
Subject: Blazing FAST!!! FFT's for IDL

Posted by [David McClain](#) on Thu, 30 Mar 2000 08:00:00 GMT

[View Forum Message](#) <> [Reply to Message](#)

I have some sources available on request to perform very, very, fast 1-D and 2-D FFT's, forward and inverse, single and double precision. The speed derives from a multithreaded manager written in C++ that calls on the Intel Math Kernel implementation of 1-D FFT's. It is multithreaded because we typically use dual and quad Pentium and Xeon machines. The manager code sniffs out how many processors you have and spawns worker threads to match. The arrays are then divided up between the different processor threads. The IDL interface is quite simple, consisting of some data prep code and a bunch of CALL_EXTERNAL's.

We have been using this system for several years now. The Intel MKL expects arrays in power of 2 size, unlike IDL, but it runs roughly 10-100 times faster than IDL's routines the last time I checked about a year ago. It properly scales as $n \log_2 n$ for 1-D and $2n \log_2 n$ for 2-D n square arrays (2-D arrays need not be square). IDL's routines scale quite dreadfully as something on the order of $n^2 (\log_2 n)^2$ which implies some kind of tree search on each Butterfly operation (???). Their routine is nice if you need arbitrary array sizes, but power of 2 can always be used anyway: the resulting spectrum is simply an interpolated spectrum at the intermediate frequencies.

On an old quad-Pentium II machine running at 200 MHz, the FFTX routines in this package performed at roughly 75 MButterflys/sec. IDL ran about 3 MButterflys/sec.

If you are interested just drop me a line. The latest Intel MKL has been sped up to roughly twice its former speed (the speed tests quoted were performed on the old version of the MKL). It is available separately, and I believe, still free, from Intel Corp.

D. McClain
Sr. Scientist
Raytheon Systems Co.
Tucson, AZ
