Subject: Re: Speed penalty using START and COUNT with HDF_SD_GETDATA
Posted by Mark Hadfield on Wed, 05 Sep 2001 05:52:35 GMT
View Forum Message <> Reply to Message

"Bob Fugate" <rqfugate@mindspring.com> wrote in message
news:B7BAF61A.2E03%rqfugate@mindspring.com...
> I have a large number of 128x128 pixel arrays stored as SDS's in
> HDF files. Since I am only interested in a 32x32 subset of each
> array, I tried using the START and COUNT keywords to read
> only that part of the array I need ---
> thinking this would be faster and less taxing on memory.
> However, I learned today that it is much faster to read
> in the entire array.
>
> ...
>
> This is a so-so Windows NT machine; IDL 5.4. The data is on a
> server. I have
> a good connection to the server.
>
> Anyone had any similar experiences

I have noticed something similar with IDL's netCDF interface: using the STRIDE keyword seems to be very inefficient. I got the impression that IDL is actually reading in the whole array then extracting a subset.

> ...suggestions on how to speed up reading
> only the part of the array I need?

Have you tried copying the file to a local disk? The local disk's caching may suit the way IDL reads the data better.

Mark Hadfield
m.hadfield@niwa.cri.nz http://katipo.niwa.cri.nz/~hadfield
National Institute for Water and Atmospheric Research

--