and assigned a value. To leave define mode and enter data mode, call `NF_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). Single values may be written to a netCDF variable with one of the members of

the `NF_PUT_VAR1` family, depending on what type of data you have to write. All the values of a variable may be written at once with one of the members of the `NF_PUT_VAR` family. Arrays or array cross-sections of a variable may be written using members of the `NF_PUT_VARA` family. Subsampled array sections may be written using members of the `NF_PUT_VARS` family. Mapped array sections may be written using members of the `NF_PUT_VARM` family. (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 `NF_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 `NF_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 `NF_SYNC` or `NF_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:

```

NF_OPEN           ! open existing netCDF dataset
...
NF_INQ_DIMID     ! get dimension IDs
...
NF_INQ_VARID     ! get variable IDs
...
NF_GET_ATT       ! get attribute values
...
NF_GET_VAR       ! get values of variables
...
NF_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 `NF_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 `NF_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 a member of the `NF_GET_ATT` family (typically `NF_GET_ATT_TEXT` or `NF_GET_ATT_DOUBLE`) for each desired attribute. Variable data values can be directly accessed from the netCDF dataset with calls to members of the `NF_GET_VAR1` family for single values, the `NF_GET_VAR` family for entire variables, or various other members of the `NF_GET_VARA`, `NF_GET_VARS`, or `NF_GET_VARM` families for array, subsampled or mapped access.

Finally, the netCDF dataset is closed with `NF_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:

```

NF_OPEN           ! open existing netCDF dataset
...
NF_INQ           ! find out what is in it
...
NF_INQ_DIM       ! get dimension names, lengths
...
NF_INQ_VAR       ! get variable names, types, shapes
...
NF_INQ_ATTNAME   ! get attribute names
...
NF_INQ_ATT       ! get attribute values
...
NF_GET_ATT       ! get attribute values
...
NF_GET_VAR       ! get values of variables
...
NF_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 `NF_INQ` 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 `NF_INQ_DIM` 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 `NF_INQ_VAR` 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 `NF_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 `NF_INQ_ATT` returns its type and length. Given the type and length, you can allocate enough space to hold the attribute values. Then a call to a member of the `NF_GET_ATT` family returns the attribute values.

Once the IDs and shapes of netCDF variables are known, data values can be accessed by calling a member of the `NF_GET_VAR1` family for single values, or members of the `NF_GET_VAR`, `NF_GET_VARA`, `NF_GET_VARS`, or `NF_GET_VARM` for various kinds of array access.

1.4 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:

```

NF_OPEN           ! open existing netCDF dataset
...
NF_REDEF         ! put it into define mode
...
NF_DEF_DIM       ! define additional dimensions (if any)
...
NF_DEF_VAR       ! define additional variables (if any)
...
NF_PUT_ATT       ! define other attributes (if any)
...
NF_ENDDEF        ! check definitions, leave define mode
...
NF_PUT_VAR       ! provide new variable values
...
NF_CLOSE         ! close netCDF dataset

```

A netCDF dataset is first opened by the `NF_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 `NF_REDEF`. In define mode, call `NF_DEF_DIM` to define new dimensions, `NF_DEF_VAR` to define new variables, and a member of the `NF_PUT_ATT` family 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 `NF_ENDDEF`. If you do not wish to reenter data mode, just call `NF_CLOSE`, which will have the effect of first calling `NF_ENDDEF`.

Until the `NF_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 `NF_ABORT`. You may also use the `NF_ABORT` call to restore the netCDF dataset to a consistent state if the call to `NF_ENDDEF` fails. If you have called `NF_CLOSE` from definition mode and the implied call to `NF_ENDDEF` fails, `NF_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 `NF_SYNC` function and the `NF_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 `NF_SYNC` before the next access.

1.5 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 `NF_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.

1.6 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 file that references netCDF functions or constants must contain an appropriate `INCLUDE` statement before the first such reference:

```
INCLUDE 'netcdf.inc'
```

Unless the `netcdf.inc` file is installed in a standard directory where the FORTRAN compiler always looks, you must use the `-I` option when invoking the compiler, to specify a directory where `netcdf.inc` is installed, for example:

```
f77 -c -I/usr/local/netcdf/include myprogram.f
```

Alternatively, you could specify an absolute path name in the `INCLUDE` statement, but then your program would not compile on another platform where netCDF is installed in a different location.

In all netCDF versions before 3.6.2, the Fortran 77, Fortran 90, and the C libraries were all built into the same library file. That is, the `libnetcdf.a` file contains the C functions, the F77 functions, and the F90 functions. (The C++ library is separate.)

Starting with version 3.6.2, another method of building the netCDF fortran libraries becomes available. With the `--enable-separate-fortran` option to configure, the user can specify that the C library should not contain the fortran functions. In these cases an additional library, `libnetcdf.f.a` (not the extra “f”) will be built. This library contains the fortran functions.

For more information about configure options, See [section “Specifying the Environment for Building”](#) in *The NetCDF Installation and Porting Guide*.

Building separate fortran libraries is required for shared library builds, but is not done, by default, for static library builds.

When linking fortran programs without a separate fortran library, programs must link to the netCDF library like this:

```
f77 -o myprogram myprogram.o -L/usr/local/netcdf/lib -lnetcdf
```

Alternatively, you could specify an absolute path name for the library:

```
f77 -o myprogram myprogram.o -l/usr/local/netcdf/lib/libnetcdf.
```

When linking fortran programs with separate fortran, the user must link to both the fortran and the C libraries.

```
f77 -o myprogram myprogram.o -L/usr/local/netcdf/lib -lnetcdff -lnetcdf
```

Unless the netCDF library (or libraries) 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.

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 function prototype 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 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 NF_STRERROR

The function `NF_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

```
CHARACTER*80 FUNCTION NF_STRERROR(INTEGER NCERR)
```

`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), `NF_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 `NF_STRERROR` to print the error message corresponding to the netCDF error status returned from any netCDF function call and then exit:

```
INCLUDE 'netcdf.inc'
...
SUBROUTINE HANDLE_ERR(STATUS)
INTEGER STATUS
IF (STATUS .NE. NF_NOERR) THEN
  PRINT *, NF_STRERROR(STATUS)
  STOP 'Stopped'
ENDIF
END
```

2.4 Get netCDF library version: NF_INQ_LIBVERS

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

Usage

```
CHARACTER*80 FUNCTION NF_INQ_LIBVERS()
```

Errors

This function takes no arguments, and thus no errors are possible in its invocation.

Example

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

```
INCLUDE 'netcdf.inc'
...
PRINT *, NF_INQ_LIBVERS()
```

2.5 NF_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

```
INTEGER FUNCTION NF_CREATE (CHARACTER*(*) PATH, INTEGER CMODE,
                           INTEGER ncid)
```

PATH The file name of the new netCDF dataset.

CMODE The creation mode flag. The following flags are available: `NF_NOCLOBBER`, `NF_SHARE`, and `NF_64BIT_OFFSET`. You can combine the affect of multiple flags in a single argument by using the bitwise OR operator. For example, to specify both `NF_NOCLOBBER` and `NF_SHARE`, you could provide the argument `OR(NF_NOCLOBBER, NF_SHARE)`.

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

The `NF_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 `NF_SHARE` flag.

Setting `NF_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 *The NetCDF Users Guide*.

A zero value (defined for convenience as `NF_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 *The NetCDF Users Guide*.

`ncid` Returned netCDF ID.

Errors

`NF_CREATE` returns the value `NF_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 `NF_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:

```
INCLUDE 'netcdf.inc'
...
INTEGER NCID, STATUS
...
STATUS = NF_CREATE('foo.nc', NF_NOCLOBBER, NCID)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
```

2.6 NF__CREATE

This function is a variant of `NF_CREATE`, `NF__CREATE` (note the double underscore) allows users to specify two tuning parameters for the file that it is creating. These tuning parameters are not written to the data file, they are only used for so long as the file remains open after an `NF_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

```
INTEGER FUNCTION NF_CREATE (CHARACTER*(*) PATH, INTEGER CMODE, INTEGER INITIALSZ,
                           INTEGER CHUNKSIZEHINT, INTEGER ncid)
```

`PATH` The file name of the new netCDF dataset.

`CMODE` The creation mode flag. The following flags are available: `NF_NOCLOBBER`, `NF_SHARE`, and `NF_64BIT_OFFSET`.

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

The `NF_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 `NF_SHARE` flag.

Setting `NF_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 *The NetCDF Users Guide*](#).

A zero value (defined for convenience as `NF_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 *The NetCDF Users Guide*](#).

`initialsz`

This parameter sets the initial size of the file at creation time.

`chunksizehintp`

The argument referenced by `chunksizehintp` 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 by reference.

Using the value `NF_SIZEHINT_DEFAULT` causes the library to choose a default. How the system chooses the default depends on the system. On many systems, the "preferred I/O block size" is available from the `stat()` system call, struct stat member `st_blksize`. If this is available it is used. Lacking that, twice the system `pagesize` is used.

Lacking a call to discover the system `pagesize`, we just set default `chunksize` to 8192.

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

`ncid` Returned netCDF ID.

Errors

`NF__CREATE` returns the value `NF_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 `NF_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:

```
INCLUDE 'netcdf.inc'
...
INTEGER NCID, STATUS, INITIALSZ, CHUNKSIZEHINT
...
INITIALSZ = 2048
CHUNKSIZEHINT = 1024
STATUS = NF_CREATE('foo.nc', NF_NOCLOBBER, INITIALSZ, CHUNKSIZEHINT, NCID)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
```

2.7 NF_OPEN

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

Usage

```
INTEGER FUNCTION NF_OPEN(CHARACTER*(*) PATH, INTEGER OMODE, INTEGER ncid)
```

<code>PATH</code>	File name for netCDF dataset to be opened.
<code>OMODE</code>	A zero value (or <code>NF_NOWRITE</code>) specifies the default behavior: open the dataset with read-only access, buffering and caching accesses for efficiency. Otherwise, the creation mode is <code>NF_WRITE</code> , <code>NF_SHARE</code> , or <code>OR(NF_WRITE, NF_SHARE)</code> . Setting the <code>NF_WRITE</code> 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 <code>NF_SHARE</code> 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 <code>NF_SHARE</code> flag.
<code>ncid</code>	Returned netCDF ID.

Errors

`NF_OPEN` returns the value `NF_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 `NF_OPEN` to open an existing netCDF dataset named `foo.nc` for read-only, non-shared access:

```

INCLUDE 'netcdf.inc'
...
INTEGER NCID, STATUS
...
STATUS = NF_OPEN('foo.nc', 0, NCID)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)

```

2.8 NF_OPEN

The function `NF_OPEN` opens an existing netCDF dataset for access, with a performance tuning parameter.

Usage

```

INTEGER FUNCTION NF_OPEN(CHARACTER*(*) PATH, INTEGER OMODE, INTEGER
CHUNKSIZEHINT, INTEGER ncid)

```

PATH File name for netCDF dataset to be opened.

OMODE A zero value (or `NF_NOWRITE`) specifies the default behavior: open the dataset with read-only access, buffering and caching accesses for efficiency. Otherwise, the creation mode is `NF_WRITE`, `NF_SHARE`, or `OR(NF_WRITE,NF_SHARE)`. Setting the `NF_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 `NF_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 `NF_SHARE` flag.

CHUNKSIZEHINT

This argument controls a space versus time tradeoff, 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 by reference.

Using the value `NF_SIZEHINT_DEFAULT` causes the library to choose a default. How the system chooses the default depends on the system. On many systems, the "preferred I/O block size" is available from the `stat()` system call, struct `stat` member `st_blksize`. If this is available it is used. Lacking that, twice the system `pagesize` is used.

Lacking a call to discover the system `pagesize`, we just set default chunksize to 8192.

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

ncid Returned netCDF ID.

Errors

NF_OPEN returns the value NF_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 NF_OPEN to open an existing netCDF dataset named foo.nc for read-only, non-shared access:

```
INCLUDE 'netcdf.inc'
...
INTEGER NCID, STATUS, CHUNKSIZEHINT
...
CHUNKSIZEHINT = 1024
STATUS = NF_OPEN('foo.nc', 0, CHUNKSIZEHINT, NCID)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
```

2.9 NF_REDEF

The function NF_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

```
INTEGER FUNCTION NF_REDEF(INTEGER NCID)
NCID      netCDF ID, from a previous call to NF_OPEN or NF_CREATE.
```

Errors

NF_REDEF returns the value NF_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 NF_REDEF to open an existing netCDF dataset named foo.nc and put it into define mode:

```
INCLUDE 'netcdf.inc'
...
INTEGER NCID, STATUS
...
STATUS = NF_OPEN('foo.nc', NF_WRITE, NCID) ! open dataset
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
...
```

```

STATUS = NF_REDEF(NCID)                                ! put in define mode
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)

```

2.10 NF_ENDDEF

The function `NF_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 2.16 \[NF_SET_FILL\]](#), page 22). 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. See [section “File Structure and Performance” in *NetCDF Users’ Guide*](#).

Usage

```

INTEGER FUNCTION NF_ENDDEF(INTEGER NCID)
NCID      NetCDF ID, from a previous call to NF_OPEN or NF_CREATE.

```

Errors

`NF_ENDDEF` returns the value `NF_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 `NF_ENDDEF` to finish the definitions of a new netCDF dataset named `foo.nc` and put it into data mode:

```

INCLUDE 'netcdf.inc'
...
INTEGER NCID, STATUS
...
STATUS = NF_CREATE('foo.nc', NF_NO_CLOBBER, NCID)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)

... ! create dimensions, variables, attributes

STATUS = NF_ENDDEF(NCID)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)

```

2.11 NF__ENDDEF

The function `NF__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 2.16 [NF_SET_FILL], page 22). 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. See section "File Structure and Performance" in *NetCDF Users' Guide*.

Caution: this function exposes internals of the netcdf version 1 file format. Users should use `nc_enddef` in most circumstances. This function may not be available on 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 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 `minfree` parameters allow one to control costs of future calls to `nc_redef`, `nc_enddef` by requesting that `minfree` bytes be available at the end of the section.

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 `ALIGN_CHUNK` tells the library to use the `chunksize` (see above) as the `align` parameter.

The file format requires mod 4 alignment, so the `align` parameters are silently rounded up to multiples of 4. The usual call,

```
nc_enddef(ncid);
```

is equivalent to

```
nc_enddef(ncid, 0, 4, 0, 4);
```

The file format does not contain a "record size" value, this is calculated from the sizes of the record variables. This unfortunate fact prevents us from providing `minfree` and alignment control of the "records" in a netcdf file. If you add a variable which has an unlimited dimension, the third section will always be copied with the new variable added.

Usage

```
INTEGER FUNCTION NF_ENDDEF(INTEGER NCID, INTEGER H_MINFREE, INTEGER V_ALIGN,
                           INTEGER V_MINFREE, INTEGER R_ALIGN)
```

NCID NetCDF ID, from a previous call to `NF_OPEN` or `NF_CREATE`.

H_MINFREE Sets the pad at the end of the "header" section.

V_ALIGN Controls the alignment of the beginning of the data section for fixed size variables.

V_MINFREE Sets the pad at the end of the data section for fixed size variables.

