Subject: How to speed up large FOR loops?
Posted by Saurav Dhital on Tue, 10 Apr 2012 15:48:19 GMT
View Forum Message <> Reply to Message

Hi,

By necessity, I have to use a long FOR loop (N >~ 300,000). I have made the operations inside the loop efficient enough that the first 40,000 iterations are relatively fast (~30 mins). Then it slows down to a crawl, with the entire process taking >24 hours. I would appreciate any advice how to speed this up.

The operation inside the loop involves matching between two (large) FITS files and assigning the matches to another (previously created) FITS structure. There are no new variable created per iteration, so there should be no memory being used (each loop writes over the variables anyways). I even create small substructures (e.g., > sub_mystruct = mystruct[0:10]) so that I don't have to index the large structures.

Any tip or advice would be much appreciated,
~Saurav

Subject: Re: How to speed up large FOR loops?
Posted by David Fanning on Fri, 20 Apr 2012 16:53:48 GMT
View Forum Message <> Reply to Message

Saurav writes:

> Thanks! I will use th profiler, probably David's as I already use a
> lot of his code.

David doesn't have a profiler. He still uses friggin' Print statements like all the other old fogies I know. :-(

Cheers,

Coyote

--
Coyote, R.N, USMC
Fanning Software Consulting, Inc.
Coyote's Guide to IDL Programming: http://www.idlcoyote.com/
Sepore ma de ni thui. ("Perhaps thou speakest truth.")