Subject: ram disk

Posted by Grady Daub on Thu, 24 Jun 1999 07:00:00 GMT

View Forum Message <> Reply to Message

I'm reading in some data that is in a weird format and I've devised a technique which involves reading the one original data file, making a temporary file for each different type of data, doing some stuff with the data through reading the temp files, then, creating a new single data file.

Does IDL have a feature where, instead of using the hard drive for the temporary files, the RAM is used? (Us amiga users call it a RAM disk.) No, I'm not talking about making a huge (huge) array

Subject: Re: ram disk

Posted by J.D. Smith on Mon, 28 Jun 1999 07:00:00 GMT

View Forum Message <> Reply to Message

Grady Daub wrote:

>

- > I'm reading in some data that is in a weird format and I've devised a
- > technique which involves reading the one original data file, making a
- > temporary file for each different type of data, doing some stuff with
- > the data through reading the temp files, then, creating a new single
- > data file.

>

- > Does IDL have a feature where, instead of using the hard drive for the
- > temporary files, the RAM is used? (Us amiga users call it a RAM disk.)
- > No, I'm not talking about making a huge (huge) array

You could use... a RAM disk. Joking aside, most OS's provide this feature. If you're using a *good* OS, however, you'll really never need this, since it will provide aggressive memory buffering of your file IO by default, if you have the memory, and performance will be the same either way.

Having said that, however, as far as I can see, there is no advantage to trying to stuff everything in a RAM disk file... The huge (huge) array you mention would not be any more huge (huge) than the file(s) you created in RAM. So long as you use a compact representation within IDL itself, you have nothing to lose and everthing to gain by working entirely in memory. It's not hard to keep the data compact in IDL, which only adds a little bookeeping to each variable, but otherwise is quite close to machine format. If your data cannot possibly all fit in memory, you could consider breaking it on input into chunks and outputting the final data file one chunk at a time.

Good Luck,

JD

-
J.D. Smith |*| WORK: (607) 255-5842

Cornell University Dept. of Astronomy |*| (607) 255-6263
304 Space Sciences Bldg. |*| FAX: (607) 255-5875

|*|

Ithaca, NY 14853