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

Errors

`NF__ENDDEF` returns the value `NF_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 `NF__ENDDEF` to finish the definitions of a new netCDF dataset named `foo.nc` and put it into data mode:

```
INCLUDE 'netcdf.inc'
...
INTEGER NCID, STATUS, H_MINFREE, V_ALIGN, V_MINFREE, R_ALIGN
...
STATUS = NF_CREATE('foo.nc', NF_NOCLOBBER, NCID)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)

... ! create dimensions, variables, attributes

H_MINFREE = 512
V_ALIGN = 512
V_MINFREE = 512
R_ALIGN = 512
STATUS = NF_ENDDEF(NCID, H_MINFREE, V_ALIGN, V_MINFREE, R_ALIGN)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
```

2.12 NF_CLOSE

The function `NF_CLOSE` closes an open netCDF dataset. If the dataset is in define mode, `NF_ENDDEF` will be called before closing. (In this case, if `NF_ENDDEF` returns an error, `NF_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

```
INTEGER FUNCTION NF_CLOSE(INTEGER NCID)
```

NCID NetCDF ID, from a previous call to `NF_OPEN` or `NF_CREATE`.

Errors

NF_CLOSE returns the value NF_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 NF_ENDDEF failed.
- The specified netCDF ID does not refer to an open netCDF dataset.

Example

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

```

INCLUDE 'netcdf.inc'
...
INTEGER NCID, STATUS
...
STATUS = NF_CREATE('foo.nc', NF_NO_CLOBBER, NCID)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)

... ! create dimensions, variables, attributes

STATUS = NF_CLOSE(NCID)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)

```

2.13 NF_INQ Family

Members of the NF_INQ family of functions return information about an open netCDF dataset, given its netCDF ID. Dataset inquire functions may be called from either define mode or data mode. The first function, NF_INQ, returns values for 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. The other functions in the family each return just one of these items of information.

For FORTRAN, these functions include NF_INQ, NF_INQ_NDIMS, NF_INQ_NVARS, NF_INQ_NATTS, and NF_INQ_UNLIMDIM. An additional function, NF_INQ_FORMAT, returns the (rarely needed) format version.

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

Usage

```

INTEGER FUNCTION NF_INQ          (INTEGER NCID, INTEGER ndims,
                                INTEGER nvars, INTEGER ngatts,
                                INTEGER unlimdimid)
INTEGER FUNCTION NF_INQ_NDIMS   (INTEGER NCID, INTEGER ndims)
INTEGER FUNCTION NF_INQ_NVARS   (INTEGER NCID, INTEGER nvars)
INTEGER FUNCTION NF_INQ_NATTS   (INTEGER NCID, INTEGER ngatts)
INTEGER FUNCTION NF_INQ_UNLIMDIM (INTEGER NCID, INTEGER unlimdimid)
INTEGER FUNCTION NF_INQ_FORMAT  (INTEGER NCID, INTEGER format)

```

<code>NCID</code>	NetCDF ID, from a previous call to <code>NF_OPEN</code> or <code>NF_CREATE</code> .
<code>ndims</code>	Returned number of dimensions defined for this netCDF dataset.
<code>nvars</code>	Returned number of variables defined for this netCDF dataset.
<code>ngatts</code>	Returned number of global attributes defined for this netCDF dataset.
<code>unlimdimid</code>	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.
<code>format</code>	Returned format version, one of <code>NF_FORMAT_CLASSIC</code> , <code>NF_FORMAT_64BIT</code> , <code>NF_FORMAT_NETCDF4</code> , <code>NF_FORMAT_NETCDF4_CLASSIC</code> .

Errors

All members of the `NF_INQ` family return the value `NF_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 `NF_INQ` to find out about a netCDF dataset named `foo.nc`:

```

INCLUDE 'netcdf.inc'
...
INTEGER STATUS, NCID, NDIMS, NVAR, NGATT, UNLIMDIMID
...
STATUS = NF_OPEN('foo.nc', NF_NOWRITE, NCID)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
...
STATUS = NF_INQ(NCID, NDIMS, NVAR, NGATT, UNLIMDIMID)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)

```

2.14 NF_SYNC

The function `NF_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 `NF_SYNC`; to accomplish this, the reading process must call `NF_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 `NF_SYNC` after writing and the readers call `NF_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 `NF_SHARE` flag, and then it will not be necessary to call `NF_SYNC` at all. However, the `NF_SYNC` function still provides finer granularity than the `NF_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 `NF_SHARE` flag. Use of the `NF_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 `NF_SYNC` before any subsequent access.

When calling `NF_SYNC`, the netCDF dataset must be in data mode. A netCDF dataset in define mode is synchronized to disk only when `NF_ENDDEF` is called. A process that is reading a netCDF dataset that another process is writing may call `NF_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

```
INTEGER FUNCTION NF_SYNC(INTEGER NCID)
```

`NCID` NetCDF ID, from a previous call to `NF_OPEN` or `NF_CREATE`.

Errors

`NF_SYNC` returns the value `NF_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 `NF_SYNC` to synchronize the disk writes of a netCDF dataset named `foo.nc`:

```
INCLUDE 'netcdf.inc'
...
INTEGER STATUS, NCID
```

```

...
STATUS = NF_OPEN('foo.nc', NF_WRITE, NCID)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
...
! write data or change attributes
...
STATUS = NF_SYNC(NCID)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)

```

2.15 NF_ABORT

You no longer need to call this function, since it is called automatically by `NF_CLOSE` in case the dataset is in define mode and something goes wrong with committing the changes. The function `NF_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 `NF_REDEF`, the netCDF dataset is restored to its state before definition mode was entered and the dataset is closed.

Usage

```
INTEGER FUNCTION NF_ABORT(INTEGER NCID)
```

`NCID` NetCDF ID, from a previous call to `NF_OPEN` or `NF_CREATE`.

Errors

`NF_ABORT` returns the value `NF_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 `NF_ABORT` to back out of redefinitions of a dataset named `foo.nc`:

```

INCLUDE 'netcdf.inc'
...
INTEGER STATUS, NCID, LATID
...
STATUS = NF_OPEN('foo.nc', NF_WRITE, NCID)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
...
STATUS = NF_REDEF(NCID)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
...
STATUS = NF_DEF_DIM(NCID, 'LAT', 18, LATID)
IF (STATUS .NE. NF_NOERR) THEN ! dimension definition failed
    CALL HANDLE_ERR(STATUS)

```

```

        STATUS = NF_ABORT(NCID) ! abort redefinitions
        IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
    ENDIF
    ...

```

2.16 NF_SET_FILL

This function is intended for advanced usage, to optimize writes under some circumstances described below. The function `NF_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 `NF_FILL` or `NF_NOFILL`. The default behavior corresponding to `NF_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 4.17 \[Fill Values\], page 60](#), for more information on the use of fill values. See [Section 5.2 \[Attribute Conventions\], page 63](#), for information about how to define your own fill values.

The behavior corresponding to `NF_NOFILL` overrides the default behavior of prefilling 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 `NF_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 `NF_SET_FILL` again to explicitly set the fill mode to `NF_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 `NF_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 `NF_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

```
INTEGER FUNCTION NF_SET_FILL(INTEGER NCID, INTEGER FILLMODE,
                             INTEGER old_mode)
```

NCID NetCDF ID, from a previous call to `NF_OPEN` or `NF_CREATE`.

FILLMODE Desired fill mode for the dataset, either `NF_NOFILL` or `NF_FILL`.

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

Errors

`NF_SET_FILL` returns the value `NF_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 `NF_NOFILL` nor `NF_FILL`.

Example

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

```
INCLUDE 'netcdf.inc'
...
INTEGER NCID, STATUS, OMODE
...
STATUS = NF_OPEN('foo.nc', NF_WRITE, NCID)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
...
! write data with default prefilling behavior
...
STATUS = NF_SET_FILL(NCID, NF_NOFILL, OMODE)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
...
! write data with no prefilling
...
...
```

2.17 NF_SET_DEFAULT_FORMAT

This function is intended for advanced users.

Starting in version 3.6, netCDF introduced a new data format, the first change in the underlying binary data format since the netCDF interface was released. The new format, 64-bit offset format, was introduced to greatly relax the limitations on creating very large files.

Users are warned that creating files in the 64-bit offset format makes them unreadable by the netCDF library prior to version 3.6.0. For reasons of compatibility, users should continue to create files in netCDF classic format.

Users who do want to use 64-bit offset format files can create them directory from `NF_CREATE`, using the proper `cmode` flag. (see [Section 2.5 \[NF_CREATE\]](#), page 9).

The function `NF_SET_DEFAULT_FORMAT` allows the user to change the format of the netCDF file to be created by future calls to `NF_CREATE` without changing the `cmode` flag.

This allows the user to convert a program to use 64-bit offset formation without changing all calls the `NF_CREATE`. See [section “Large File Support”](#) in *The NetCDF Users Guide*.

Once the default format is set, all future created files will be in the desired format.

Two constants are provided in the `netcdf.inc` file to be used with this function, `nf_format_64bit` and `nf_format_classic`.

Using `NF_CREATE` with a `cmode` including `nf_64bit_offset` overrides the default format, and creates a 64-bit offset file.

Usage

```
INTEGER FUNCTION NF_SET_DEFAULT_FORMAT(INTEGER FORMAT, INTEGER OLD_FORMAT)
FORMAT      Either nf_format_64bit or nf_format_classic.
OLD_FORMAT
            The default format at the time the function is called is returned here.
```

Errors

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

- An `NC_EINVAL` error is returned if an invalid default format is specified.

Example

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

```
INCLUDE 'netcdf.inc'
...
INTEGER STATUS, OLD_FORMAT
...
STATUS = NF_SET_DEFAULT_FORMAT(nf_format_64bit, OLD_FORMAT)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
...

```

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 (100) to the number of dimensions that can be defined in a single netCDF dataset. The limit is the value of the predefined macro `NF_MAX_DIMS`. The purpose of the limit is to make writing generic applications simpler. They need only provide an array of `NF_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 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 `NF_DEF_DIM`

The function `NF_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

```
INTEGER FUNCTION NF_DEF_DIM (INTEGER NCID, CHARACTER*(*) NAME,
                             INTEGER LEN, INTEGER dimid)
```

NCID NetCDF ID, from a previous call to `NF_OPEN` or `NF_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>NF_UNLIMITED</code> .
dimid	Returned dimension ID.

Errors

`NF_DEF_DIM` returns the value `NF_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 `NF_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`:

```

INCLUDE 'netcdf.inc'
...
INTEGER STATUS, NCID, LATID, RECID
...
STATUS = NF_CREATE('foo.nc', NF_NOCLOBBER, NCID)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
...
STATUS = NF_DEF_DIM(NCID, 'lat', 18, LATID)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
STATUS = NF_DEF_DIM(NCID, 'rec', NF_UNLIMITED, RECID)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)

```

3.3 NF_INQ_DIMID

The function `NF_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

```

INTEGER FUNCTION NF_INQ_DIMID (INTEGER NCID, CHARACTER*(*) NAME,
                              INTEGER dimid)

```

NCID	NetCDF ID, from a previous call to <code>NF_OPEN</code> or <code>NF_CREATE</code> .
NAME	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.
dimid	Returned dimension ID.

Errors

NF_INQ_DIMID returns the value NF_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 NF_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:

```
INCLUDE 'netcdf.inc'
...
INTEGER STATUS, NCID, LATID
...
STATUS = NF_OPEN('foo.nc', NF_NOWRITE, NCID)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
...
STATUS = NF_INQ_DIMID(NCID, 'lat', LATID)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
```

3.4 NF_INQ_DIM Family

This family of functions returns 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.

The functions in this family include NF_INQ_DIM, NF_INQ_DIMNAME, and NF_INQ_DIMLEN. The function NF_INQ_DIM returns all the information about a dimension; the other functions each return just one item of information.

Usage

```
INTEGER FUNCTION NF_INQ_DIM      (INTEGER NCID, INTEGER DIMID,
                                CHARACTER*(*) name, INTEGER len)
INTEGER FUNCTION NF_INQ_DIMNAME (INTEGER NCID, INTEGER DIMID,
                                CHARACTER*(*) name)
INTEGER FUNCTION NF_INQ_DIMLEN  (INTEGER NCID, INTEGER DIMID,
                                INTEGER len)
```

NCID	NetCDF ID, from a previous call to NF_OPEN or NF_CREATE.
DIMID	Dimension ID, from a previous call to NF_INQ_DIMID or NF_DEF_DIM.
NAME	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 NF_MAX_NAME.
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 `NF_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 `NF_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`:

```

INCLUDE 'netcdf.inc'
...
INTEGER STATUS, NCID, LATID, LATLEN, RECID, NRECS
CHARACTER*(NF_MAX_NAME) LATNAM, RECNAME
...
STATUS = NF_OPEN('foo.nc', NF_NOWRITE, NCID)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
! get ID of unlimited dimension
STATUS = NF_INQ_UNLIMDIM(NCID, RECID)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
...
STATUS = NF_INQ_DIMID(NCID, 'lat', LATID)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
! get lat length
STATUS = NF_INQ_DIMLEN(NCID, LATID, LATLEN)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
! get unlimited dimension name and current length
STATUS = NF_INQ_DIM(NCID, RECID, RECNAME, NRECS)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)

```

3.5 NF_RENAME_DIM

The function `NF_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

```

INTEGER FUNCTION NF_RENAME_DIM (INTEGER NCID, INTEGER DIMID,
                               CHARACTER*(*) NAME)

```

<code>NCID</code>	NetCDF ID, from a previous call to <code>NF_OPEN</code> or <code>NF_CREATE</code> .
<code>DIMID</code>	Dimension ID, from a previous call to <code>NF_INQ_DIMID</code> or <code>NF_DEF_DIM</code> .
<code>NAME</code>	New dimension name.

