
Subject: reading/writing large files

Posted by [Russell Ryan](#) on Fri, 01 Feb 2013 15:15:25 GMT

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

Okay gang I've been working on this for a few days and have given up.

I've got this simulation that outputs an array of floating point numbers (roughly 5000 or so), which I want to put into a file. If the file exists, I want to append to it; if not, I want to create it. I want to do this of order a million times (at least append of order a million times). When the simulation finishes, I want to read these numbers and do some post-processing. I don't want to read the entire file at once because I'm afraid I'll run into memory problems (especially since I can envision doing the appending 10^7 or even 10^8 times). So, instead I'd like to read say all 10^6 (or 10^7 or 10^8) trials of the k-th element of the array and get a single floating-point array of 10^6 elements (or what have you). Basically, I'm envisioning a table with say 5000ish columns but the number of rows is variable, and I want to read the k-th column.

Any ideas on the most efficient way of doing this? Obviously, my idea of a table is just illustrative and I don't actually care the format of the data in the file --- or even the file type. Currently, I'm opening a file with open and writing unformatted data to it with writeu. 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.

Any ideas?

Russell

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
