Subject: Re: converting floats to doubles
Posted by Peter Mason on Sun, 22 May 2005 23:01:23 GMT
View Forum Message <> Reply to Message

Benjamin Hornberger wrote:
<...>
> Actually I meant dindgen(n) vs. double(findgen(n)) in the first
> place. I always mix them up ... But I think I got the point.

Hi Benjamin,

A single-precision float (as in findgen()) has a mantissa that's effectively 24 bits. Consequentially it represents every integer in the range [-16777215, 16777215] *exactly*. So if your "n" in "findgen(n)" is smaller than 1677216, it doesn't matter whether you use double(findgen(n)) or dindgen(n).
For bigger "n", you will start getting repeat values if you use double(findgen(n)) as the float mantissa will have run out of bits and the double conversion happens after the fact. Here, dindgen() is better. The double mantissa is much bigger (effectively 53 bits) so dindgen ~should~ only start hitting repeats for "n" >= 2^53. (I can't actually check if the algorithm within dindgen() behaves properly outside the 32-bit unsigned integer range. Dindgen(2LL^32LL) alone would be a 32GB array! :-))

Peter Mason