
Subject: Re: Error in reading large Fortran unformatted files

Posted by [Nigel Wade](#) on Fri, 18 Feb 2011 17:15:35 GMT

[View Forum Message](#) <> [Reply to Message](#)

On 18/02/11 15:14, OM wrote:

>
> So I take it the only viable solution you can think of is as suggested
> by Ken - to break down the file into manageable bits?
>

I don't know of any other. I think it would be easier to write, and to read back, as 1024 records. Programming the loop parameters would be simpler.

I've had a play with a little FORTRAN program, compiled using gfortran, which writes files >2GB. It seems the first 32bit word is always 2147483657, regardless of the length of the record actually output. Maybe this is some special flag to the FORTRAN I/O library to indicate a large record, I don't know. I'm not about to load a file that size into a binary editor to look at it.

Also, the file is actually 8bytes longer than it should be for a single record. So it may be that FORTRAN is actually splitting the record into two, and it knows this because of that special record length indicator. Presumably IDL doesn't understand this new "feature" of the GNU FORTRAN compiler, and fails to read the file.

--

Nigel Wade
