Subject: Re: mean() function
Posted by Foldy Lajos on Tue, 10 Jan 2006 22:53:22 GMT
View Forum Message <> Reply to Message

Hi,

have you tried mean(y,/double)?

regards,
lajos

On Tue, 10 Jan 2006 biocpu@yahoo.com wrote:

> The following looks very odd. Have any clues?
>
> IDL> y = fltarr(1008879)+35
> IDL> id = where(y ne 35, cc)
> IDL> print, cc
> 0
> ; so y is strictly 35.0 BUT
> IDL> print, mean(y)
> 35.5249
>
> IDL> print, mean(y(0:400000))
> 35.0000
>
> Thanks,
>
>

Subject: Re: mean() function
Posted by K. Bowman on Tue, 10 Jan 2006 23:00:39 GMT
View Forum Message <> Reply to Message

In article <1136932449.216202.42760@g44g2000cwa.googlegroups.com>,
biocpu@yahoo.com wrote:

> y = fltarr(1008879)+35

Looks like roundoff error to me

IDL> y = fltarr(1008879)+35
IDL> print, mean(y)
35.0497
IDL> print, mean(y, /double)
35.000000

```
IDL> y = dblarr(1008879)+35
IDL> print, mean(y)
      35.000000
```

Ken Bowman

Subject: Re: mean() function
Posted by David Fanning on Tue, 10 Jan 2006 23:07:03 GMT
View Forum Message <> Reply to Message

biocpu@yahoo.com writes:

> The following looks very odd. Have any clues?
>
> IDL> y = fltarr(1008879)+35
> IDL> id = where(y ne 35, cc)
> IDL> print, cc
> 0
> ; so y is strictly 35.0 BUT
> IDL> print, mean(y)
> 35.5249
>
> IDL> print, mean(y(0:400000))
> 35.0000

Goodness! It is the 10th already and this is the first
"The Sky is Falling" post of the year. I was beginning
to wonder what was going on! :-)

You might want to have a look at this:

http://www.dfanning.com/math_tips/sky_is_falling.html

Cheers,

David
--
David Fanning, Ph.D.
Fanning Software Consulting, Inc.
Coyote's Guide to IDL Programming: http://www.dfanning.com/

Subject: Re: mean() function
Posted by Paul Van Delst[1] on Tue, 10 Jan 2006 23:09:02 GMT
View Forum Message <> Reply to Message

Fï¿½LDY Lajos wrote:
> Hi,
>
> have you tried mean(y,/double)?

Even regular old single worked for me:

```
IDL> y = fltarr(1008879)+35
IDL> id = where(y ne 35, cc)
IDL> print, cc
           0
IDL> print, mean(y)
% Compiled module: MEAN.
% Compiled module: MOMENT.
      35.0000
IDL> print, mean(y[0:400000])
      35.0000
```

Maybe it's version related? I noticed that biocpu used () rather than [] for array indexing. Maybe an earlier version of IDL had a MEAN() function that didn't use a compensated summation algorithm? The OP mean from 0->400000 that worked suggests that's not the case, but who knows? Anyway....

```
IDL> print, !version
{ x86 linux unix linux 6.0.3 Feb 26 2004      32      64}
```

paulv

>
> regards,
> lajos
>
>
> On Tue, 10 Jan 2006 biocpu@yahoo.com wrote:
>
>
>> The following looks very odd. Have any clues?
>>
>> IDL> y = fltarr(1008879)+35
>> IDL> id = where(y ne 35, cc)
>> IDL> print, cc
>>   0
>> ; so y is strictly 35.0 BUT
>> IDL> print, mean(y)
>>   35.5249
>>
>> IDL> print, mean(y(0:400000))

>>   35.0000
>>
>> Thanks,
>>
>>

--
Paul van Delst
CIMSS @ NOAA/NCEP/EMC

Subject: Re: mean() function
Posted by Paul Van Delst[1] on Tue, 10 Jan 2006 23:13:29 GMT
View Forum Message <> Reply to Message

