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

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