Errors

NF_RENAME_DIM returns the value NF_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 NF_RENAME_DIM to rename the dimension lat to latitude in an existing netCDF dataset named foo.nc:

```
INCLUDE 'netcdf.inc'
...
INTEGER STATUS, NCID, LATID
...
STATUS = NF_OPEN('foo.nc', NF_WRITE, NCID)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
...
! put in define mode to rename dimension
STATUS = NF_REDEF(NCID)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
STATUS = NF_INQ_DIMID(NCID, 'lat', LATID)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
STATUS = NF_RENAME_DIM(NCID, LATID, 'latitude')
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
! leave define mode
STATUS = NF_ENDDEF(NCID)
IF (STATUS .NE. NF_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 (see [Chapter 5 \[Attributes\], page 63](#)) 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	NF_BYTE	8

char	NF_CHAR	8
short	NF_SHORT	16
int	NF_INT	32
float	NF_FLOAT	32
double	NF_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 parameter for use in netCDF functions (the parameters are defined in the netCDF FORTRAN include-file netcdf.inc). 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: NF_DEF_VAR

The function NF_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

```
INTEGER FUNCTION NF_DEF_VAR(INTEGER NCID, CHARACTER*(*) NAME,
                           INTEGER XTYPE, INTEGER NVDIMS,
                           INTEGER VDIMS(*), INTEGER varid)
```

NCID	NetCDF ID, from a previous call to NF_OPEN or NF_CREATE.
NAME	Variable name. Must begin with an alphabetic character, followed by zero or more alphanumeric characters including the underscore ('_'). Case is significant.
XTYPE	One of the set of predefined netCDF external data types. The type of this parameter, NF_TYPE, is defined in the netCDF header file. The valid netCDF external data types are NF_BYTE, NF_CHAR, NF_SHORT, NF_INT, NF_FLOAT, and NF_DOUBLE.
NVDIMS	Number of dimensions for the variable. For example, 2 specifies a matrix, 1 specifies a vector, and 0 means the variable is a scalar with no dimensions. Must not be negative or greater than the predefined constant NF_MAX_VAR_DIMS.
VDIMS	Vector of ndims dimension IDs corresponding to the variable dimensions. If the ID of the unlimited dimension is included, it must be last. This argument is ignored if ndims is 0.
varid	Returned variable ID.

Errors

NF_DEF_VAR returns the value NF_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 NF_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 NF_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 NF_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:

```

INCLUDE 'netcdf.inc'
...
INTEGER STATUS, NCID
INTEGER LATDIM, LONDIM, TIMDIM ! dimension IDs
INTEGER RHID ! variable ID
INTEGER RHDIMS(3) ! variable shape
...
STATUS = NF_CREATE ('foo.nc', NF_NO_CLOBBER, NCID)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
...
                                ! define dimensions
STATUS = NF_DEF_DIM(NCID, 'lat', 5, LATDIM)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
STATUS = NF_DEF_DIM(NCID, 'lon', 10, LONDIM)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
STATUS = NF_DEF_DIM(NCID, 'time', NF_UNLIMITED, TIMDIM)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
...
                                ! define variable
RHDIMS(1) = LONDIM
RHDIMS(2) = LATDIM
RHDIMS(3) = TIMDIM
STATUS = NF_DEF_VAR (NCID, 'rh', NF_DOUBLE, 3, RHDIMS, RHID)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)

```

4.4 Get a Variable ID from Its Name: `NF_INQ_VARID`

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

Usage

```
INTEGER FUNCTION NF_INQ_VARID(INTEGER NCID, CHARACTER*(*) NAME,
                             INTEGER varid)
```

`NCID` NetCDF ID, from a previous call to `NF_OPEN` or `NF_CREATE`.

`NAME` Variable name for which ID is desired.

`varid` Returned variable ID.

Errors

`NF_INQ_VARID` returns the value `NF_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 `NF_INQ_VARID` to find out the ID of a variable named `rh` in an existing netCDF dataset named `foo.nc`:

```
INCLUDE 'netcdf.inc'
...
INTEGER STATUS, NCID, RHID
...
STATUS = NF_OPEN ('foo.nc', NF_NOWRITE, NCID)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
...
STATUS = NF_INQ_VARID (NCID, 'rh', RHID)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
```

4.5 Get Information about a Variable from Its ID: `NF_INQ_VAR` family

A family of functions that 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.

The function `NF_INQ_VAR` returns all the information about a netCDF variable, given its ID. The other functions each return just one item of information about a variable.

These other functions include `NF_INQ_VARNAME`, `NF_INQ_VARTYPE`, `NF_INQ_VARNDIMS`, `NF_INQ_VARDIMID`, and `NF_INQ_VARNATTS`.

Usage

<code>INTEGER FUNCTION NF_INQ_VAR</code>	<code>(INTEGER NCID, INTEGER VARID, CHARACTER*(*) name, INTEGER xtype, INTEGER ndims, INTEGER dimids(*), INTEGER natts)</code>
<code>INTEGER FUNCTION NF_INQ_VARNAME</code>	<code>(INTEGER NCID, INTEGER VARID, CHARACTER*(*) name)</code>
<code>INTEGER FUNCTION NF_INQ_VARTYPE</code>	<code>(INTEGER NCID, INTEGER VARID, INTEGER xtype)</code>
<code>INTEGER FUNCTION NF_INQ_VARNDIMS</code>	<code>(INTEGER NCID, INTEGER VARID, INTEGER ndims)</code>
<code>INTEGER FUNCTION NF_INQ_VARDIMID</code>	<code>(INTEGER NCID, INTEGER VARID, INTEGER dimids(*))</code>
<code>INTEGER FUNCTION NF_INQ_VARNATTS</code>	<code>(INTEGER NCID, INTEGER VARID, INTEGER natts)</code>
NCID	NetCDF ID, from a previous call to <code>NF_OPEN</code> or <code>NF_CREATE</code> .
VARID	Variable ID.
NAME	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>NF_MAX_NAME</code> .
xtype	Returned variable type, one of the set of predefined netCDF external data types. The type of this parameter, <code>NF_TYPE</code> , is defined in the netCDF header file. The valid netCDF external data types are <code>NF_BYTE</code> , <code>NF_CHAR</code> , <code>NF_SHORT</code> , <code>NF_INT</code> , <code>NF_FLOAT</code> , AND <code>NF_DOUBLE</code> .
ndims	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.
dimids	Returned vector of *ndimsp dimension IDs corresponding to the variable dimensions. The caller must allocate enough space for a vector of at least *ndimsp integers to be returned. The maximum possible number of dimensions for a variable is given by the predefined constant <code>NF_MAX_VAR_DIMS</code> .
natts	Returned number of variable attributes assigned to this variable.

These functions return the value `NF_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 `NF_INQ_VAR` to find out about a variable named `rh` in an existing netCDF dataset named `foo.nc`:

```
INCLUDE 'netcdf.inc'
...
```

```

INTEGER STATUS, NCID
INTEGER RHID           ! variable ID
CHARACTER*31 RHNAME    ! variable name
INTEGER RHTYPE         ! variable type
INTEGER RHN           ! number of dimensions
INTEGER RHDIMS(NF_MAX_VAR_DIMS) ! variable shape
INTEGER RHNATT         ! number of attributes
...
STATUS = NF_OPEN ('foo.nc', NF_NOWRITE, NCID)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
...
STATUS = NF_INQ_VARID (NCID, 'rh', RHID) ! get ID
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
STATUS = NF_INQ_VAR (NCID, RHID, RHNAME, RHTYPE, RHN, RHDIMS, RHNATT)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)

```

4.6 Write a Single Data Value: `NF_PUT_VAR1_` type

The functions `NF_PUT_VAR1_` type put a single data value of the specified type into a variable of an open netCDF dataset that is in data mode. Inputs are the netCDF ID, the variable ID, an index that specifies which value to add or alter, and the data value. The value is converted to the external data 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

```

INTEGER FUNCTION  NF_PUT_VAR1_TEXT(INTEGER NCID, INTEGER VARID,
                                   INTEGER INDEX(*), CHARACTER CHVAL)
INTEGER FUNCTION  NF_PUT_VAR1_INT1(INTEGER NCID, INTEGER VARID,
                                   INTEGER INDEX(*), INTEGER*1 I1VAL)
INTEGER FUNCTION  NF_PUT_VAR1_INT2(INTEGER NCID, INTEGER VARID,
                                   INTEGER INDEX(*), INTEGER*2 I2VAL)
INTEGER FUNCTION  NF_PUT_VAR1_INT (INTEGER NCID, INTEGER VARID,
                                   INTEGER INDEX(*), INTEGER IVAL)
INTEGER FUNCTION  NF_PUT_VAR1_REAL(INTEGER NCID, INTEGER VARID,
                                   INTEGER INDEX(*), REAL      RVAL)
INTEGER FUNCTION  NF_PUT_VAR1_DOUBLE(INTEGER NCID, INTEGER VARID,
                                     INTEGER INDEX(*), DOUBLE   DVAL)

```

NCID NetCDF ID, from a previous call to `NF_OPEN` or `NF_CREATE`.

VARID Variable ID.

INDEX The index of the data value to be written. The indices are relative to 1, so for example, the first data value of a two-dimensional variable would have index (1,1). The elements of index must correspond to the variable's dimensions. Hence, if

the variable uses the unlimited dimension, the last index would correspond to the record number.

CHVAL
 I1VAL
 I2VAL
 IVAL
 RVAL
 DVAL Pointer to the data value to be written. If the type of data values differs from the netCDF variable type, type conversion will occur. See [section “Type Conversion”](#) in *The NetCDF Users Guide*.

Errors

NF_PUT_VAR1_ type returns the value NF_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 NF_PUT_VAR1_DOUBLE 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:

```
INCLUDE 'netcdf.inc'
...
INTEGER STATUS           ! error status
INTEGER NCID
INTEGER RHID             ! variable ID
INTEGER RHINDX(3)       ! where to put value
DATA RHINDX /4, 3, 2/
...
STATUS = NF_OPEN ('foo.nc', NF_WRITE, NCID)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
...
STATUS = NF_INQ_VARID (NCID, 'rh', RHID) ! get ID
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
STATUS = NF_PUT_VAR1_DOUBLE (NCID, RHID, RHINDX, 0.5)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
```

4.7 Write an Entire Variable: `NF_PUT_VAR_` type

The `NF_PUT_VAR_` type family of functions write all the values of a variable into a netCDF variable of an open netCDF dataset. This is the simplest interface to use for writing a value in a scalar variable or whenever all the values of a multidimensional variable can all be written at once. The values to be written are associated with the netCDF variable by assuming that the last dimension of the netCDF variable varies fastest in the C interface. The values are converted to the external data 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

```

INTEGER FUNCTION NF_PUT_VAR_TEXT (INTEGER NCID, INTEGER VARID,
                                CHARACTER*(*) TEXT)
INTEGER FUNCTION NF_PUT_VAR_INT1 (INTEGER NCID, INTEGER VARID,
                                INTEGER*1 I1VALS(*))
INTEGER FUNCTION NF_PUT_VAR_INT2 (INTEGER NCID, INTEGER VARID,
                                INTEGER*2 I2VALS(*))
INTEGER FUNCTION NF_PUT_VAR_INT (INTEGER NCID, INTEGER VARID,
                                INTEGER IVALS(*))
INTEGER FUNCTION NF_PUT_VAR_REAL (INTEGER NCID, INTEGER VARID,
                                REAL RVALS(*))
INTEGER FUNCTION NF_PUT_VAR_DOUBLE (INTEGER NCID, INTEGER VARID,
                                DOUBLE DVALS(*))

```

NCID NetCDF ID, from a previous call to `NF_OPEN` or `NF_CREATE`.

VARID Variable ID.

TEXT

I1VALS

I2VALS

IVALS

RVALS

DVALS

The block of data values to be written. The data should be of the type appropriate for the function called. 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 *The NetCDF Users Guide*). The order in which the data will be written into the specified variable is with the first dimension varying fastest (like the ordinary FORTRAN convention).

Errors

Members of the `NF_PUT_VAR_` *type* family return the value `NF_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.
- One or more of the specified values are out of the range of values representable by the external data type of the variable.
- One or more of the specified values are out of the range of values representable by the external data type of the variable.
- The specified netCDF dataset 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 `NF_PUT_VAR_DOUBLE` to add or change all the values 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`, and that there are ten `lon` values, five `lat` values, and three `time` values.

```

INCLUDE 'netcdf.inc'
...
PARAMETER (TIMES=3, LATS=5, LONS=10) ! dimension lengths
INTEGER STATUS, NCID, TIMES
INTEGER RHID ! variable ID
DOUBLE RHVALS(LONS, LATS, TIMES)
...
STATUS = NF_OPEN ('foo.nc', NF_WRITE, NCID)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
...
STATUS = NF_INQ_VARID (NCID, 'rh', RHID)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
DO 10 ILO = 1, LONS
  DO 10 ILAT = 1, LATS
    DO 10 ITIME = 1, TIMES
      RHVALS(ILO, ILAT, ITIME) = 0.5
10 CONTINUE
STATUS = NF_PUT_var_DOUBLE (NCID, RHID, RHVALS)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)

```

4.8 Write an Array of Values: `NF_PUT_VARA_` *type*

The function `NF_PUT_VARA_` *type* writes values into a netCDF variable of an open netCDF dataset. The part of the netCDF variable to write is specified by giving a corner and a vector of edge lengths that refer to an array section of the netCDF variable. 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 interface. The netCDF dataset must be in data mode.

Usage

```

INTEGER FUNCTION NF_PUT_VARA_TEXT(INTEGER NCID, INTEGER VARID,
                                INTEGER START(*), INTEGER COUNT(*),
                                CHARACTER*(*) TEXT)
INTEGER FUNCTION NF_PUT_VARA_INT1(INTEGER NCID, INTEGER VARID,
                                INTEGER START(*), INTEGER COUNT(*),
                                INTEGER*1 I1VALS(*))
INTEGER FUNCTION NF_PUT_VARA_INT2(INTEGER NCID, INTEGER VARID,
                                INTEGER START(*), INTEGER COUNT(*),
                                INTEGER*2 I2VALS(*))
INTEGER FUNCTION NF_PUT_VARA_INT (INTEGER NCID, INTEGER VARID,
                                INTEGER START(*), INTEGER COUNT(*),
                                INTEGER IVALS(*))
INTEGER FUNCTION NF_PUT_VARA_REAL(INTEGER NCID, INTEGER VARID,
                                INTEGER START(*), INTEGER COUNT(*),
                                REAL RVALS(*))
INTEGER FUNCTION NF_PUT_VARA_DOUBLE(INTEGER NCID, INTEGER VARID,
                                INTEGER START(*), INTEGER COUNT(*),
                                DOUBLE DVALS(*))

```

NCID NetCDF ID, from a previous call to `NF_OPEN` or `NF_CREATE`.

VARID Variable ID.

START A vector of integers specifying the index in the variable where the first 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 length of `START` must be the same as the number of dimensions of the specified variable. The elements of `START` must correspond to the variable's dimensions in order. Hence, if the variable is a record variable, the last index would correspond to the starting record number for writing the data values.

COUNT A vector of integers specifying the edge lengths along each dimension of the block of data values to be written. To write a single value, for example, specify `COUNT` as (1, 1, ..., 1). The length of `COUNT` is the number of dimensions of the specified variable. The elements of `COUNT` correspond 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.

TEXT

I1VALS

I2VALS

IVALS

RVALS

DVALS

The block of data values to be written. The data should be of the type appropriate for the function called. 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 *The NetCDF Users Guide*).

Errors

`NF_PUT_VARA_` *type* returns the value `NF_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 corner 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 edge lengths added to the specified corner would have referenced data out of range for the rank of the specified variable. For example, an edge length that is larger than the corresponding dimension length minus the corner index will cause an error.
- One or more of the specified values are out of the range of values representable by the external data type of the variable.
- The specified netCDF dataset 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 `NF_PUT_VARA_DOUBLE` to add or change all the values 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 time, lat, and lon, and that there are three time values, five lat values, and ten lon values.

