Subject: Multi-core techniques

Posted by Tim B on Fri, 16 Apr 2010 01:15:50 GMT

View Forum Message <> Reply to Message

I am working with various satellite datasets (e.g. Pathfinder SST) where

most of the prior work has been producing 50km resolution analysis.

Even

an intensive piece of code over a full 30-odd year data set could still be

finished overnight. The algorithms that I've seen don't excessively use

array processing (no FFT's or such) and there is a lot of data independence i.e.

values over the course of a few days are used rather than values over a complete year.

With a move to higher resolution, i.e. 4km, there is significantly more data (approx 150 4km

pixels in a 50km pixel). Given the data independence, my usual approach would be to

create a thread pool around the number of available cores and calculate each data

'piece' independently. However IDL doesn't expose threads to the programmer. So I'd

value any thoughts about how to take advantage of multicore CPU's (heck, even my laptop

is a dual core machine). My thoughts are:

- use C(?) to manage forking separate processes that start IDL and pass parameters to the appropriate procedure to run in IDL
- run a number of IDL programs in parallel, the same code but processing different

temporal regions of the dataset. I can start up IDL in different windows on an 8-core

machine and each seems to be a separate process.

- use a different language/architecture completely :-)

I'd be interested to hear from anyone else trying to take advantage of multicore CPU's..

Tim Burgess