
Subject: Re: Read Total lines in an ASCII file
Posted by [wmconnolley](#) on Wed, 18 Dec 2002 22:12:34 GMT
[View Forum Message](#) <> [Reply to Message](#)

Mark Hadfield <m.hadfield@niwa.co.nz> wrote:

> I think you mean my method 1 (ie two passes: read the file once to count the
> lines, create the result array, then read the file again to get the data
> into the array).

Oops yes.

> So I revisited method 1, comparing three ways of counting the lines in the
> file:

> 1a - Count lines with IDL readf statements in a while loop
> 1b - Count lines by spawning "wc -l"
> 1c - Count lines with IDL 5.6 FILE_LINES function

> and here are the times taken to read the same 20,000-line, uncompressed file
> on my hard drive

> 1a 0.24 s
> 1b 0.32 s
> 1c 0.09 s

> FILE_LINES is the clear winner. (Isn't it a pity it doesn't accept a
> COMPRESS keyword!)

> Spawning "wc -l" is the slowest. Note that this is on Windows 2000 with the
> Cygwin "wc" command. Unix is much faster at spawning subprocesses than
> Windows, so method 1b may be competitive there.

How interesting. Thanks for doing all this. I would certainly hope that
spawn would be a bit faster under unix. I have a slight feeling that there
are options to spawn that can make it a bit faster (something about not
copying the env perhaps?).

But also intereting to see that wc -l is not quite as o/s specific as I'd
assumed!

-W.

--

William M Connolley | wmc@bas.ac.uk | <http://www.nerc-bas.ac.uk/icd/wmc/>
Climate Modeller, British Antarctic Survey | Disclaimer: I speak for myself
I'm a .signature virus! copy me into your .signature file & help me spread!