```

INCLUDE 'netcdf.inc'
...
PARAMETER (NDIMS=3)           ! number of dimensions
PARAMETER (TIMES=3, LATS=5, LONS=10) ! dimension lengths
INTEGER STATUS, NCID, TIMES
INTEGER RHID                   ! variable ID
INTEGER START(NDIMS), COUNT(NDIMS)
DOUBLE RHVALS(LONS, LATS, TIMES)
DATA START /1, 1, 1/          ! start at first value
DATA COUNT /LONS, LATS, TIMES/
...
STATUS = NF_OPEN ('foo.nc', NF_WRITE, NCID)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
...
STATUS = NF_INQ_VARID (NCID, 'rh', RHID)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
DO 10 ILO = 1, LONS
  DO 10 ILAT = 1, LATS
    DO 10 ITIME = 1, TIMES
      RHVALS(ILO, ILAT, ITIME) = 0.5
10 CONTINUE
STATUS = NF_PUT_VARA_DOUBLE (NCID, RHID, START, COUNT, RHVALS)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)

```

4.9 NF_PUT_VARS_ *type*

Each member of the family of functions `NF_PUT_VARS_ type` writes a subsampled (strided) array section of values into a netCDF variable of an open netCDF dataset. The subsampled array section is specified by giving a corner, a vector of counts, and a stride vector. The netCDF dataset must be in data mode.

Usage

```

INTEGER FUNCTION NF_PUT_VARS_TEXT (INTEGER NCID, INTEGER VARID,
                                   INTEGER START(*), INTEGER COUNT(*),
                                   INTEGER STRIDE(*), CHARACTER*(*) TEXT)
INTEGER FUNCTION NF_PUT_VARS_INT1 (INTEGER NCID, INTEGER VARID,
                                   INTEGER START(*), INTEGER COUNT(*),
                                   INTEGER STRIDE(*), INTEGER*1 I1VALS(*))
INTEGER FUNCTION NF_PUT_VARS_INT2 (INTEGER NCID, INTEGER VARID,
                                   INTEGER START(*), INTEGER COUNT(*),
                                   INTEGER STRIDE(*), INTEGER*2 I2VALS(*))
INTEGER FUNCTION NF_PUT_VARS_INT (INTEGER NCID, INTEGER VARID,
                                   INTEGER START(*), INTEGER COUNT(*),
                                   INTEGER STRIDE(*), INTEGER IVALS(*))
INTEGER FUNCTION NF_PUT_VARS_REAL (INTEGER NCID, INTEGER VARID,
                                   INTEGER START(*), INTEGER COUNT(*),
                                   INTEGER STRIDE(*), REAL RVALS(*))
INTEGER FUNCTION NF_PUT_VARS_DOUBLE (INTEGER NCID, INTEGER VARID,
                                     INTEGER START(*), INTEGER COUNT(*),
                                     INTEGER STRIDE(*), DOUBLE DVALS(*))

```

- NCID** NetCDF ID, from a previous call to `NF_OPEN` or `NF_CREATE`.
- VARID** Variable ID.
- START** A vector of integers specifying the index in the variable where the first 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.
- 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.
- 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.).

TEXT

I1VALS

I2VALS

IVALS

RVALS

DVALS The block of data values to be written. The data should be of the type appropriate for the function called. 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 *The NetCDF Users Guide*).

Errors

NF_PUT_VARS_ *type* returns the value NF_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 start, count and stride generate an index which is out of range.
- One or more of the specified values are 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 of using NF_PUT_VARS_REAL to write – from an internal array – every other point of a netCDF variable named rh which is described by the FORTRAN declaration REAL RH(6,4) (note the size of the dimensions):

```

INCLUDE 'netcdf.inc'
...
PARAMETER (NDIM=2)    ! rank of netCDF variable
INTEGER NCID          ! netCDF dataset ID
INTEGER STATUS        ! return code
INTEGER RHID          ! variable ID
INTEGER START(NDIM)  ! netCDF variable start point
INTEGER COUNT(NDIM)  ! size of internal array
INTEGER STRIDE(NDIM) ! netCDF variable subsampling intervals
REAL RH(3,2)          ! note subsampled sizes for netCDF variable
                      ! dimensions
DATA START    /1, 1/ ! start at first netCDF variable value
DATA COUNT    /3, 2/ ! size of internal array: entire (subsampled)
                      ! netCDF variable
DATA STRIDE   /2, 2/ ! access every other netCDF element
...
STATUS = NF_OPEN('foo.nc', NF_WRITE, NCID)

```

```

IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
...
STATUS = NF_INQ_VARID(NCID, 'rh', RHID)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
...
STATUS = NF_PUT_VARS_REAL(NCID, RHID, START, COUNT, STRIDE, RH)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)

```

4.10 NF_PUT_VARM_ *type*

The `NF_PUT_VARM_` *type* family of functions writes a mapped array section of values into a netCDF variable of an open netCDF dataset. The mapped array section is specified by giving a corner, a vector of counts, a stride vector, and an index mapping vector. The index mapping vector is 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. No assumptions are made about the ordering or length of the dimensions of the data array. The netCDF dataset must be in data mode.

Usage

```

INTEGER FUNCTION NF_PUT_VARM_TEXT (INTEGER NCID, INTEGER VARID,
                                   INTEGER START(*), INTEGER COUNT(*),
                                   INTEGER STRIDE(*), INTEGER IMAP(*),
                                   CHARACTER*(*) TEXT)

INTEGER FUNCTION NF_PUT_VARM_INT1 (INTEGER NCID, INTEGER VARID,
                                   INTEGER START(*), INTEGER COUNT(*),
                                   INTEGER STRIDE(*), INTEGER IMAP(*),
                                   INTEGER*1 I1VALS(*))

INTEGER FUNCTION NF_PUT_VARM_INT2 (INTEGER NCID, INTEGER VARID,
                                   INTEGER START(*), INTEGER COUNT(*),
                                   INTEGER STRIDE(*), INTEGER IMAP(*),
                                   INTEGER*2 I2VALS(*))

INTEGER FUNCTION NF_PUT_VARM_INT (INTEGER NCID, INTEGER VARID,
                                   INTEGER START(*), INTEGER COUNT(*),
                                   INTEGER STRIDE(*), INTEGER IMAP(*),
                                   INTEGER IVALS(*))

INTEGER FUNCTION NF_PUT_VARM_REAL (INTEGER NCID, INTEGER VARID,
                                   INTEGER START(*), INTEGER COUNT(*),
                                   INTEGER STRIDE(*), INTEGER IMAP(*),
                                   REAL RVALS(*))

INTEGER FUNCTION NF_PUT_VARM_DOUBLE(INTEGER NCID, INTEGER VARID,
                                   INTEGER START(*), INTEGER COUNT(*),
                                   INTEGER STRIDE(*), INTEGER IMAP(*),
                                   DOUBLE DVALS(*))

```

`NCID` NetCDF ID, from a previous call to `NF_OPEN` or `NF_CREATE`.

VARID	Variable ID.
START	A vector of integers specifying the index in the variable where the first 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.
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.
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.).
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 (IMAP(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 (the distance between internal elements that occupy adjacent memory locations is 1 and not the element's byte-length as in netCDF 2).
TEXT	
I1VALS	
I2VALS	
IVALS	
RVALS	
DVALS	The data values to be written. The data should be of the type appropriate for the function called. 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 <i>The NetCDF Users Guide</i>).

Errors

NF_PUT_VARM_ type returns the value NF_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 START, COUNT, and STRIDE generate an index which is out of range. Note that no error checking is possible on the imap vector.

- One or more of the specified values are 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

The following IMAP vector maps in the trivial way a 2x3x4 netCDF variable and an internal array of the same shape:

```
REAL A(2,3,4)      ! same shape as netCDF variable
INTEGER IMAP(3)
DATA IMAP /1, 2, 6/ ! netCDF dimension      inter-element distance
                   ! -----
                   ! most rapidly varying   1
                   ! intermediate           2 (=IMAP(1)*2)
                   ! most slowly varying    6 (=IMAP(2)*3)
```

Using the IMAP vector above with `NF_PUT_VARM_REAL` obtains the same result as simply using `NF_PUT_VAR_REAL`.

Here is an example of using `NF_PUT_VARM_REAL` to write – from a transposed, internal array – a netCDF variable named `rh` which is described by the FORTRAN declaration `REAL RH(4,6)` (note the size and order of the dimensions):

```
INCLUDE 'netcdf.inc'
...
PARAMETER (NDIM=2) ! rank of netCDF variable
INTEGER NCID       ! netCDF ID
INTEGER STATUS     ! return code
INTEGER RHID       ! variable ID
INTEGER START(NDIM) ! netCDF variable start point
INTEGER COUNT(NDIM) ! size of internal array
INTEGER STRIDE(NDIM) ! netCDF variable subsampling intervals
INTEGER IMAP(NDIM) ! internal array inter-element distances
REAL RH(6,4)       ! note transposition of netCDF variable dimensions
DATA START /1, 1/ ! start at first netCDF variable element
DATA COUNT /4, 6/ ! entire netCDF variable; order corresponds
                  ! to netCDF variable -- not internal array
DATA STRIDE /1, 1/ ! sample every netCDF element
DATA IMAP /6, 1/ ! would be /1, 4/ if not transposing

STATUS = NF_OPEN('foo.nc', NF_WRITE, NCID)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
...
STATUS = NF_INQ_VARID(NCID, 'rh', RHID)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
...
STATUS = NF_PUT_VARM_REAL(NCID, RHID, START, COUNT, STRIDE, IMAP, RH)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
```



```

INTEGER FUNCTION  NF_GET_VAR1_REAL(INTEGER NCID, INTEGER VARID,
                                   INTEGER INDEX(*), REAL      RVAL)
INTEGER FUNCTION  NF_GET_VAR1_DOUBLE(INTEGER NCID, INTEGER VARID,
                                     INTEGER INDEX(*), DOUBLE   DVAL)

```

NCID NetCDF ID, from a previous call to `NF_OPEN` or `NF_CREATE`.

VARID Variable ID.

INDEX The index of the data value to be read. The indices are relative to 1, so for example, the first data value of a two-dimensional variable has index (1,1). The elements of index correspond to the variable's dimensions. Hence, if the variable is a record variable, the last index is the record number.

chval

i1val

i2val

ival

rval

dval The location into which the data value will be read. You cannot get CHARACTER data from a numeric variable or numeric data from a character variable. For numeric data, if the type of data differs from the netCDF variable type, type conversion will occur. (see section “Type Conversion” in *The NetCDF Users Guide*).

Errors

`NF_GET_VAR1_` type returns the value `NF_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 value is out of the range of values representable by the desired data type.
- 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 `NF_GET_VAR1_DOUBLE` to get the (4,3,2) element of the variable named `rh` 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 get the value of `rh` that corresponds to the fourth `lon` value, the third `lat` value, and the second `time` value:

```

INCLUDE 'netcdf.inc'
...
INTEGER STATUS, NCID
INTEGER RHID            ! variable ID
INTEGER RHINDX(3)       ! where to get value

```

```

DOUBLE PRECISION RHVAL ! put it here
DATA RHINDX /4, 3, 2/
...
STATUS = NF_OPEN ('foo.nc', NF_NOWRITE, NCID)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
...
STATUS = NF_INQ_VARID (NCID, 'rh', RHID)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
STATUS = NF_GET_VAR1_DOUBLE (NCID, RHID, RHINDX, RHVAL)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)

```

4.12 NF_GET_VAR_ type

The members of the `NF_GET_VAR_ type` family of functions read all the values from a netCDF variable of an open netCDF dataset. This is the simplest interface to use for reading the value of a scalar variable or when all the values of a multidimensional variable can be read at once. The values are read into consecutive locations with the first dimension varying fastest. The netCDF dataset must be in data mode.

Usage

```

INTEGER FUNCTION NF_GET_VAR_TEXT (INTEGER NCID, INTEGER VARID,
                                CHARACTER*(*) text)
INTEGER FUNCTION NF_GET_VAR_INT1 (INTEGER NCID, INTEGER VARID,
                                INTEGER*1 i1vals(*))
INTEGER FUNCTION NF_GET_VAR_INT2 (INTEGER NCID, INTEGER VARID,
                                INTEGER*2 i2vals(*))
INTEGER FUNCTION NF_GET_VAR_INT (INTEGER NCID, INTEGER VARID,
                                INTEGER ivals(*))
INTEGER FUNCTION NF_GET_VAR_REAL (INTEGER NCID, INTEGER VARID,
                                REAL rvals(*))
INTEGER FUNCTION NF_GET_VAR_DOUBLE (INTEGER NCID, INTEGER VARID,
                                DOUBLE dvals(*))

```

NCID NetCDF ID, from a previous call to `NF_OPEN` or `NF_CREATE`.

VARID Variable ID.

text

i1vals

i2vals

ivals

rvals

dvals

The block of data values to be read. The data should be of the type appropriate for the function called. 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 *The NetCDF Users Guide*).


```

                                INTEGER*2 i2vals(*)
INTEGER FUNCTION NF_GET_VARA_INT (INTEGER NCID, INTEGER VARID,
                                INTEGER START(*), INTEGER COUNT(*),
                                INTEGER ival(*))
INTEGER FUNCTION NF_GET_VARA_REAL(INTEGER NCID, INTEGER VARID,
                                INTEGER START(*), INTEGER COUNT(*),
                                REAL rvals(*))
INTEGER FUNCTION NF_GET_VARA_DOUBLE(INTEGER NCID, INTEGER VARID,
                                INTEGER START(*), INTEGER COUNT(*),
                                DOUBLE dvals(*))

```

NCID NetCDF ID, from a previous call to `NF_OPEN` or `NF_CREATE`.

VARID Variable ID.

START A vector of integers specifying the index in the variable where the first 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 length of `START` must be the same as the number of dimensions of the specified variable. 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 reading the data values.

COUNT A vector of integers specifying the edge lengths along each dimension of the block of data values to be read. To read a single value, for example, specify `COUNT` as (1, 1, ..., 1). The length of `COUNT` is the number of dimensions of the specified variable. 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.

text

i1vals

i2vals

ival

rvals

dvals

The block of data values to be read. The data should be of the type appropriate for the function called. 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 *The NetCDF Users Guide*).

Errors

`NF_GET_VARA_` *type* returns the value `NF_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 corner 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 edge lengths added to the specified corner would have referenced data out of range for the rank of the specified variable. For example, an edge length that is larger than the corresponding dimension length minus the corner index will cause an error.
- One or more of the 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.

Example

Here is an example using `NF_GET_VARA_DOUBLE` to read all the values 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`, and that there are ten `lon` values, five `lat` values, and three `time` values.