Kenneth Bowman wrote:
> In article <1136932449.216202.42760@g44g2000cwa.googlegroups.com>,
>   biocpu@yahoo.com wrote:
>
>
>> y = fltarr(1008879)+35
>
>
> Looks like roundoff error to me
>
> IDL> y = fltarr(1008879)+35
> IDL> print, mean(y)
>   35.0497
> IDL> print, mean(y, /double)
>   35.000000
> IDL> y = dblarr(1008879)+35
> IDL> print, mean(y)
>   35.000000

Huh. I don't see this in single precision (see other post). What version of IDL did you use?

paulv

--
Paul van Delst
CIMSS @ NOAA/NCEP/EMC

Subject: Re: mean() function
Posted by biocpu on Wed, 11 Jan 2006 00:18:21 GMT
View Forum Message <> Reply to Message

{ mipseb IRIX unix IRIX 6.0 Jun 27 2003 64 64}

Thanks everybody!

Subject: Re: mean() function
Posted by Kenneth P. Bowman on Wed, 11 Jan 2006 03:22:54 GMT
View Forum Message <> Reply to Message

In article <dq1gv6$v05$1@news.nems.noaa.gov>,
Paul Van Delst <Paul.vanDelst@noaa.gov> wrote:

> Kenneth Bowman wrote:
>> In article <1136932449.216202.42760@g44g2000cwa.googlegroups.com>,
>> biocpu@yahoo.com wrote:
>>
>>
>>> y = fltarr(1008879)+35
>>
>>
>> Looks like roundoff error to me
>>
>> IDL> y = fltarr(1008879)+35
>> IDL> print, mean(y)
>> 35.0497
>> IDL> print, mean(y, /double)
>> 35.000000
>> IDL> y = dblarr(1008879)+35
>> IDL> print, mean(y)
>> 35.000000
>
> Huh. I don't see this in single precision (see other post). What version of
> IDL did you use?
>
> paulv

The version I posted (quoted above) is

{ ppc darwin unix Mac OS X 6.2 Jun 20 2005 32 32}

If I run it on my PowerBook (now sadly obsolete ;-)), which is running
6.1

{ ppc darwin unix Mac OS X 6.1 Jul 14 2004 32 32}

I get exactly what was in the original post (he was running 6.0 on IRIX)

```
IDL> y = fltarr(1008879) + 35
IDL> print, mean(y)
      35.5249
IDL> print, mean(y, /double)
       35.000000
IDL> print, version
```

I suppose math libraries or compilers changed between 6.1 and 6.2.

Once the values you are adding differ by 6-7 orders of magnitude,
precision is completely lost for single precision floats.

```
IDL> print, total(replicate(1.0, 10^7))
  1.00000e+07
IDL> print, total(replicate(1.0, 10^8))
  1.67772e+07
IDL> print, total(replicate(1.0D0, 10^8))
   1.0000000e+08
```

Cheers, Ken

Subject: Re: mean() function
Posted by Nigel Wade on Wed, 11 Jan 2006 09:49:11 GMT
View Forum Message <> Reply to Message

biocpu@yahoo.com wrote:

> { mipseb IRIX unix IRIX 6.0 Jun 27 2003 64 64}
>
>
> Thanks everybody!

```
IDL> print,!version
{ mipseb IRIX unix IRIX 6.0 Jun 27 2003      32      64}
IDL>
IDL>  y = fltarr(1008879)+35
IDL> print, mean(y)
% Compiled module: MEAN.
% Compiled module: MOMENT.
      35.0000
```

The only difference I see is that this binary is the n32 version whereas yours
looks like the 64bit version. Maybe that's the difference.

Do you have the n32 version you can test? I don't have the 64bit version
installed, and I can't locate it for download.

