Subject: little and big endian -- once more
Posted by Martin Schultz on Wed, 12 May 1999 07:00:00 GMT
View Forum Message <> Reply to Message

Hi all,

took me a while to realize that it is the machine architecture not the OS that determines the byte swapping -- in fact I needed to have IDL on linux (after using it on an SGI and with Windows) to figure that out ... Anyway, I now devised the following test for byte swapping which is applied in my open_file routine and handles everything transparently. Please tell me if there are any other machine architectures that are big endian and what their !version.arch tag would be.

Thanks,
Martin.

```
   ; Extra attention to binary (f77 type) files:
   ; Set keywords for SWAP_IF_LITTLE_ENDIAN and SWAP_IF_BIG_ENDIAN
   ; depending on operating system
   ; do it only for read only opening

   if not (keyword_set(write)) then begin
      Swap_If_Big_Endian = (strupcase(!version.arch) eq 'X86')
      Swap_If_Little_Endian = 1 - Swap_If_Big_Endian
   endif

 ; debug uses little_endian function from Robert Mallozi
 print,'#DEBUG: little_endian = ',little_endian()
 help,Swap_If_Little_Endian,Swap_If_Big_Endian
```

--

 ||||||||||||||||\\\\\\\\\\\\\\\----------////////////// //||||||||||||||||
Martin Schultz, DEAS, Harvard University, 29 Oxford St., Pierce 109,
Cambridge, MA 02138 phone (617) 496 8318 fax (617) 495 4551
e-mail mgs@io.harvard.edu web http://www-as/people/staff/mgs/

Subject: Re: little and big endian -- once more
Posted by Liam Gumley on Thu, 13 May 1999 07:00:00 GMT
View Forum Message <> Reply to Message

Dick Jackson wrote:
> I'm not sure if I'm missing some subtlety here, but wouldn't the (newish)
> keywords /SWAP_IF_[BIG|LITTLE]_ENDIAN to IDL's OPEN and BYTEORDER routines
> work in all cases? I've used them to good effect for cross-platform work.

>
> From the help file:
>
> [...] it only takes effect if the current system has [big|little] endian
> byte ordering. This keyword does not refer to the byte ordering of the input
> data, but to the computer hardware.

This method assumes that you keep track of which platform was used to create the datafile. I'd rather just interrogate the datafile.

Or of course, you can use netCDF. We're working an a specification right now for a simplified set of 3 functions to perform netCDF read, write, and inquire operations. I'll post them when they're ready.

Liam E. Gumley
Space Science and Engineering Center, UW-Madison
http://cimss.ssec.wisc.edu/~gumley