```

INCLUDE 'netcdf.inc'
...
PARAMETER (NDIMS=3)           ! number of dimensions
PARAMETER (TIMES=3, LATS=5, LONS=10) ! dimension lengths
INTEGER STATUS, NCID
INTEGER RHID                   ! variable ID
INTEGER START(NDIMS), COUNT(NDIMS)
DOUBLE RHVALS(LONS, LATS, TIMES)
DATA START /1, 1, 1/           ! start at first value
DATA COUNT /LONS, LATS, TIMES/ ! get all the values
...
STATUS = NF_OPEN ('foo.nc', NF_NOWRITE, NCID)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
...
STATUS = NF_INQ_VARID (NCID, 'rh', RHID)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
STATUS = NF_GET_VARA_DOUBLE (NCID, RHID, START, COUNT, RHVALS)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)

```

4.14 NF_GET_VARS_ *type*

The `NF_GET_VARS_ type` family of functions read a subsampled (strided) array section of values from a netCDF variable of an open netCDF dataset. The subsampled array section is specified by giving a corner, a vector of edge lengths, and a stride vector. The values are read with the first dimension of the netCDF variable varying fastest. The netCDF dataset must be in data mode.

Usage

```

INTEGER FUNCTION NF_GET_VARS_TEXT (INTEGER NCID, INTEGER VARID,
                                  INTEGER START(*), INTEGER COUNT(*),
                                  INTEGER STRIDE(*), CHARACTER*(*) text)

```

```

INTEGER FUNCTION NF_GET_VARS_INT1 (INTEGER NCID, INTEGER VARID,
                                   INTEGER START(*), INTEGER COUNT(*),
                                   INTEGER STRIDE(*), INTEGER*1 i1vals(*))
INTEGER FUNCTION NF_GET_VARS_INT2 (INTEGER NCID, INTEGER VARID,
                                   INTEGER START(*), INTEGER COUNT(*),
                                   INTEGER STRIDE(*), INTEGER*2 i2vals(*))
INTEGER FUNCTION NF_GET_VARS_INT  (INTEGER NCID, INTEGER VARID,
                                   INTEGER START(*), INTEGER COUNT(*),
                                   INTEGER STRIDE(*), INTEGER ival(*))
INTEGER FUNCTION NF_GET_VARS_REAL (INTEGER NCID, INTEGER VARID,
                                   INTEGER START(*), INTEGER COUNT(*),
                                   INTEGER STRIDE(*), REAL rvals(*))
INTEGER FUNCTION NF_GET_VARS_DOUBLE(INTEGER NCID, INTEGER VARID,
                                    INTEGER START(*), INTEGER COUNT(*),
                                    INTEGER STRIDE(*), DOUBLE dvals(*))

```

NCID NetCDF ID, from a previous call to `NF_OPEN` or `NF_CREATE`.

VARID Variable ID.

START A vector of integers specifying the index in the variable from which the first 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 reading the data values.

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.

STRIDE A vector of integers specifying, for each dimension, the interval between selected indices or the value 0. The elements of the vector correspond, in order, to the variable's dimensions. A value of 1 accesses adjacent values of the netCDF variable in the corresponding dimension; a value of 2 accesses every other value of the netCDF variable in the corresponding dimension; and so on. A 0 argument is treated as (1, 1, ..., 1).

text
i1vals
i2vals
ival
rvals
dvals

The block of data values to be read. The data should be of the type appropriate for the function called. 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 *The NetCDF Users Guide*).

Errors

`NF_GET_VARS_` *type* returns the value `NF_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 start, count and stride generate an index which is out of range.
- One or more of the 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.

Example

Here is an example using `NF_GET_VARS_DOUBLE` to read every other value in each dimension of the variable named `rh` from an existing netCDF dataset named `foo.nc`. Values are assigned, using the same dimensional strides, to a 2-parameter array. For simplicity in this example, we assume that we know that `rh` is dimensioned with `lon`, `lat`, and `time`, and that there are ten `lon` values, five `lat` values, and three `time` values.

```

INCLUDE 'netcdf.inc'
...
PARAMETER (NDIMS=3)           ! number of dimensions
PARAMETER (TIMES=3, LATS=5, LONS=10) ! dimension lengths
INTEGER STATUS, NCID
INTEGER RHID ! variable ID
INTEGER START(NDIMS), COUNT(NDIMS), STRIDE(NDIMS)
DOUBLE DATA(LONS, LATS, TIMES)
DATA START /1, 1, 1/           ! start at first value
DATA COUNT /LONS, LATS, TIMES/
DATA STRIDE /2, 2, 2/
...
STATUS = NF_OPEN ('foo.nc', NF_NOWRITE, NCID)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
...
STATUS = NF_INQ_VARID (NCID, 'rh', RHID)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
STATUS = NF_GET_VARS_DOUBLE(NCID,RHID,START,COUNT,STRIDE,DATA(1,1,1))
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)

```

4.15 `NF_GET_VARM_` *type*

The `NF_GET_VARM_` *type* family of functions reads a mapped array section of values from a netCDF variable of an open netCDF dataset. The mapped array section is specified by giving a corner, a vector of edge lengths, a stride vector, and an index mapping vector. The index mapping vector is 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. No assumptions are made about the ordering or length of the dimensions of the data array. The netCDF dataset must be in data mode.

Usage

```

INTEGER FUNCTION NF_GET_VARM_TEXT (INTEGER NCID, INTEGER VARID,
                                   INTEGER START(*), INTEGER COUNT(*),
                                   INTEGER STRIDE(*), INTEGER IMAP(*),
                                   CHARACTER*(*) text)

INTEGER FUNCTION NF_GET_VARM_INT1 (INTEGER NCID, INTEGER VARID,
                                   INTEGER START(*), INTEGER COUNT(*),
                                   INTEGER STRIDE(*), INTEGER IMAP(*),
                                   INTEGER*1 i1vals(*))

INTEGER FUNCTION NF_GET_VARM_INT2 (INTEGER NCID, INTEGER VARID,
                                   INTEGER START(*), INTEGER COUNT(*),
                                   INTEGER STRIDE(*), INTEGER IMAP(*),
                                   INTEGER*2 i2vals(*))

INTEGER FUNCTION NF_GET_VARM_INT (INTEGER NCID, INTEGER VARID,
                                   INTEGER START(*), INTEGER COUNT(*),
                                   INTEGER STRIDE(*), INTEGER IMAP(*),
                                   INTEGER ival(*))

INTEGER FUNCTION NF_GET_VARM_REAL (INTEGER NCID, INTEGER VARID,
                                   INTEGER START(*), INTEGER COUNT(*),
                                   INTEGER STRIDE(*), INTEGER IMAP(*),
                                   REAL rvals(*))

INTEGER FUNCTION NF_GET_VARM_DOUBLE(INTEGER NCID, INTEGER VARID,
                                   INTEGER START(*), INTEGER COUNT(*),
                                   INTEGER STRIDE(*), INTEGER IMAP(*),
                                   DOUBLE dvals(*))

```

- NCID** NetCDF ID, from a previous call to `NF_OPEN` or `NF_CREATE`.
- VARID** Variable ID.
- START** A vector of integers specifying the index in the variable from which the first 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 reading the data values.
- 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.
- STRIDE** A vector of integers specifying, for each dimension, the interval between selected indices or the value 0. The elements of the vector correspond, in order, to the variable's dimensions. A value of 1 accesses adjacent values of the netCDF variable in the corresponding dimension; a value of 2 accesses every other value of the netCDF variable in the corresponding dimension; and so on. A 0 argument is treated as (1, 1, ..., 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. `IMAP(1)` gives the distance between elements of the internal array corresponding to the most rapidly varying dimension of the netCDF variable. `IMAP(N)` (where `N` is the rank of the netCDF variable) gives the distance between elements of the internal array corresponding to the most slowly varying dimension of the netCDF variable. Intervening `IMAP` elements correspond to other dimensions of the netCDF variable in the obvious way. Distances between elements are specified in units of elements (the distance between internal elements that occupy adjacent memory locations is 1 and not the element's byte-length as in netCDF 2).

text

i1vals

i2vals

ivals

rvals

dvals The block of data values to be read. The data should be of the type appropriate for the function called. 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 *The NetCDF Users Guide*).

Errors

`NF_GET_VARS` type returns the value `NF_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 `START`, `COUNT`, and `STRIDE` generate an index which is out of range. Note that no error checking is possible on the `imap` vector.
- One or more of the 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.

Example

The following `IMAP` vector maps in the trivial way a 2x3x4 netCDF variable and an internal array of the same shape:

```

REAL A(2,3,4)      ! same shape as netCDF variable
INTEGER IMAP(3)
DATA IMAP /1, 2, 6/ ! netCDF dimension      inter-element distance
                   ! -----
                   ! most rapidly varying      1
                   ! intermediate              2 (=IMAP(1)*2)
                   ! most slowly varying      6 (=IMAP(2)*3)

```

Using the IMAP vector above with `NF_GET_VARM_REAL` obtains the same result as simply using `NF_GET_VAR_REAL`.

Here is an example of using `NF_GET_VARM_REAL` to transpose a netCDF variable named `rh` which is described by the FORTRAN declaration `REAL RH(4,6)` (note the size and order of the dimensions):

```

INCLUDE 'netcdf.inc'
...
PARAMETER (NDIM=2)    ! rank of netCDF variable
INTEGER NCID          ! netCDF dataset ID
INTEGER STATUS        ! return code
INTEGER RHID          ! variable ID
INTEGER START(NDIM)  ! netCDF variable start point
INTEGER COUNT(NDIM)  ! size of internal array
INTEGER STRIDE(NDIM) ! netCDF variable subsampling intervals
INTEGER IMAP(NDIM)   ! internal array inter-element distances
REAL    RH(6,4)       ! note transposition of netCDF variable dimensions
DATA START  /1, 1/    ! start at first netCDF variable element
DATA COUNT  /4, 6/    ! entire netCDF variable; order corresponds
                  ! to netCDF variable -- not internal array
DATA STRIDE /1, 1/    ! sample every netCDF element
DATA IMAP   /6, 1/    ! would be /1, 4/ if not transposing
...
STATUS = NF_OPEN('foo.nc', NF_NOWRITE, NCID)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
...
STATUS = NF_INQ_VARID(NCID, 'rh', RHID)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
...
STATUS = NF_GET_VARM_REAL(NCID, RHID, START, COUNT, STRIDE, IMAP, RH)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)

```

Here is another example of using `NF_GET_VARM_REAL` to simultaneously transpose and subsample the same netCDF variable, by accessing every other point of the netCDF variable:

```

INCLUDE 'netcdf.inc'
...
PARAMETER (NDIM=2)    ! rank of netCDF variable
INTEGER NCID          ! netCDF dataset ID
INTEGER STATUS        ! return code
INTEGER RHID          ! variable ID
INTEGER START(NDIM)  ! netCDF variable start point
INTEGER COUNT(NDIM)  ! size of internal array
INTEGER STRIDE(NDIM) ! netCDF variable subsampling intervals
INTEGER IMAP(NDIM)   ! internal array inter-element distances
REAL    RH(3,2)       ! note transposition of (subsampled) dimensions
DATA START  /1, 1/    ! start at first netCDF variable value
DATA COUNT  /2, 3/    ! order of (subsampled) dimensions corresponds

```

```

                                ! to netCDF variable -- not internal array
DATA STRIDE /2, 2/      ! sample every other netCDF element
DATA IMAP   /3, 1/     ! would be '1, 2' if not transposing
...
STATUS = NF_OPEN('foo.nc', NF_NOWRITE, NCID)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
...
STATUS = NF_INQ_VARID(NCID, 'rh', RHID)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
...
STATUS = NF_GET_VARM_REAL(NCID, RHID, START, COUNT, STRIDE, IMAP, RH)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)

```

4.16 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). 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. 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, 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 programs.

The FORTRAN interface for reading and writing strings requires the use of different functions for accessing string values and numeric values, because standard FORTRAN does not permit the same formal parameter to be used for both character values and numeric values. An additional argument, specifying the declared length of the character string passed as a value, is required for `NF_PUT_VARA_TEXT` and `NF_GET_VARA_TEXT`. The actual length of the string is specified as the value of the edge-length vector corresponding to the character-position dimension.

Here is an example that defines a record variable, `tx`, for character strings and stores a character-string value into the third record using `NF_PUT_VARA_TEXT`. In this example, we assume the string variable and data are to be added to an existing netCDF dataset named `foo.nc` that already has an unlimited record dimension `time`.

```

INCLUDE 'netcdf.inc'
...
INTEGER  TDIMS, TXLEN
PARAMETER (TDIMS=2)    ! number of TX dimensions
PARAMETER (TXLEN = 15) ! length of example string
INTEGER  NCID
INTEGER  CHID          ! char position dimension id
INTEGER  TIMEID        ! record dimension id
INTEGER  TXID          ! variable ID
INTEGER  TXDIMS(TDIMS) ! variable shape
INTEGER  TSTART(TDIMS), TCOUNT(TDIMS)
CHARACTER*40 TXVAL     ! max length 40
DATA TXVAL /'example string'/
...
TXVAL(TXLEN:TXLEN) = CHAR(0)    ! null terminate
...
STATUS = NF_OPEN('foo.nc', NF_WRITE, NCID)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
STATUS = NF_REDEF(NCID)        ! enter define mode
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
...
! define character-position dimension for strings of max length 40
STATUS = NF_DEF_DIM(NCID, "chid", 40, CHID)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
...
! define a character-string variable
TXDIMS(1) = CHID    ! character-position dimension first
TXDIMS(2) = TIMEID
STATUS = NF_DEF_VAR(NCID, "tx", NF_CHAR, TDIMS, TXDIMS, TXID)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
...
STATUS = NF_ENDDEF(NCID) ! leave define mode
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
...
! write txval into tx netCDF variable in record 3

```


NCID NetCDF ID, from a previous call to `NF_OPEN` or `NF_CREATE`.
 VARID Variable ID.
 NAME New name for the specified variable.

Errors

`NF_RENAME_VAR` returns the value `NF_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 `NF_RENAME_VAR` to rename the variable `rh` to `rel_hum` in an existing netCDF dataset named `foo.nc`:

```
INCLUDE 'netcdf.inc'
...
INTEGER STATUS, NCID
INTEGER RHID           ! variable ID
...
STATUS = NF_OPEN ('foo.nc', NF_WRITE, NCID)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
...
STATUS = NF_REDEF (NCID) ! enter definition mode
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
STATUS = NF_INQ_VARID (NCID, 'rh', RHID)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
STATUS = NF_RENAME_VAR (NCID, RHID, 'rel_hum')
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
STATUS = NF_ENDDEF (NCID) ! leave definition mode
IF (STATUS .NE. NF_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 `NF_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 `NF_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 "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 `NF_SET_FILL`. See [Section 2.16 \[NF_SET_FILL\]](#), page 22. 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 4.17 \[Fill Values\]](#), page 60.

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.


```

                                LEN, INTEGER*2 I2VALS(*))
INTEGER FUNCTION  NF_PUT_ATT_INT  (INTEGER NCID, INTEGER VARID,
                                CHARACTER*(*) NAME, INTEGER XTYPE,
                                LEN, INTEGER IVALS(*))
INTEGER FUNCTION  NF_PUT_ATT_REAL (INTEGER NCID, INTEGER VARID,
                                CHARACTER*(*) NAME, INTEGER XTYPE,
                                LEN, REAL RVALS(*))
INTEGER FUNCTION  NF_PUT_ATT_DOUBLE(INTEGER NCID, INTEGER VARID,
                                CHARACTER*(*) NAME, INTEGER XTYPE,
                                LEN, DOUBLE DVALS(*))

```

NCID NetCDF ID, from a previous call to `NF_OPEN` or `NF_CREATE`.

VARID Variable ID of the variable to which the attribute will be assigned or `NF_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 see [Section 5.2 \[Attribute Conventions\]](#), page 63.

XTYPE One of the set of predefined netCDF external data types. The type of this parameter, `NF_TYPE`, is defined in the netCDF header file. The valid netCDF external data types are `NF_BYTE`, `NF_CHAR`, `NF_SHORT`, `NF_INT`, `NF_FLOAT`, and `NF_DOUBLE`. Although it's possible to create attributes of all types, `NF_CHAR` and `NF_DOUBLE` attributes are adequate for most purposes.

LEN Number of values provided for the attribute.

TEXT

I1VALS

I2VALS

IVALS

RVALS

DVALS An array of `LEN` attribute values. The data should be of a type appropriate for the function called. You cannot write `CHARACTER` data into a numeric attribute or numeric data into a text attribute. For numeric data, if the type of data differs from the attribute type, type conversion will occur. See [section "Type Conversion" in *The NetCDF Users Guide*](#).

