Subject: Re: How Computers Represent Floats Posted by colinr on Fri, 01 Dec 2000 08:00:00 GMT

View Forum Message <> Reply to Message

```
On Thu, 30 Nov 2000 13:23:41 -0700,
William B. Clodius < wclodius@lanl.gov> wrote:
>
>
> "William B. Clodius" wrote:
>> <snip> IEEE 754 requires that all intermediate calculations
>> be performed a higher precision so
> Ignore the above incomplete sentence. What I originally attempted to
> write was covered later.
```

>>

>> <snip>

>

> Some other surprises.

- > The definition of the IEEE 754 mantisa, an integer with values from
- > 2^n_mant to 2*2^n_mant-1, where n_mant is the number of bits available
- > for the mantisa, is termed a normalized number. This is error prone for
- > very small numbers. IEEE 754 mandates that there be available for such
- > small numbers what are termed denorms where the mantissa is interpreted
- > as an integer from 0 to 2^n_mant, so that accuracy degrades gradually
- > for such values. However, this complicates the implementation of the
- > floating point, so some processors, e.g., the DEC Alpha make this
- > available only in software at a greatly reduced performance.

This sounds like it relates to my most recent problem - generating real input files with IDL for a DEC Alpha fortran program. The fortran program had big problems manipulating small numbers generated by the IDL and I had to pepper the IDL code with WHERE statements to set all very small numbers to zero. Has anyone else seen stuff like this?

Colin Rosenthal Astrophysics Institute University of Oslo