
Subject: Re: Floating Underflow/Overflow

Posted by [Paul van Delst](#) on Mon, 15 Oct 2001 15:55:01 GMT

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

Kay wrote:

>
> Hi,
>
> i get Floating Overflow/Underflow error messages during my
> calculations, but the result seems to be correct, can these warnings
> be ignored then?
> I'm calculating a Fermi Distribution (I want a sphere with smooth
> edges and this seemed to be the easiest way) I think, that the results
> get to low for larger radiuses so IDL makes this error message. Is it
> possible to tell IDL to round to zero then or what do i have to do?

First test your code after setting !EXCEPT = 2. this will tell you on what lines of code you are getting the over/underflow. Then you can alter your algorithm to avoid them depending on your needs.

E.g. if the numerical precision is a good enough tolerance level, you can do stuff like

```
; set some tolerance level, with double keyword just in case.  
tolerance = ( machar(double=double) ).eps
```

```
if ( x LT tolerance ) then $  
    result = KEYWORD_SET( double ) ? 0.0d0 : 0.0 $  
else $  
    result = y / x
```

I think a lot of posters will probably say this is overkill, and maybe it is for what you want to do. My recent experience has shown that, in general, one can ignore the underflow errors if:

- 1) you are sure your code will always be executed in the same regime, and
- 2) you don't care if you code crashes and burns, or produces a crappy number every now and again.

My e.g.:

the global forecast model here used to run up to about the, oh, around 2hPa pressure level. At these high altitudes the amount of water vapour is negligible. But it's not negligible compared to numerical precision. I prototyped, in IDL, the forward, then tangent-linear (TL), then adjoint (AD) model of the radiative transfer code used to simulate the satellite measurements in the forecast model - which uses up the FOURTH power of the integrated water vapour amount. To test the adjoint model, you expect agreement between the TL and AD models to numerical precision (in double precision). I was getting huge errors at the top of the atmosphere for some satellite channels in the AD model. The reason was that my test code went up to 0.005hPa where the integrated water vapour amounts are next to nothing and (next to nothing)⁴ in a denominator was catastrophic. I put in a test (similar to the above code) and the errors went

away, no more under/overflow errors. The forecast model doesn't go up as high as I tested the radiative transfer code....yet. But in the future it may (it recently got bumped up to the 0.1hPa level)

So, like I said before, it depends on what you want to do. The operational forecast model always has to run. It can't crash with an error or produce a crappy number unless it's flagged as crappy.

Just my opinion of course, but I treat all fp errors as serious.

paulv

--

Paul van Delst Religious and cultural
CIMSS @ NOAA/NCEP purity is a fundamentalist
Ph: (301)763-8000 x7274 fantasy
Fax:(301)763-8545 V.S.Naipaul