Errors

`NF_PUT_ATT_` *type* returns the value `NF_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 `NF_MAX_ATTRS`.

Example

Here is an example using `NF_PUT_ATT_DOUBLE` 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`:

```

INCLUDE 'netcdf.inc'
...
INTEGER STATUS, NCID
INTEGER RHID           ! variable ID
DOUBLE RHRNGE(2)
DATA RHRNGE /0.0D0, 100.0D0/
...
STATUS = NF_OPEN ('foo.nc', NF_WRITE, NCID)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
...
STATUS = NF_REDEF (NCID)      ! enter define mode
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
STATUS = NF_INQ_VARID (NCID, 'rh', RHID)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
...
STATUS = NF_PUT_ATT_DOUBLE (NCID, RHID, 'valid_range', NF_DOUBLE, &
                           2, RHRNGE)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
STATUS = NF_PUT_ATT_TEXT (NCID, NF_GLOBAL, 'title', 19,
                          'example netCDF dataset')
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
...
STATUS = NF_ENDDEF (NCID)    ! leave define mode
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)

```

5.4 NF_INQ_ATT Family

This family of functions returns information about a netCDF attribute. All but one of these functions require the variable ID and attribute name; the exception is `NF_INQ_ATTNAME`. Information about an attribute includes its type, length, name, and number. See the `NF_GET_ATT` family for getting attribute values.

The function `NF_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.

The function `NF_INQ_ATT` returns the attribute's type and length. The other functions each return just one item of information about an attribute.

Usage

```

INTEGER FUNCTION NF_INQ_ATT      (INTEGER NCID, INTEGER VARID,
                                CHARACTER*(*) NAME, INTEGER xtype,
                                INTEGER len)

INTEGER FUNCTION NF_INQ_ATTTYPE (INTEGER NCID, INTEGER VARID,
                                CHARACTER*(*) NAME, INTEGER xtype)

INTEGER FUNCTION NF_INQ_ATTLEN  (INTEGER NCID, INTEGER VARID,
                                CHARACTER*(*) NAME, INTEGER len)

INTEGER FUNCTION NF_INQ_ATTNAME (INTEGER NCID, INTEGER VARID,
                                INTEGER ATTNUM, CHARACTER*(*) name)

INTEGER FUNCTION NF_INQ_ATTID   (INTEGER NCID, INTEGER VARID,
                                CHARACTER*(*) NAME, INTEGER attnum)

```

NCID NetCDF ID, from a previous call to `NF_OPEN` or `NF_CREATE`.

VARID Variable ID of the attribute's variable, or `NF_GLOBAL` for a global attribute.

NAME Attribute name. For `NF_INQ_ATTNAME`, this is a pointer to the location for the returned attribute name.

xtype Returned attribute type, one of the set of predefined netCDF external data types. The valid netCDF external data types are `NF_BYTE`, `NF_CHAR`, `NF_SHORT`, `NF_INT`, `NF_FLOAT`, and `NF_DOUBLE`.

len 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 `NF_INQ_ATTNAME`, the input attribute number; for `NF_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 `NF_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 `NF_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 `NF_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 `NF_INQ_ATT` to find out the type and length of a variable attribute named `valid_range` for a netCDF variable named `rh` and a global attribute named `title` in an existing netCDF dataset named `foo.nc`:

```

INCLUDE 'netcdf.inc'
...
INTEGER STATUS, NCID
INTEGER RHID           ! variable ID
INTEGER VRLEN, TLEN    ! attribute lengths
...
STATUS = NF_OPEN ('foo.nc', NF_NOWRITE, NCID)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
...
STATUS = NF_INQ_VARID (NCID, 'rh', RHID)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
...
STATUS = NF_INQ_ATTLEN (NCID, RHID, 'valid_range', VRLEN)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
STATUS = NF_INQ_ATTLEN (NCID, NF_GLOBAL, 'title', TLEN)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)

```

5.5 `NF_GET_ATT_type`

Members of the `NF_GET_ATT_type` family of functions get the value(s) of a netCDF attribute, given its variable ID and name.

Usage

```

INTEGER FUNCTION NF_GET_ATT_TEXT (INTEGER NCID, INTEGER VARID,
CHARACTER*(*) NAME,
CHARACTER*(*) text)
INTEGER FUNCTION NF_GET_ATT_INT1 (INTEGER NCID, INTEGER VARID,
CHARACTER*(*) NAME,
INTEGER*1 i1vals(*))
INTEGER FUNCTION NF_GET_ATT_INT2 (INTEGER NCID, INTEGER VARID,
CHARACTER*(*) NAME,
INTEGER*2 i2vals(*))
INTEGER FUNCTION NF_GET_ATT_INT (INTEGER NCID, INTEGER VARID,
CHARACTER*(*) NAME,
INTEGER ival(*))
INTEGER FUNCTION NF_GET_ATT_REAL (INTEGER NCID, INTEGER VARID,
CHARACTER*(*) NAME,
REAL rvals(*))
INTEGER FUNCTION NF_GET_ATT_DOUBLE (INTEGER NCID, INTEGER VARID,
CHARACTER*(*) NAME,
DOUBLE dvals(*))

```

`NCID` NetCDF ID, from a previous call to `NF_OPEN` or `NF_CREATE`.

VARID	Variable ID of the attribute's variable, or <code>NF_GLOBAL</code> for a global attribute.
NAME	Attribute name.
TEXT	
I1VALS	
I2VALS	
IVALS	
RVALS	
DVALS	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 <code>NF_INQ_ATTLEN</code> first to find out the length of the attribute. 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 <i>The NetCDF Users Guide</i> .

Errors

`NF_GET_ATT_` *type* returns the value `NF_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 `NF_GET_ATT_DOUBLE` to determine the values of a variable 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, but that we do know the types of the attributes. Hence, to allocate enough space to store them, we must first inquire about the length of the attributes.

```

INCLUDE 'netcdf.inc'
...
PARAMETER (MVRLEN=3)           ! max number of "valid_range" values
PARAMETER (MTLEN=80)          ! max length of "title" attribute
INTEGER STATUS, NCID
INTEGER RHID                   ! variable ID
INTEGER VRLEN, TLEN           ! attribute lengths
DOUBLE PRECISION VRVAL(MVRLEN) ! vr attribute values
CHARACTER*80 TITLE            ! title attribute values
...
STATUS = NF_OPEN ('foo.nc', NF_WRITE, NCID)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
...

```

```

STATUS = NF_INQ_VARID (NCID, 'rh', RHID)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
...
! find out attribute lengths, to make sure we have enough space
STATUS = NF_INQ_ATTLEN (NCID, RHID, 'valid_range', VRLEN)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
STATUS = NF_INQ_ATTLEN (NCID, NF_GLOBAL, 'title', TLEN)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
! get attribute values, if not too big
IF (VRLEN .GT. MVRLEN) THEN
    WRITE (*,*) 'valid_range attribute too big!'
    CALL EXIT
ELSE
    STATUS = NF_GET_ATT_DOUBLE (NCID, RHID, 'valid_range', VRVAL)
    IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
ENDIF
IF (TLEN .GT. MTLEN) THEN
    WRITE (*,*) 'title attribute too big!'
    CALL EXIT
ELSE
    STATUS = NF_GET_ATT_TEXT (NCID, NF_GLOBAL, 'title', TITLE)
    IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
ENDIF

```

5.6 NF_COPY_ATT

The function `NF_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.

Usage

```

INTEGER FUNCTION NF_COPY_ATT (INTEGER NCID_IN, INTEGER VARID_IN,
                             CHARACTER*(*) NAME, INTEGER NCID_OUT,
                             INTEGER VARID_OUT)

```

- NCID_IN** The netCDF ID of an input netCDF dataset from which the attribute will be copied, from a previous call to `NF_OPEN` or `NF_CREATE`.
- VARID_IN** ID of the variable in the input netCDF dataset from which the attribute will be copied, or `NF_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 `NF_OPEN` or `NF_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 `NF_GLOBAL` to copy to a global attribute.

Errors

`NF_COPY_ATT` returns the value `NF_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 `NF_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:

```

INCLUDE 'netcdf.inc'
...
INTEGER STATUS           ! error status
INTEGER NCID1, NCID2     ! netCDF IDs
INTEGER RHID, AVRHID     ! variable IDs
...
STATUS = NF_OPEN ('foo.nc', NF_NOWRITE, NCID1)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
STATUS = NF_OPEN ('bar.nc', NF_WRITE, NCID2)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
...
STATUS = NF_INQ_VARID (NCID1, 'rh', RHID)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
STATUS = NF_INQ_VARID (NCID2, 'avgrh', AVRHID)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
...
STATUS = NF_REDEF (NCID2) ! enter define mode
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
! copy variable attribute from "rh" to "avgrh"
STATUS = NF_COPY_ATT (NCID1, RHID, 'units', NCID2, AVRHID)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
...
STATUS = NF_ENDDEF (NCID2) ! leave define mode
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)

```

5.7 NF_RENAME_ATT

The function `NF_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

```

      INTEGER FUNCTION NF_RENAME_ATT (INTEGER NCID, INTEGER VARID,
                                     CHARACTER*(*) NAME,
                                     CHARACTER*(*) NEWNAME)

```

NCID NetCDF ID, from a previous call to `NF_OPEN` or `NF_CREATE`
VARID ID of the attribute's variable, or `NF_GLOBAL` for a global attribute
NAME 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

`NF_RENAME_ATT` returns the value `NF_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 `NF_RENAME_ATT` to rename the variable attribute units to Units for a variable `rh` in an existing netCDF dataset named `foo.nc`:

```

      INCLUDE "netcdf.inc"
      ...
      INTEGER STATUS    ! error status
      INTEGER NCID     ! netCDF ID
      INTEGER RHID     ! variable ID
      ...
      STATUS = NF_OPEN ("foo.nc", NF_NOWRITE, NCID)
      IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
      ...
      STATUS = NF_INQ_VARID (NCID, "rh", RHID)
      IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
      ...
      ! rename attribute
      STATUS = NF_RENAME_ATT (NCID, RHID, "units", "Units")
      IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)

```

5.8 NF_DEL_ATT

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

Usage

INTEGER FUNCTION `NF_DEL_ATT` (INTEGER `NCID`, INTEGER `VARID`, CHARACTER*(*) `NAME`)

`NCID` NetCDF ID, from a previous call to `NF_OPEN` or `NF_CREATE`.

`VARID` ID of the attribute's variable, or `NF_GLOBAL` for a global attribute.

`NAME` The name of the attribute to be deleted.

Errors

`NF_DEL_ATT` returns the value `NF_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 `NF_DEL_ATT` to delete the variable attribute `Units` for a variable `rh` in an existing netCDF dataset named `foo.nc`:

```

INCLUDE 'netcdf.inc'
...
INTEGER STATUS           ! error status
INTEGER NCID             ! netCDF ID
INTEGER RHID             ! variable ID
...
STATUS = NF_OPEN ('foo.nc', NF_WRITE, NCID)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
...
STATUS = NF_INQ_VARID (NCID, 'rh', RHID)
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
...
! delete attribute
STATUS = NF_REDEF (NCID) ! enter define mode
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
STATUS = NF_DEL_ATT (NCID, RHID, 'Units')
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)
STATUS = NF_ENDDEF (NCID) ! leave define mode
IF (STATUS .NE. NF_NOERR) CALL HANDLE_ERR(STATUS)

```


Appendix A Summary of FORTRAN 77 Interface

Input parameters are in upper case, output parameters are in lower case. The FORTRAN types of all the parameters are listed alphabetically by parameter name below the function declarations.

```

CHARACTER*80 FUNCTION  NF_INQ_LIBVERS()
CHARACTER*80 FUNCTION  NF_STRERROR  (NCERR)
INTEGER FUNCTION  NF_CREATE          (PATH, CMODE, ncid)
INTEGER FUNCTION  NF_OPEN            (PATH, MODE, ncid)
INTEGER FUNCTION  NF_SET_FILL        (NCID, FILLMODE, old_mode)
INTEGER FUNCTION  NF_REDEF            (NCID)
INTEGER FUNCTION  NF_ENDDEF          (NCID)
INTEGER FUNCTION  NF_SYNC            (NCID)
INTEGER FUNCTION  NF_ABORT           (NCID)
INTEGER FUNCTION  NF_CLOSE           (NCID)
INTEGER FUNCTION  NF_INQ              (NCID, ndims, nvars, ngatts,
                                       unlimdimid)
INTEGER FUNCTION  NF_INQ_NDIMS        (NCID, ndims)
INTEGER FUNCTION  NF_INQ_NVARS        (NCID, nvars)
INTEGER FUNCTION  NF_INQ_NATTS        (NCID, ngatts)
INTEGER FUNCTION  NF_INQ_UNLIMDIM     (NCID, unlimdimid)
INTEGER FUNCTION  NF_DEF_DIM          (NCID, NAME, LEN, dimid)
INTEGER FUNCTION  NF_INQ_DIMID        (NCID, NAME, dimid)
INTEGER FUNCTION  NF_INQ_DIM          (NCID, DIMID, name, len)
INTEGER FUNCTION  NF_INQ_DIMNAME      (NCID, DIMID, name)
INTEGER FUNCTION  NF_INQ_DIMLEN       (NCID, DIMID, len)
INTEGER FUNCTION  NF_RENAME_DIM       (NCID, DIMID, NAME)

INTEGER FUNCTION  NF_DEF_VAR          (NCID, NAME, XTYPE, NDIMS, DIMIDS,
                                       varid)
INTEGER FUNCTION  NF_INQ_VAR          (NCID, VARID, name, xtype, ndims,
                                       dimids, natts)
INTEGER FUNCTION  NF_INQ_VARID        (NCID, NAME, varid)
INTEGER FUNCTION  NF_INQ_VARNAME      (NCID, VARID, name)
INTEGER FUNCTION  NF_INQ_VARTYPE      (NCID, VARID, xtype)
INTEGER FUNCTION  NF_INQ_VARNDIMS     (NCID, VARID, ndims)
INTEGER FUNCTION  NF_INQ_VARDIMID     (NCID, VARID, DIMIDS)
INTEGER FUNCTION  NF_INQ_VARNATTS     (NCID, VARID, natts)
INTEGER FUNCTION  NF_RENAME_VAR       (NCID, VARID, NAME)
INTEGER FUNCTION  NF_PUT_VAR_TEXT     (NCID, VARID, TEXT)
INTEGER FUNCTION  NF_GET_VAR_TEXT     (NCID, VARID, text)
INTEGER FUNCTION  NF_PUT_VAR_INT1     (NCID, VARID, I1VAL)
INTEGER FUNCTION  NF_GET_VAR_INT1     (NCID, VARID, i1val)
INTEGER FUNCTION  NF_PUT_VAR_INT2     (NCID, VARID, I2VAL)
INTEGER FUNCTION  NF_GET_VAR_INT2     (NCID, VARID, i2val)
INTEGER FUNCTION  NF_PUT_VAR_INT      (NCID, VARID, IVAL)
INTEGER FUNCTION  NF_GET_VAR_INT      (NCID, VARID, ival)

