
Subject: Re: HDF File Question

Posted by [James Kuyper](#) on Thu, 02 Aug 2007 09:26:41 GMT

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

David Fanning wrote:

- > Folks,
- >
- > Is it possible to both read from and write to the *same*
- > HDF file? In other words, could I read a scientific
- > data set, make some changes to it, and save it back
- > in the same file? If so, would I have to write it
- > back as a different SD, or could I just replace the
- > one I read.

Yes, it is perfectly possible. That's what the /RDWR option of HDF_SD_Start() is for. I've used this feature frequently.

However, keep in mind that the HDF format doesn't allow a very sophisticated management of file space. What I'm about to say doesn't apply to SDSs, but it applies to most of the other things you can write to an HDF file, such as file or SDS attributes. If you write a new value for an existing attribute, the current value will remain in the file, but is simply marked as no longer in use. The new value is added to the file, rather than overwriting the first one. This is invisible if you locate the attribute by name; you always locate the last-written value. However, the earlier values can be retrieved by using low-level HDF routines to walk through the file one component at a time. This is an understanding that I've picked up somewhere along the line of more than a decade of working with HDF files - I can't say that I've ever put this to a direct test. However, I have watched the file size get steadily bigger when the same file attribute gets re-written multiple times.