--
Nigel Wade, System Administrator, Space Plasma Physics Group,
University of Leicester, Leicester, LE1 7RH, UK
E-mail : nmw@ion.le.ac.uk
Phone : +44 (0)116 2523548, Fax : +44 (0)116 2523555

Subject: Re: mean() function
Posted by Paolo Grigis on Wed, 11 Jan 2006 14:19:24 GMT
View Forum Message <> Reply to Message

>
> [...]
>
> I suppose math libraries or compilers changed between 6.1 and 6.2.
>
> Once the values you are adding differ by 6-7 orders of magnitude,
> precision is completely lost for single precision floats.
>
> IDL> print, total(replicate(1.0, 10^7))
> 1.00000e+07
> IDL> print, total(replicate(1.0, 10^8))
> 1.67772e+07
> IDL> print, total(replicate(1.0D0, 10^8))
> 1.0000000e+08

I have no problem with that, understanding the argument of finite
precision and all that, but there is one thing which has me puzzling
(all commands issued from the same linux machine)

IDL Version 5.4 (linux x86). (c) 2000, Research Systems, Inc.

IDL> help,total(replicate(1.,1d8))
<Expression> FLOAT = 1.67772e+07
IDL> exit

IDL Version 5.5a (linux x86). (c) 2001, Research Systems, Inc.

IDL> help,total(replicate(1.,1d8))
<Expression> FLOAT = 6.71089e+07
IDL> exit

IDL Version 5.6 (linux x86 m32). (c) 2002, Research Systems, Inc.

IDL> help,total(replicate(1.,1d8))

<Expression>   FLOAT   =   1.00000e+08
IDL> exit

I was thinking that the result of the operation would depend
on the hardware used, but I would have guessed that on the same
machine no difference would be seen between different versions
of IDL, since all of them should represent floats in the same way...

So what's happening here? Different compiler optimizations?

Ciao,
Paolo

>
> Cheers, Ken

Subject: Re: mean() function
Posted by Maarten[1] on Thu, 12 Jan 2006 11:07:20 GMT
View Forum Message <> Reply to Message

And the reason you need that page, is in part because IDL uses the
moment routine described in Numerical Recipes (take total first, divide
later), instead of a proper running average, like the GNU scientific
library does.

However, since looping is slow in IDL, you don't want to implement that
in IDL, so the next best thing is to have that page.

Maarten

Subject: Re: mean() function
Posted by savoie on Thu, 12 Jan 2006 16:49:40 GMT
View Forum Message <> Reply to Message

"Maarten" <maarten.sneep@knmi.nl> writes:

> And the reason you need that page, is in part because IDL uses the
> moment routine described in Numerical Recipes (take total first, divide
> later), instead of a proper running average, like the GNU scientific
> library does.
>
> However, since looping is slow in IDL, you don't want to implement that
> in IDL, so the next best thing is to have that page.

Maarten,

Would you mind explaining this a bit for me? What's a proper running average? And why is it better in general?

Thanks
Matt

--
Matthew Savoie - Scientific Programmer
National Snow and Ice Data Center
(303) 735-0785 http://nsidc.org

Subject: Re: mean() function
Posted by R.Bauer on Sun, 15 Jan 2006 08:58:06 GMT
View Forum Message <> Reply to Message

David Fanning wrote:

> biocpu@yahoo.com writes:
>
>> The following looks very odd. Have any clues?
>>
>> IDL> y = fltarr(1008879)+35
>> IDL> id = where(y ne 35, cc)
>> IDL> print, cc
>> 0
>> ; so y is strictly 35.0 BUT
>> IDL> print, mean(y)
>> 35.5249
>>
>> IDL> print, mean(y(0:400000))
>> 35.0000
>
> Goodness! It is the 10th already and this is the first
> "The Sky is Falling" post of the year. I was beginning
> to wonder what was going on! :-)
>
> You might want to have a look at this:
>
> http://www.dfanning.com/math_tips/sky_is_falling.html
>
> Cheers,
>
> David

mean is not useable if it results in this

```
IDL> print,mean( make_array(500000,val=35,/float) )
      35.0413
IDL> print,mean( make_array(400000,val=35,/float) )
      35.0000
```

I prefer a slower routine if this is right.

no one would accept 1.0 + 1.0 result = 1.5

The article is good to understand why mean should not be resolved this way!

Does one have implemented the gnu's version?

cheers
Reimar