```

```

INTEGER FUNCTION  NF_PUT_VAR_REAL    (NCID, VARID, RVAL)
INTEGER FUNCTION  NF_GET_VAR_REAL    (NCID, VARID, rval)
INTEGER FUNCTION  NF_PUT_VAR_DOUBLE (NCID, VARID, DVAL)
INTEGER FUNCTION  NF_GET_VAR_DOUBLE (NCID, VARID, dval)
INTEGER FUNCTION  NF_PUT_VAR1_TEXT  (NCID, VARID, INDEX, TEXT)
INTEGER FUNCTION  NF_GET_VAR1_TEXT  (NCID, VARID, INDEX, text)
INTEGER FUNCTION  NF_PUT_VAR1_INT1   (NCID, VARID, INDEX, I1VAL)
INTEGER FUNCTION  NF_GET_VAR1_INT1   (NCID, VARID, INDEX, i1val)
INTEGER FUNCTION  NF_PUT_VAR1_INT2   (NCID, VARID, INDEX, I2VAL)
INTEGER FUNCTION  NF_GET_VAR1_INT2   (NCID, VARID, INDEX, i2val)
INTEGER FUNCTION  NF_PUT_VAR1_INT    (NCID, VARID, INDEX, IVAL)
INTEGER FUNCTION  NF_GET_VAR1_INT    (NCID, VARID, INDEX, ival)
INTEGER FUNCTION  NF_PUT_VAR1_REAL   (NCID, VARID, INDEX, RVAL)
INTEGER FUNCTION  NF_GET_VAR1_REAL   (NCID, VARID, INDEX, rval)
INTEGER FUNCTION  NF_PUT_VAR1_DOUBLE(NCID, VARID, INDEX, DVAL)
INTEGER FUNCTION  NF_GET_VAR1_DOUBLE(NCID, VARID, INDEX, dval)
INTEGER FUNCTION  NF_PUT_VARA_TEXT   (NCID, VARID, START, COUNT, TEXT)
INTEGER FUNCTION  NF_GET_VARA_TEXT   (NCID, VARID, START, COUNT, text)
INTEGER FUNCTION  NF_PUT_VARA_INT1   (NCID, VARID, START, COUNT, I1VALS)
INTEGER FUNCTION  NF_GET_VARA_INT1   (NCID, VARID, START, COUNT, i1vals)
INTEGER FUNCTION  NF_PUT_VARA_INT2   (NCID, VARID, START, COUNT, I2VALS)
INTEGER FUNCTION  NF_GET_VARA_INT2   (NCID, VARID, START, COUNT, i2vals)
INTEGER FUNCTION  NF_PUT_VARA_INT    (NCID, VARID, START, COUNT, IVALS)
INTEGER FUNCTION  NF_GET_VARA_INT    (NCID, VARID, START, COUNT, ivals)
INTEGER FUNCTION  NF_PUT_VARA_REAL   (NCID, VARID, START, COUNT, RVALS)
INTEGER FUNCTION  NF_GET_VARA_REAL   (NCID, VARID, START, COUNT, rvals)
INTEGER FUNCTION  NF_PUT_VARA_DOUBLE(NCID, VARID, START, COUNT, DVALS)
INTEGER FUNCTION  NF_GET_VARA_DOUBLE(NCID, VARID, START, COUNT, dvals)
INTEGER FUNCTION  NF_PUT_VARS_TEXT   (NCID, VARID, START, COUNT, STRIDE,
                                     TEXT)
INTEGER FUNCTION  NF_GET_VARS_TEXT   (NCID, VARID, START, COUNT, STRIDE,
                                     text)
INTEGER FUNCTION  NF_PUT_VARS_INT1   (NCID, VARID, START, COUNT, STRIDE,
                                     I1VALS)
INTEGER FUNCTION  NF_GET_VARS_INT1   (NCID, VARID, START, COUNT, STRIDE,
                                     i1vals)
INTEGER FUNCTION  NF_PUT_VARS_INT2   (NCID, VARID, START, COUNT, STRIDE,
                                     I2VALS)
INTEGER FUNCTION  NF_GET_VARS_INT2   (NCID, VARID, START, COUNT, STRIDE,
                                     i2vals)
INTEGER FUNCTION  NF_PUT_VARS_INT    (NCID, VARID, START, COUNT, STRIDE,
                                     IVALS)
INTEGER FUNCTION  NF_GET_VARS_INT    (NCID, VARID, START, COUNT, STRIDE,
                                     ivals)
INTEGER FUNCTION  NF_PUT_VARS_REAL   (NCID, VARID, START, COUNT, STRIDE,
                                     RVALS)
INTEGER FUNCTION  NF_GET_VARS_REAL   (NCID, VARID, START, COUNT, STRIDE,

```

```

                                rvals)
INTEGER FUNCTION  NF_PUT_VARS_DOUBLE(NCID, VARID, START, COUNT, STRIDE,
                                DVALS)
INTEGER FUNCTION  NF_GET_VARS_DOUBLE(NCID, VARID, START, COUNT, STRIDE,
                                dvals)
INTEGER FUNCTION  NF_PUT_VARM_TEXT  (NCID, VARID, START, COUNT, STRIDE,
                                IMAP, TEXT)
INTEGER FUNCTION  NF_GET_VARM_TEXT  (NCID, VARID, START, COUNT, STRIDE,
                                IMAP, text)
INTEGER FUNCTION  NF_PUT_VARM_INT1  (NCID, VARID, START, COUNT, STRIDE,
                                IMAP, I1VALS)
INTEGER FUNCTION  NF_GET_VARM_INT1  (NCID, VARID, START, COUNT, STRIDE,
                                IMAP, i1vals)
INTEGER FUNCTION  NF_PUT_VARM_INT2  (NCID, VARID, START, COUNT, STRIDE,
                                IMAP, I2VALS)
INTEGER FUNCTION  NF_GET_VARM_INT2  (NCID, VARID, START, COUNT, STRIDE,
                                IMAP, i2vals)
INTEGER FUNCTION  NF_PUT_VARM_INT   (NCID, VARID, START, COUNT, STRIDE,
                                IMAP, IVALS)
INTEGER FUNCTION  NF_GET_VARM_INT   (NCID, VARID, START, COUNT, STRIDE,
                                IMAP, ivalS)
INTEGER FUNCTION  NF_PUT_VARM_REAL   (NCID, VARID, START, COUNT, STRIDE,
                                IMAP, RVALS)
INTEGER FUNCTION  NF_GET_VARM_REAL   (NCID, VARID, START, COUNT, STRIDE,
                                IMAP, rvals)
INTEGER FUNCTION  NF_PUT_VARM_DOUBLE(NCID, VARID, START, COUNT, STRIDE,
                                IMAP, DVALS)
INTEGER FUNCTION  NF_GET_VARM_DOUBLE(NCID, VARID, START, COUNT, STRIDE,
                                IMAP, dvals)

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)
INTEGER FUNCTION  NF_COPY_ATT        (NCID_IN, VARID_IN, NAME,
                                NCID_OUT, VARID_OUT)

INTEGER FUNCTION  NF_RENAME_ATT      (NCID, VARID, CURNAME, NEWNAME)
INTEGER FUNCTION  NF_DEL_ATT         (NCID, VARID, NAME)
INTEGER FUNCTION  NF_PUT_ATT_TEXT    (NCID, VARID, NAME, LEN, TEXT)
INTEGER FUNCTION  NF_GET_ATT_TEXT    (NCID, VARID, NAME, text)
INTEGER FUNCTION  NF_PUT_ATT_INT1    (NCID, VARID, NAME, XTYPE, LEN,
                                I1VALS)
INTEGER FUNCTION  NF_GET_ATT_INT1    (NCID, VARID, NAME, i1vals)
INTEGER FUNCTION  NF_PUT_ATT_INT2    (NCID, VARID, NAME, XTYPE, LEN,
                                I2VALS)
INTEGER FUNCTION  NF_GET_ATT_INT2    (NCID, VARID, NAME, i2vals)

```

```

INTEGER FUNCTION  NF_PUT_ATT_INT      (NCID, VARID, NAME, XTYPE, LEN,
                                       IVALS)
INTEGER FUNCTION  NF_GET_ATT_INT      (NCID, VARID, NAME, ivals)
INTEGER FUNCTION  NF_PUT_ATT_REAL     (NCID, VARID, NAME, XTYPE, LEN,
                                       RVALS)
INTEGER FUNCTION  NF_GET_ATT_REAL     (NCID, VARID, NAME, rvals)
INTEGER FUNCTION  NF_PUT_ATT_DOUBLE   (NCID, VARID, NAME, XTYPE, LEN,
                                       DVALS)
INTEGER FUNCTION  NF_GET_ATT_DOUBLE   (NCID, VARID, NAME, dvals)

INTEGER          ATTNUM                ! attribute number
INTEGER          attnum                 ! returned attribute number
INTEGER          CMODE                  ! NF_NOCLOBBER, NF_SHARE flags expression
INTEGER          COUNT                  ! array of edge lengths of block of values
CHARACTER(*)     CURNAME                ! current name (before renaming)
INTEGER          DIMID                  ! dimension ID
INTEGER          dimid                  ! returned dimension ID
INTEGER          DIMIDS                 ! list of dimension IDs
INTEGER          dimids                 ! list of returned dimension IDs
DOUBLEPRECISION DVAL                   ! single data value
DOUBLEPRECISION dval                   ! returned single data value
DOUBLEPRECISION DVALS                  ! array of data values
DOUBLEPRECISION dvals                  ! array of returned data values
INTEGER          FILLMODE               ! NF_NOFILL or NF_FILL, for setting fill mode
INTEGER*1        I1VAL                  ! single data value
INTEGER*1        i1val                  ! returned single data value
INTEGER*1        I1VALS                 ! array of data values
INTEGER*1        i1vals                 ! array of returned data values
INTEGER*2        I2VAL                  ! single data value
INTEGER*2        i2val                  ! returned single data value
INTEGER*2        I2VALS                 ! array of data values
INTEGER*2        i2vals                 ! array of returned data values
INTEGER          IMAP                   ! index mapping vector
INTEGER          INDEX                  ! variable array index vector
INTEGER          IVAL                   ! single data value
INTEGER          ival                   ! returned single data value
INTEGER          IVALS                  ! array of data values
INTEGER          ivals                  ! array of returned data values
INTEGER          LEN                    ! dimension or attribute length
INTEGER          len                    ! returned dimension or attribute length
INTEGER          MODE                   ! open mode, one of NF_WRITE or NF_NOWRITE
CHARACTER(*)     NAME                   ! dimension, variable, or attribute name
CHARACTER(*)     name                   ! returned dim, var, or att name
INTEGER          natts                  ! returned number of attributes
INTEGER          NCERR                  ! error returned from NF_xxx function call
INTEGER          NCID                   ! netCDF ID of an open netCDF dataset
INTEGER          ncid                   ! returned netCDF ID

```

INTEGER	NCID_IN	! netCDF ID of open source netCDF dataset
INTEGER	NCID_OUT	! netCDF ID of open destination netCDF dataset
INTEGER	NDIMS	! number of dimensions
INTEGER	ndims	! returned number of dimensions
CHARACTER(*)	NEWNAME	! new name for dim, var, or att
INTEGER	ngatts	! returned number of global attributes
INTEGER	nvars	! returned number of variables
INTEGER	old_mode	! previous fill mode, NF_NOFILL or NF_FILL,
CHARACTER(*)	PATH	! name of netCDF dataset
REAL	RVAL	! single data value
REAL	rval	! returned single data value
REAL	RVALS	! array of data values
REAL	rvals	! array of returned data values
INTEGER	START	! variable array indices of first value
INTEGER	STRIDE	! variable array dimensional strides
CHARACTER(*)	TEXT	! input text value
CHARACTER(*)	text	! returned text value
INTEGER	unlimdimid	! returned ID of unlimited dimension
INTEGER	VARID	! variable ID
INTEGER	varid	! returned variable ID
INTEGER	VARID_IN	! variable ID
INTEGER	VARID_OUT	! variable ID
INTEGER	XTYPE	! external type: NF_BYTE, NF_CHAR, ... ,
INTEGER	xtype	! returned external type

Appendix B Overview of C interface changes from NetCDF 2 to NetCDF 3

NetCDF version 3 includes a complete rewrite of the netCDF library. It is about twice as fast as the previous version. The netCDF file format is unchanged, so files written with version 3 can be read with version 2 code and vice versa.

The core library is now written in ANSI C. For example, prototypes are used throughout as well as const qualifiers where appropriate. You must have an ANSI C compiler to compile this version.

Rewriting the library offered an opportunity to implement improved C and FORTRAN interfaces that provide some significant benefits:

- type safety, by eliminating the need to use generic void* pointers;
- automatic type conversions, by eliminating the undesirable coupling between the language-independent external netCDF types (NF_BYTE, ..., NF_DOUBLE) and language-dependent internal data types (char, ..., double);
- support for future enhancements, by eliminating obstacles to the clean addition of support for packed data and multithreading;
- more standard error behavior, by uniformly communicating an error status back to the calling program in the return value of each function.

It is not necessary to rewrite programs that use the version 2 C interface, because the netCDF-3 library includes a backward compatibility interface that supports all the old functions, globals, and behavior. We are hoping that the benefits of the new interface will be an incentive to use it in new netCDF applications. It is possible to convert old applications to the new interface incrementally, replacing netCDF-2 calls with the corresponding netCDF-3 calls one at a time. If you want to check that only netCDF-3 calls are used in an application, a preprocessor macro (NO_NETCDF_2) is available for that purpose.

Other changes in the implementation of netCDF result in improved portability, maintainability, and performance on most platforms. A clean separation between I/O and type layers facilitates platform-specific optimizations. The new library no longer uses a vendor-provided XDR library, which simplifies linking programs that use netCDF and speeds up data access significantly in most cases.

B.1 The NetCDF-3 C Interface

First, here's an example of C code that uses the netCDF-2 interface:

```
void *bufferp;
NF_TYPE xtype;
ncvarinq(ncid, varid, ..., &xtype, ...
...
/* allocate bufferp based on dimensions and type */
...
if (ncvarget(ncid, varid, start, count, bufferp) == -1) {
    fprintf(stderr, "Can't get data, error code = %d\n", ncerr);
    /* deal with it */
    ...
}
```

```

switch(xtype) {
    /* deal with the data, according to type */
    ...
case NF_FLOAT:
    fanalyze((float *)bufferp);
    break;
case NF_DOUBLE:
    danalyze((double *)bufferp);
    break;
}

```

Here's how you might handle this with the new netCDF-3 C interface:

```

/*
 * I want to use doubles for my analysis.
 */
double dbuf[NDOUBLES];
int status;

/* So, I use the function that gets the data as doubles. */
status = NF_GET_VARA_DOUBLE(NCID, varid, start, count, dbuf)
if (status != NF_NOERR) {
    fprintf(stderr, "Can't get data: %s\n", NF_STRERROR(STATUS));
    /* deal with it */
    ...
}
danalyze(dbuf);

```

The example above illustrates changes in function names, data type conversion, and error handling, discussed in detail in the sections below.

B.2 Function Naming Conventions

The netCDF-3 C library employs a new naming convention, intended to make netCDF programs more readable. For example, the name of the function to rename a variable is now `NF_RENAME_VAR` instead of the previous `ncvarrename`.

