Subject: Why is MEAN so slow?
Posted by Matt Francis on Wed, 19 Jan 2011 00:52:45 GMT
View Forum Message <> Reply to Message

I've been using the PROFILER to track down why some code was a bit slow and found that it was spending most of its time in the MEAN function (and then within that, in the MOMENT function called by MEAN).

I had a look at the code for this and it looked very inefficient. It took just moments to come up with this straight forward version:

function mean_quick,data,nan=nan,double=double

```
if keyword_set(nan) then begin
  indx = where(finite(data),count)
  if count EQ 0 then return, keyword_set(double) ? !values.D_nan : !
values.f_nan
  return,total(data[indx],double=double)/count
endif else begin
  return,total(data,double=double)/n_elements(data)
endelse
```

end

This is very simple minded, but as far as I can tell matches the functionality of the library function MEAN and is at least 5 times faster (for my usage in this code).

Why would the Library function be written so inefficiently?