Subject: Re: Add a new variable to netCDF file?
Posted by Paul Van Delst[1] on Tue, 25 Oct 2011 15:02:31 GMT

hello,

Kenneth P. Bowman wrote:
> In article <j86e3h$5of$1@speranza.aioe.org>,
>   Paul van Delst <paul.vandelst@noaa.gov> wrote:
>
>> Hello,
>>
>> Ashley Berg wrote:
>>> So, I know how to overwrite a variable in a netCDF file, but how do I
>>> add a new one?
>> 0) Open the existing file for writing (NCDF_OPEN)
> .
> .
> .
>> 8) Close the file (NCDF_CLOSE)
>>
>> cheers,
>>
>> paulv
>
> And, by the way, this usually requires creating a new file and
> copying the contents of the original file. So, it may not be very
> efficient, and you probably want to operate on a *copy* of your input
> file in case something goes wrong during the process.

Huh. I have never encountered this.

From the "Adding new dimensions, variables, and attributes" section of the "netCDF Fortran90
interface guide":

<quote>
An existing netCDF dataset can be extensively altered. New dimensions, variables, and attributes
can be added and
existing ones renamed, and existing attributes can be deleted.
</quote>

Also, from the "Variables" section of the same document:

<quote>
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.
</quote>

Seems quite clear to me. I've been doing it in v3.x of netCDF for years now. I would consider *any* failure of an API to
do this correctly as a critical bug.

I will second Ken's comment about efficiency: adding a new dimension, or a variable attribute, or a global attribute, to
a large existing netCDF dataset can be quite slow since all of the following data has to be "pushed down" to make way
for it. At least, that's my perception of the process for netcdf v3.x. NetCDF 4.x may be supa-efficient regardless.

cheers,

paulv