Subject: Re: performing multiple histograms without loops
Posted by Jeremy Bailin on Thu, 04 Feb 2010 19:21:59 GMT
View Forum Message <> Reply to Message

On Feb 4, 7:05 am, chris <rog...@googlemail.com> wrote:
> Hi,
> maybe you can speed up your vectorized approach by avoding the
> transpose step and by using the SAMPLE keyword within rebin to:
> rebin(ul64lindgen(1,ndatasets)*dataspan,size.... ,/SAMPLE) For large
> vectors (replicate({ : })).(0) may work sometimes faster.
>
> Greets
>
> CR

Indeed! Changing the rebin makes a huge difference... for example,
with ndatasets=10000000, datalen=10, datarange=100, I get a factor of
3 improvement.

10000000 10   100   11.5   9.1

turns into

10000000 10   100   11.6   3.2

I didn't try the full grid of parameters, but if that's consistent
across the board, that brings the vectorized approach up to at least
as good as the loop approach over almost the entire grid, and quite a
bit better over most of it.

-Jeremy.