Subject: Re: IDL sorting
Posted by Karl Schultz on Thu, 18 Oct 2007 18:30:14 GMT
View Forum Message <> Reply to Message

wlandsman <wlandsman@gmail.com> wrote:
> On Oct 18, 8:20 am, Wox <nom...@hotmail.com> wrote:
>
>> Sorting the resulting 64-bit longs would do the trick, wouldn't it?
>>
>> function sort,array
>> array64=ishft(long64(array),32)+lindgen(n_elements(array))
>> return,sort(array64)
>> end
>
> That is clever. My tests on my V6.4 Linux box find that it is
> usually faster than bsort.pro. It is somewhat slower when there are
> only a few duplicate values
>
>
>> Btw, I didn't want to ask this, but
>> why is IDL's sort doing this?

Because it is not a *stable* sort. Stable sorting algorithms preserve
the order of equal keys.

> IDL just uses the sort algorithm of the underlying OS. As far as I
> am aware, the SORT function on Linux boxes *does* preserve the order
> of equal values, but that on Mac and Windows machines does not. I
> would be interested to hear if anyone finds any exceptions to this
> rule.

Are you using this SORT function from the command line? If so, you
are using a shell function or a sort program in your PATH. Someone
probably decided that a stable sort made more sense for people sorting
things from the command line or from shell scripts. Reasonable.

IDL uses the C lib function qsort() which is usually an implmentation of
QuickSort, a good overall sort function for general purpose sorting.
Since IDL has no idea what you are sorting, it is actually a pretty good
choice. However, it is not stable. Speed may be more important to some
people than stability.

>
>> Is there any situation where mixing up
>> the order of equal values has a benefit?
>
> None that I can think of. But if you just want the fastest SORT
> possible, you might not care what happens to the equal values.

Exactly. Or your application may not care about equal values, regardless
of speed issues.

> Actually, I think a good suggestion to ITTVIS would be to add a /
> preserve_equal keyword or something similar to SORT(). This topic
> comes up repeatedly.

Yep, perhaps /STABLE

--
Karl Schultz kws@frii.com
There are 844,739 ways to enjoy a Waffle House hamburger.
