Subject: Specifying a particular LAPACK implementation using the standard IDL7 LAPACK DLM?

Posted by Matt Francis on Mon, 07 Jan 2013 03:32:51 GMT

View Forum Message <> Reply to Message

I have a task that has a bottleneck inverting large matrices. My understanding is that the LAPACK DLM that comes with IDL7 that is loaded when LA_INVERT is called for the first time is a serial implementation? I am running my code on an 8 core machine so would expect a reasonable speed up in the matrix inversion routine could utilise multiple cores to do the operation in parallel.

There are multi threaded LAPACK implementations, such as ATLAS, that I could use. If possible though, I'd like to avoid creating a new DLM from scratch. It would seem to me to be conceivable that if I could point IDL to the ATLAS (or whatever) multi-threaded LAPACK implementation instead of wherever it currently points to that this might neatly achieve what I want to do.

So, does anyone know if this is possible and if so how to do it? If not, does anyone have any suggestions for the simplest way to achieve what I want, which is parallel (dense) matrix inversion in IDL?