Hello IDL wizards,

I am trying to speed up a routine whose PROFILER looks like this (sorted by total time):

Module Type Count Only(s) Avg.(s) Time(s) Avg.(s)
REBIN (S) 2158 285.788439 0.132432 285.788439 0.132432
MIN (S) 272 39.719054 0.146026 39.719054 0.146026
FILE_SEARCH (S) 4 21.07632 5.26908 21.07632 5.26908
REFORM (S) 2591 12.59025 0.004859 12.59025 0.004859

The heart of the calculation is a
MINARRAY = MIN(BIGARRAY,DIM=1), where
BIGARRAY is [M*N,A,B,C,D] and so
MINARRAY is [A,B,C,D].
M=~10,000
N=~200
A,B,C,D are all <5

In order to get to BIGARRAY, several steps of REBIN are required. And the result is a calculation that is too slow; it takes 6-20 seconds, depending on the particular machine we run it on. My instinct says that this is not a calculation that should be this slow, though I guess I could be wrong.

Note that 1) I don't think memory is an obstacle, we have 16GB of RAM and the routine has peak usage <3 GB (I would know exactly if there was a working MEMTEST for 64bit IDL); 2) Threading is not really an option, as we intend to multiplex this process with 1 job per processor once we get it tuned.

Does the collective wisdom of the newsgroup have any suggestions as to why this routine might be spending so much time REBINning, and how we might speed it up?

In supplication,

--Edward H.