Subject: Re: readu sees unexpected records
Posted by David Fanning on Wed, 04 May 2005 13:27:17 GMT
View Forum Message <> Reply to Message

Hjalti Sig writes:

> I found out that when reading unformatted data, written with a fortran
> 90 program, the readu command detects a couple of extra records before
> and after each line written with the fortran program. For example, I
> wrote a 10x10 array, containing the numbers from 1-100 in a file.
> Reading the file using readu, the content of the file appears as a
> 14x10 array, with very small numbers in columns 1, 2, 13 and 14,
> (producing floating underflow), and the correct numbers in the columns
> 3-12.
> In the data file I actually had to read there are three extra records
> before and two after each line.
> These extra records do certainly not appear when reading the files
> with a fortran program.
>
> I was curious if anyone had an explanation of this.

I'd have a look at the F77_UNFORMATTED keyword to the OPEN
command.

Cheers,

David

--
David Fanning, Ph.D.
Fanning Software Consulting, Inc.
Coyote's Guide to IDL Programming: http://www.dfanning.com/

Subject: Re: readu sees unexpected records
Posted by Paul Van Delst[1] on Wed, 04 May 2005 13:49:04 GMT
View Forum Message <> Reply to Message

Hjalti Sig wrote:
> Hello group
> I found out that when reading unformatted data, written with a fortran
> 90 program, the readu command detects a couple of extra records before
> and after each line written with the fortran program. For example, I
> wrote a 10x10 array, containing the numbers from 1-100 in a file.
> Reading the file using readu, the content of the file appears as a
> 14x10 array, with very small numbers in columns 1, 2, 13 and 14,
> (producing floating underflow), and the correct numbers in the columns
> 3-12.

> In the data file I actually had to read there are three extra records
> before and two after each line.
> These extra records do certainly not appear when reading the files
> with a fortran program.

Fortran (currently) only really supports record-based I/O. For *sequential* access output in Fortran, each variable length output record is preceded and followed with a a record marker indicating the length of the record (the prefix marker is for checking the length, the suffix marker for facilitating BACKSPACE-ing). To get IDL to emulate this type of I/O, do as DFanning suggested and use the /F77_UNFORMATTED keyword in your OPENR statement. If
your output records are all the same length, you can modify your Fortran code to output in *direct* access and you will no longer have the record marker problem.

FWIW, the Fortran2003 standard has a new type of I/O access called "STREAM" (along with the venerable SEQUENTIAL and DIRECT) which produces the output that just about every non-Fortran programming type expects. A few Fortran compilers already support this feature (some fully standard, others only sorta.)

paulv

--
Paul van Delst
CIMSS @ NOAA/NCEP/EMC

Subject: Re: readu sees unexpected records
Posted by K. Bowman on Wed, 04 May 2005 13:55:39 GMT
View Forum Message <> Reply to Message

In article <e1330fff.0505040322.67613e37@posting.google.com>,
hjalti@vatnaskil.is (Hjalti Sig) wrote:

> Hello group
> I found out that when reading unformatted data, written with a fortran
> 90 program, the readu command detects a couple of extra records before
> and after each line written with the fortran program. For example, I
> wrote a 10x10 array, containing the numbers from 1-100 in a file.
> Reading the file using readu, the content of the file appears as a
> 14x10 array, with very small numbers in columns 1, 2, 13 and 14,
> (producing floating underflow), and the correct numbers in the columns
> 3-12.
> In the data file I actually had to read there are three extra records
> before and two after each line.
> These extra records do certainly not appear when reading the files
> with a fortran program.
>

> I was curious if anyone had an explanation of this.
>
> Regards, Hjalti

Fortran I/O is record oriented. Each Fortran unformatted WRITE produces one record in the output file. The number of data bytes in each record is written as a 4 byte integer at the beginning and end of each record. You can read the record lengths yourself, or use the F77_UNFORMATTED keyword.

Ken Bowman

Subject: Re: readu sees unexpected records
Posted by hjalti on Fri, 06 May 2005 13:12:44 GMT
View Forum Message <> Reply to Message

Thank you guys. The f77_unformatted fixed the problem. The IDL documentation however says this keyword is UNIX specific (see below), and I am running on a windows XP. So the documentation does not seem quite accurate about this...?

F77_UNFORMATTED
Unformatted variable-length record files produced by UNIX FORTRAN programs contain extra information along with the data in order to allow the data to be properly recovered. This method is necessary because FORTRAN input/output is based on record-oriented files, while UNIX files are simple byte streams that do not impose any record structure. Set the F77_UNFORMATTED keyword to read and write this extra information in the same manner as f77(1), so that data to be processed by both IDL and FORTRAN. See UNIX-Specific Information for further details.

Regards, Hjalti

Subject: Re: readu sees unexpected records
Posted by David Fanning on Fri, 06 May 2005 13:21:01 GMT
View Forum Message <> Reply to Message

Hjalti Sig writes:

> Thank you guys. The f77_unformatted fixed the problem. The IDL
> documentation however says this keyword is UNIX specific (see below),
> and I am running on a windows XP. So the documentation does not seem
> quite accurate about this...?

You'll get used to this. :-)

Cheers,

David

--
David Fanning, Ph.D.
Fanning Software Consulting, Inc.
Coyote's Guide to IDL Programming: http://www.dfanning.com/

Subject: Re: readu sees unexpected records
Posted by Paul Van Delst[1] on Fri, 06 May 2005 15:30:47 GMT
View Forum Message <> Reply to Message

Hjalti Sig wrote:
> Thank you guys. The f77_unformatted fixed the problem. The IDL
> documentation however says this keyword is UNIX specific (see below),
> and I am running on a windows XP. So the documentation does not seem
> quite accurate about this...?

I'd be more forgiving about the IDL documentation in this case because the common Fortran behaviour for sequential, unformatted I/O (4-byte markers at begin and end of record) is not stipulated by the Fortran standard (at least, anywhere I can find it). I think most vendors decided that the current system was the most efficient to satisfy those things that *are* standardised. It is not guaranteed that every Fortran compiler on every platform will produce sequential, unformatted files the same way (e.g. some older systems only use 2-byte record markers. Some only put markers at the beginning of the record). Over time, fortuitously, a "defacto" standard for sequential, unformatted file formats has come to pass - most likely due to customer complaints to maverick Fortran compiler vendors to pull their finger out and get their act together.

The Unix reference is maybe due to people sharing files between SGIs, Suns, IBMs, etc back when those workstation type of machines were all the rage and customers wanted the file sharing to be independent of the vendor platform/compiler? Then again, I've never really used a Windows machine (for anything important) so who knows?

paulv

>
> F77_UNFORMATTED
> Unformatted variable-length record files produced by UNIX FORTRAN
> programs contain extra information along with the data in order to
> allow the data to be properly recovered. This method is necessary
> because FORTRAN input/output is based on record-oriented files, while
> UNIX files are simple byte streams that do not impose any record
> structure. Set the F77_UNFORMATTED keyword to read and write this

> extra information in the same manner as f77(1), so that data to be
> processed by both IDL and FORTRAN. See UNIX-Specific Information for
> further details.
> ---------------
>
> Regards, Hjalti

--
Paul van Delst
CIMSS @ NOAA/NCEP/EMC
