Subject: Re: fix(4.70*100) is... 469
Posted by Jean H. on Thu, 19 Apr 2007 00:58:08 GMT
View Forum Message <> Reply to Message

>> .... so how can it be the float accuracy problem if the difference
>> between the expected and the real value is 256 times bigger than the
>> float error?
>>
> Careful here. The smallest float provides relative accuracy, meaning
> the difference between exact and stored value X doesn't exceed
> X*(machar()).eps. This is well satisfied here.
>
> Mati Meron | "When you argue with a fool,
> meron@cars.uchicago.edu | chances are he is doing just the same"

I don't understand why one should multiply epsilon by X .... why would
the acceptable difference between expect and real value be a function of
the value? ... X = 100.0 Y = 900.0 .... they both have the same number
of significant digits, so why would the max acceptable difference be
IDL> print, 100.0 * epsilon
  1.19209e-005
IDL> print, 900.0 * epsilon
  0.000107288

Also, if one must really multiply epsilon by X, does it mean that there
is an error on http://www.dfanning.com/math_tips/razoredge.html, at the
last line of the page?

IDL> print,abs(0.9 - (0.6+0.3)) lt (machar()).eps

should it be
IDL> print,abs(0.9 - (0.6+0.3)) lt 0.9 * (machar()).eps

???

I admit to be lost on this issue... and it scares me as I might have to
check/change all my codes!!!!!

Do you have a reference at hand on this?

Thanks,
Jean