Subject: Reading HDF5 signed bytes gives strange results
Posted by Maarten[1] on Tue, 22 Aug 2006 16:38:52 GMT
View Forum Message <> Reply to Message

Hi,

When reading a signed byte from and HDF5 file I get all kinds of nasty results. Of course, IDL doesn't know about signed bytes, but I expect it to be nice to the bit-images, and read -127 (signed) as 129 (unsigned). However, none of that seems to happen with the routines I use.

When reading with h5_parse(file, /read), I get the fields, with -127 (the fill value, in case you're wondering) replaced by 129, as expected in 2-complement notation. All would be well, If the rest of my software could use h5_parse. But I can't use it for various reasons.

When using h5f_open, h5d_open, h5d_read & friends, the value -127 is repleaced by 0. The fill value (an attribute) is replaced by 0. And 0 is a perfectly valid data-value. When I then try to filter for fill values, I throw out quite a few valid values.

How can I force the field to be read differently so that values < 0 end up at values > 127, while not resorting to h5_parse?

Maarten

Subject: Re: Reading HDF5 signed bytes gives strange results
Posted by Maarten[1] on Wed, 23 Aug 2006 15:17:45 GMT
View Forum Message <> Reply to Message

Hi,

This is a follow-up to my own question.

Maarten wrote:

> When reading a signed byte from and HDF5 file I get all kinds of nasty
> results. Of course, IDL doesn't know about signed bytes, but I expect
> it to be nice to the bit-images, and read -127 (signed) as 129
> (unsigned). However, none of that seems to happen with the routines I
> use.
>
> When reading with h5_parse(file, /read), I get the fields, with -127
> (the fill value, in case you're wondering) replaced by 129, as expected
> in 2-complement notation. All would be well, If the rest of my software
> could use h5_parse. But I can't use it for various reasons.

I can use h5_parse as a work around, but further investigation shows
even stranger things. Perhaps some real wizards can halp me here.

> When using h5f_open, h5d_open, h5d_read & friends, the value -127 is
> repleaced by 0. The fill value (an attribute) is replaced by 0. And 0
> is a perfectly valid data-value. When I then try to filter for fill
> values, I throw out quite a few valid values.

I tried to create a minimal example using h5f_open, h5d_open, h5d_read
& friends working on a small hdf5 sample file. After I had created that
file, I noticed that the values were read correctly in this minimal
example.

Some further investigation showed that using this test software to read
data from the "real" file, also created the expected results (fill
value at 129). This is a bit strange, as the sequence of commands is
the same in both cases. As a further test, I tried to use the "real"
function I have to read the data in this test file. This works when
calling the function directly, but fails when the call is made several
levels deep.

If you're interested: the software can be found at [1], and it is meant
to read data from OMI, for instance OMI DOAS Ozone columns [2] and [3].
If you want to reproduce the effect: I set a breakpoint at line 98 of
the file read_hdfeos5_data_or_geo_field.pro, and can see a min/max of 0
and 100, when I expect 0 and 129 when reading the CloudFraction field
in the dataproduct mentioned above.

Any clues, hints and other details to deal with this are appreciated.

Maarten

[1] http://www.knmi.nl/omi/research/validation/cama/
[2]
http://avdc.gsfc.nasa.gov/Data/Aura/OMI/OMDOAO3/OMDOAO3_READ ME_File.html

[3] http://disc.sci.gsfc.nasa.gov/data/datapool/OMI/Level2/OMDOA O3

Subject: Re: Reading HDF5 signed bytes gives strange results
Posted by David Fanning on Wed, 23 Aug 2006 15:23:12 GMT
View Forum Message <> Reply to Message

Maarten writes:

> I tried to create a minimal example using h5f_open, h5d_open, h5d_read
> & friends working on a small hdf5 sample file. After I had created that

> file, I noticed that the values were read correctly in this minimal
> example.

Ain't that always the way! :-)

Cheers,

David
--
David Fanning, Ph.D.
Fanning Software Consulting, Inc.
Coyote's Guide to IDL Programming: http://www.dfanning.com/
Sepore ma de ni thui. (Opata Indian saying, meaning "Perhaps thou
speakest truth.")

Subject: Re: Reading HDF5 signed bytes gives strange results
Posted by Maarten[1] on Thu, 24 Aug 2006 06:36:07 GMT
View Forum Message <> Reply to Message

David Fanning wrote:
> Maarten writes:
>
>> I tried to create a minimal example using h5f_open, h5d_open, h5d_read
>> & friends working on a small hdf5 sample file. After I had created that
>> file, I noticed that the values were read correctly in this minimal
>> example.
>
> Ain't that always the way! :-)

And here I am, hoping someone has a bright idea :-(

I _suspect_ that there is some missing balancing of
h5(?)_open/h5\1_close pair, but that is only a suspicion. Without
proper tracing functionality in IDL (i.e. record all calls to functions
in the order they are called), I don't think it is going to be easy to
find. It is reproducible over multiple machines though (Linux and
Windows), so I should be able to find something, eventually.

Maarten
