
Subject: Re: Why is MEAN so slow?

Posted by [Matt Francis](#) on Wed, 19 Jan 2011 02:14:04 GMT

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

Interesting. I've only tested this for 1 dimensional arrays on IDL 7.1, not for matrices (the application that I was trying to speed up only used MEAN on 1D arrays).

Here is how I am comparing the two with a test code:

```
pro test_mean_speed
```

```
    ; Test whether the inbuilt MEAN function is slower than MEAN_QUICK
```

```
    n = 100000
```

```
    data = randomu(seed,100000)
```

```
    indx = where(data GT 0.9,count)
```

```
    data[indx] = !values.f_nan
```

```
    for i=0,10000 do begin
```

```
        ;foo=mean_quick(data,/nan)
```

```
        foo=mean(data,/nan)
```

```
    endfor
```

```
end
```

I ran this once using MEAN and once using MEAN_QUICK. Here are the full results from PROFILER

For MEAN:

Module	Type	Count	Only(s)	Avg.(s)	Time(s)	Avg.
(s)						
ABS	(S)	10001	5.040137	0.000504	5.040137	
0.000504						
ARG_PRESENT	(S)	10001	0.005722	0.000001	0.005722	
0.000001						
FINITE	(S)	10001	3.328896	0.000333	3.328896	
0.000333						
KEYWORD_SET	(S)	30003	0.014247	0.000000	0.014247	
0.000000						
MEAN	(U)	10001	0.025580	0.000003	61.421387	
0.006142						
MOMENT	(U)	20002	32.259783	0.001613	108.917069	

```

0.005445
N_ELEMENTS (S) 20002 0.008905 0.000000 0.008905
0.000000
ON_ERROR (S) 30003 0.012295 0.000000 0.012295
0.000000
RANDOMU (S) 1 0.001750 0.001750 0.001750
0.001750
SIZE (S) 10001 0.011161 0.000001 0.011161
0.000001
SQRT (S) 10001 0.007367 0.000001 0.007367
0.000001
TEST_MEAN_SPEED (U) 1 0.006986 0.006986 61.430769
61.430769
TOTAL (S) 60006 18.553728 0.000309 18.553728
0.000309
WHERE (S) 10002 2.154211 0.000215 2.154211
0.000215

```

For MEAN_QUICK:

```

Module      Type Count  Only(s) Avg.(s)  Time(s) Avg.
(s)
FINITE (S) 10001 3.247247 0.000325 3.247247
0.000325
KEYWORD_SET (S) 10001 0.004842 0.000000 0.004842
0.000000
MEAN_QUICK (U) 10001 1.432719 0.000143 9.902029
0.000990
RANDOMU (S) 1 0.001732 0.001732 0.001732
0.001732
TEST_MEAN_SPEED (U) 1 0.011449 0.011449 9.915830
9.915830
TOTAL (S) 10001 3.133684 0.000313 3.133684
0.000313
WHERE (S) 10002 2.084156 0.000208 2.084156
0.000208

```

According to PROFILER (as well as the obvious difference in how long they ran for), MEAN_QUICK is much faster, for this specific problem.