Subject: Re: reading/writing large files

Posted by ben.bighair on Sun, 03 Feb 2013 15:06:25 GMT

View Forum Message <> Reply to Message

On Friday, February 1, 2013 10:15:25 AM UTC-5, rr...@stsci.edu wrote:

Then to read it, I create an assoc and loop over the rows. I've estimated it takes about 1.7e-4 s (on my 4 year-old laptop) so I'm guessing of order 3 minutes for 10^6. I can live with 3 min (if I have to), it just seems that given all the file types that IDL can read/write there's a way to do just this.

Hi,

It sounds like you are doing some cumulative processing of successive rows - not random access stuff across rows. So, what if you made your associated variable a bigger chunk of rows? Instead of creating a per-row associated variable, maybe you could create one of 10,000 rows? That way you'll pack more umph into each file I/O. Your output routine would have to save the data in the same sized chunks that your input routine would be reading in - that could mean the tail end of the file has a number of dummy rows. Keeping track of your navigation within the file should not be too terribly hard - maybe something like Mike Galloy's Collection objects would be helpful.

Cheers, Ben