All netCDF-3 C function names begin with the `NF_` prefix. The second part of the name is a verb, like `get`, `put`, `inq` (for inquire), or `open`. The third part of the name is typically the object of the verb: for example `dim`, `var`, or `att` for functions dealing with dimensions, variables, or attributes. To distinguish the various I/O operations for variables, a single character modifier is appended to `var`:

`var` entire variable access `var1` single value access `vara` array or array section access `vars` strided access to a subsample of values `varm` mapped access to values not contiguous in memory

At the end of the name for variable and attribute functions, there is a component indicating the type of the final argument: `text`, `uchar`, `schar`, `short`, `int`, `long`, `float`, or `double`. This part of the function name indicates the type of the data container you are using in your program: character string, unsigned char, signed char, and so on.

Also, all macro names in the public C interface begin with the prefix `NF_`. For example, the macro which was formerly `MAX_NF_NAME` is now `NF_MAX_NAME`, and the former `FILL_FLOAT` is now `NF_FILL_FLOAT`.

As previously mentioned, all the old names are still supported for backward compatibility.

B.3 Type Conversion

With the new interface, users need not be aware of the external type of numeric variables, since automatic conversion to or from any desired numeric type is now available. You can use this feature to simplify code, by making it independent of external types. The elimination of `void*` pointers provides detection of type errors at compile time that could not be detected with the previous interface. Programs may be made more robust with the new interface, because they need not be changed to accommodate a change to the external type of a variable.

If conversion to or from an external numeric type is necessary, it is handled by the library. This automatic conversion and separation of external data representation from internal data types will become even more important in netCDF version 4, when new external types will be added for packed data for which there is no natural corresponding internal type, for example, arrays of 11-bit values.

Converting from one numeric type to another may result in an error if the target type is not capable of representing the converted value. (In netCDF-2, such overflows can only happen in the XDR layer.) For example, a float may not be able to hold data stored externally as an `NF_DOUBLE` (an IEEE floating-point number). When accessing an array of values, an `NF_ERANGE` error is returned if one or more values are out of the range of representable values, but other values are converted properly.

Note that mere loss of precision in type conversion does not return an error. Thus, if you read double precision values into an int, for example, no error results unless the magnitude of the double precision value exceeds the representable range of ints on your platform. Similarly, if you read a large integer into a float incapable of representing all the bits of the integer in its mantissa, this loss of precision will not result in an error. If you want to avoid such precision loss, check the external types of the variables you access to make sure you use an internal type that has a compatible precision.

The new interface distinguishes arrays of characters intended to represent text strings from arrays of 8-bit bytes intended to represent small integers. The interface supports the internal types `text`, `uchar`, and `schar`, intended for text strings, unsigned byte values, and signed byte values.

The `_uchar` and `_schar` functions were introduced in netCDF-3 to eliminate an ambiguity, and support both signed and unsigned byte data. In netCDF-2, whether the external `NF_BYTE` type represented signed or unsigned values was left up to the user. In netcdf-3, we treat `NF_BYTE` as signed for the purposes of conversion to `short`, `int`, `long`, `float`, or `double`. (Of course, no conversion takes place when the internal type is signed char.) In the `_uchar` functions, we treat `NF_BYTE` as if it were unsigned. Thus, no `NF_ERANGE` error can occur converting between `NF_BYTE` and unsigned char.

B.4 Error handling

The new interface handles errors differently than netCDF-2. In the old interface, the default behavior when an error was detected was to print an error message and exit. To get control of error handling, you had to set flag bits in a global variable, `ncopts`, and to determine the cause of an error, you had to test the value of another global variable `ncerr`.

In the new interface, functions return an integer status that indicates not only success or failure, but also the cause of the error. The global variables `ncerr` and `ncopt` have been eliminated. The library will never try to print anything, nor will it call `exit` (unless you are using the netCDF version 2 compatibility functions). You will have to check the function return status and do this yourself. We eliminated these globals in the interest of supporting parallel (multiprocessor) execution cleanly, as well as reducing the number of assumptions about the environment where netCDF is used. The new behavior should provide better support for using netCDF as a hidden layer in applications that have their own GUI interface.

B.5 NF_LONG and NF_INT

Where the netCDF-2 interface used `NF_LONG` to identify an external data type corresponding to 32-bit integers, the new interface uses `NF_INT` instead. `NF_LONG` is defined to have the same value as `NF_INT` for backward compatibility, but it should not be used in new code. With new 64-bit platforms using `long` for 64-bit integers, we would like to reduce the confusion caused by this name clash. Note that there is still no netCDF external data type corresponding to 64-bit integers.

B.6 What's Missing?

The new C interface omits three "record I/O" functions, `ncrecput`, `ncrecget`, and `ncrecinq`, from the netCDF-2 interface, although these functions are still supported via the netCDF-2 compatibility interface.

This means you may have to replace one record-oriented call with multiple type-specific calls, one for each record variable. For example, a single call to `ncrecput` can always be replaced by multiple calls to the appropriate `NF_PUT_VAR` functions, one call for each variable accessed. The record-oriented functions were omitted, because there is no simple way to provide type-safety and automatic type conversion for such an interface.

There is no function corresponding to the `nctypelen` function from the version 2 interface. The separation of internal and external types and the new type-conversion interfaces make `nctypelen` unnecessary. Since users read into and write out of native types, the `sizeof` operator is perfectly adequate to determine how much space to allocate for a value.

In the previous library, there was no checking that the characters used in the name of a netCDF object were compatible with CDL restrictions. The `ncdump` and `ncgen` utilities that use CDL permit only alphanumeric characters, `"_"` and `"-"` in names. Now this restriction is also enforced by the library for creation of new dimensions, variables, and attributes. Previously existing components with less restrictive names will still work OK.

B.7 Other Changes

There are two new functions in netCDF-3 that don't correspond to any netCDF-2 functions: `NF_INQ_LIBVERS` and `NF_STRERROR`. The version of the netCDF library in use is returned as a string by `NF_INQ_LIBVERS`. An error message corresponding to the status returned by a netCDF function call is returned as a string by the `NF_STRERROR` function.

A new `NF_SHARE` flag is available for use in an `NF_OPEN` or `NF_CREATE CALL`, to suppress the default buffering of accesses. The use of `NF_SHARE` for concurrent access to a netCDF dataset means you don't have to call `NF_SYNC` after every access to make sure that disk updates are synchronous. It is important to note that changes to ancillary data, such as attribute values, are not propagated automatically by use of the `NF_SHARE` flag. Use of the `NF_SYNC` function is still required for this purpose.

The version 2 interface had a single inquiry function, `ncvarinq` for getting the name, type, and shape of a variable. Similarly, only a single inquiry function was available for getting information about a dimension, an attribute, or a netCDF dataset. When you only wanted a subset of this information, you had to provide `NULL` arguments as placeholders for the unneeded information. The new interface includes additional inquire functions that return each item separately, so errors are less likely from miscounting arguments.

The previous implementation returned an error when 0-valued count components were specified in `ncvarput` and `ncvarget` calls. This restriction has been removed, so that now functions in the `NF_PUT_VAR` AND `NF_GET_VAR` families may be called with 0-valued count components, resulting in no data being accessed. Although this may seem useless, it simplifies some programs to not treat 0-valued counts as a special case.

The previous implementation returned an error when the same dimension was used more than once in specifying the shape of a variable in `ncvardef`. This restriction is relaxed in the netCDF-3 implementation, because an autocorrelation matrix is a good example where using the same dimension twice makes sense.

In the new interface, units for the `imap` argument to the `NF_PUT_VARM` AND `NF_GET_VARM` families of functions are now in terms of the number of data elements of the desired internal type, not in terms of bytes as in the netCDF version-2 mapped access interfaces.

Following is a table of netCDF-2 function names and names of the corresponding netCDF-3 functions. For parameter lists of netCDF-2 functions, see the netCDF-2 User's Guide.

<code>ncabort</code>	<code>NF_ABORT</code>
<code>ncattcopy</code>	<code>NF_COPY_ATT</code>
<code>ncattdel</code>	<code>NF_DEL_ATT</code>
<code>ncattget</code>	<code>NF_GET_ATT_DOUBLE</code> , <code>NF_GET_ATT_FLOAT</code> , <code>NF_GET_ATT_INT</code> , <code>NF_GET_ATT_LONG</code> , <code>NF_GET_ATT_SCHAR</code> , <code>NF_GET_ATT_SHORT</code> , <code>NF_GET_ATT_TEXT</code> , <code>NF_GET_ATT_UCHAR</code>
<code>ncattinq</code>	<code>NF_INQ_ATT</code> , <code>NF_INQ_ATTID</code> , <code>NF_INQ_ATTLEN</code> , <code>NF_INQ_ATTTYPE</code>
<code>ncattname</code>	<code>NF_INQ_ATTNAME</code>

```

ncattput  NF_PUT_ATT_DOUBLE, NF_PUT_ATT_FLOAT, NF_PUT_ATT_INT,
          NF_PUT_ATT_LONG, NF_PUT_ATT_SCHAR, NF_PUT_ATT_SHORT,
          NF_PUT_ATT_TEXT, NF_PUT_ATT_UCHAR

ncattrename
          NF_RENAME_ATT

ncclose  NF_CLOSE

nccreate  NF_CREATE

ncdimdef  NF_DEF_DIM

ncdimid  NF_INQ_DIMID

ncdiminq  NF_INQ_DIM, NF_INQ_DIMLEN, NF_INQ_DIMNAME

ncdimrename
          NF_RENAME_DIM

ncendef  NF_ENDDEF

ncinquire
          NF_INQ, NF_INQ_NATTS, NF_INQ_NDIMS, NF_INQ_NVARS,
          NF_INQ_UNLIMDIM

ncopen  NF_OPEN

ncrecget  (none)

ncrecinq  (none)

ncrecput  (none)

ncredef  NF_REDEF

ncsetfill
          NF_SET_FILL

ncsync  NF_SYNC

nctypelen
          (none)

ncvardef  NF_DEF_VAR

ncvarget  NF_GET_VARA_DOUBLE, NF_GET_VARA_FLOAT, NF_GET_VARA_INT,
          NF_GET_VARA_LONG, NF_GET_VARA_SCHAR, NF_GET_VARA_SHORT,
          NF_GET_VARA_TEXT, NF_GET_VARA_UCHAR

ncvarget1
          NF_GET_VAR1_DOUBLE, NF_GET_VAR1_FLOAT, NF_GET_VAR1_INT,
          NF_GET_VAR1_LONG, NF_GET_VAR1_SCHAR, NF_GET_VAR1_SHORT,
          NF_GET_VAR1_TEXT, NF_GET_VAR1_UCHAR

ncvargetg
          NF_GET_VARM_DOUBLE, NF_GET_VARM_FLOAT, NF_GET_VARM_INT,
          NF_GET_VARM_LONG, NF_GET_VARM_SCHAR, NF_GET_VARM_SHORT,
          NF_GET_VARM_TEXT, NF_GET_VARM_UCHAR, NF_GET_VARS_DOUBLE,

```

```

        NF_GET_VARS_FLOAT, NF_GET_VARS_INT, NF_GET_VARS_LONG,
        NF_GET_VARS_SCHAR, NF_GET_VARS_SHORT, NF_GET_VARS_TEXT,
        NF_GET_VARS_UCHAR
ncvarid  NF_INQ_VARID
ncvarinq NF_INQ_VAR,          NF_INQ_VARDIMID,          NF_INQ_VARNAME,
        NF_INQ_VARNATTS, NF_INQ_VARNDIMS, NF_INQ_VARTYPE
ncvarput NF_PUT_VARA_DOUBLE, NF_PUT_VARA_FLOAT, NF_PUT_VARA_INT,
        NF_PUT_VARA_LONG, NF_PUT_VARA_SCHAR, NF_PUT_VARA_SHORT,
        NF_PUT_VARA_TEXT, NF_PUT_VARA_UCHAR
ncvarput1
        NF_PUT_VAR1_ double, NF_PUT_VAR1_ float, NF_PUT_VAR1_ int,
        NF_PUT_VAR1_ long, NF_PUT_VAR1_ schar, NF_PUT_VAR1_ short,
        NF_PUT_VAR1_ text, NF_PUT_VAR1_ uchar
ncvarputg
        NF_PUT_VARM_DOUBLE, NF_PUT_VARM_FLOAT, NF_PUT_VARM_INT,
        NF_PUT_VARM_LONG, NF_PUT_VARM_SCHAR, NF_PUT_VARM_SHORT,
        NF_PUT_VARM_TEXT, NF_PUT_VARM_UCHAR, NF_PUT_VARS_DOUBLE,
        NF_PUT_VARS_FLOAT, NF_PUT_VARS_INT, NF_PUT_VARS_LONG,
        NF_PUT_VARS_SCHAR, NF_PUT_VARS_SHORT, NF_PUT_VARS_TEXT,
        NF_PUT_VARS_UCHAR
ncvarrename
        NF_rename_var
(none)  NF_INQ_libvers
(none)  NF_strerror

```


Index

A

attributes, adding 4

C

compiling with netCDF library 5
 creating dataset 1

D

datasets, introduction 7
 dimensions, adding 4

H

HANDLE_ERR 8

I

interface descriptions 7

L

linking to netCDF library 5

N

NF__CREATE 10
 NF__ENDDEF 15
 NF__OPEN 13
 NF_ABORT 21
 NF_CLOSE 17
 NF_CLOSE, typical use 1
 NF_COPY_ATT 72
 NF_CREATE 9
 NF_CREATE, typical use 1
 NF_DEF_DIM 25
 NF_DEF_DIM, typical use 1
 NF_DEF_VAR 32
 NF_DEF_VAR, typical use 1
 NF_DEL_ATT 75
 NF_ENDDEF 15
 NF_ENDDEF, typical use 1
 NF_GET_ATT, typical use 2
 NF_GET_ATT_ type 70
 NF_GET_VAR, typical use 2

NF_GET_VAR_ type 49
 NF_GET_VAR1_ type 47
 NF_GET_VARA_ type 50
 NF_GET_VARM_ type 54
 NF_GET_VARS_ type 52
 NF_INQ Family 18
 NF_INQ, typical use 3
 NF_INQ_ATT Family 68
 NF_INQ_ATTNAME, typical use 3
 NF_INQ_DIM Family 27
 NF_INQ_DIMID 26
 NF_INQ_DIMID, typical use 2
 NF_INQ_FORMAT 18
 NF_INQ_LIBVERS 8
 NF_INQ_NATTS 18
 NF_INQ_NDIMS 18
 NF_INQ_NVARS 18
 NF_INQ_UNLIMDIM 18
 NF_INQ_VAR family 34
 NF_INQ_VARID 34
 NF_INQ_VARID, typical use 2
 NF_LONG and NF_INT 86
 NF_OPEN 12
 NF_PUT_ATT, typical use 1
 NF_PUT_ATT_ type 66
 NF_PUT_VAR, typical use 1
 NF_PUT_VAR_ type 38
 NF_PUT_VAR1_ type 36
 NF_PUT_VARA_ type 39
 NF_PUT_VARM_ type 44
 NF_PUT_VARS_ type 42
 NF_REDEF 14
 NF_RENAME_ATT 74
 NF_RENAME_DIM 28
 NF_RENAME_VAR 60
 NF_SET_DEFAULT_FORMAT 23
 NF_SET_FILL 22
 NF_STRERROR 8
 NF_SYNC 19

R

reading dataset with unknown names 3
 reading datasets with known names 2

V

variables, adding 4

