Subject: Re: IDL 8.2.2 released
Posted by timothyja123 on Sat, 09 Feb 2013 10:44:01 GMT
View Forum Message <> Reply to Message

> This bug is on the plate for IDL 8.3.

This is great to hear.

> I just did some quick tests, and it looks like some of the speed difference is because the IDL save files are saved in "big endian" format. When you read this on a Windows, Linux, or Mac machine, it needs to convert all of the arrays to little endian. Something we are doing in the C code must be slower than the Python big->little endian conversion.

The thing I noticed was that each array was restored one at a time, as I'm running a 6 core Xeon processor I guess I was hoping that this would become multithreaed so that multiple arrays would be restored in parallel although I guess I'm making assumptions about the inner-workings of IDL.

> Just out of curiosity, are you using scipy's "idlsave" package to read IDL save files into Python?

Yeah I used the original version of the idlsave package (which scipy added to their library)

Thanks to everyone for the suggestions on using a different format to speed things up. I'm not sure if this is an option right now as the data is coming from a different part of the orginisation and its up to them to change the format but is definitely something to consider.

Tim