Subject: Re: Array has too many elements?
Posted by JD Smith on Mon, 22 Dec 2003 22:01:30 GMT
View Forum Message <> Reply to Message

On Mon, 22 Dec 2003 14:35:55 -0700, Jonathan Greenberg wrote:

- > I'm working with a fuzzy set classifier I found an article on that
- > subdivides a dataspace into subspaces of smaller and smaller divisions.
- > So, for instance, if you have 2 attributes and 2 divisions per
- > attribute, you have an array that has 2^2 elements. Most datasets
- > really need about 25 subdivisions per attribute, and, say, 6 attributes
- > to start getting good classification, so this problem blows up VERY
- > quickly (e.g. 25^6 elements for the problem I just described). One of
- > the issues is that I need to extract a maximum value from the 25^6
- > array, which, it now looks like, i'll have to do in stages (e.g.
- > subdivide the array into fixed size subsets, check each subset for max
- > value and write to a new array, and then determine the max value for
- > this new array). Doable, but obviously involves rewriting a lot of
- > code.

>

- > This, by the way, is why I brought up the supercomputer thread --
- > manipulating arrays (even if I do get the programming bugs sorted out)
- > of this size will be silly to do with a desktop PC. I hope IDL modifies
- > how they deal with large arrays in the future -- since remote sensing
- > images are getting bigger as the spatial, extent and spectral resolution
- > gets higher, this problem is only going to get worse.

Hmmm, I'm not sure about your desktop PC, but that's only a .9GB array... nothing to get bothered by:

IDL> big=randomu(sd,replicate(25L,6)) & print,size(big,/DIMENSIONS) & t=systime(1) & m=max(big,pos) & print,systime(1)-t & print,m,pos

25 25 25 25 25 25 1.7971461 1.000000 68392329

on a two-year old dual processor Athlon system with 1GB RAM. Give me a fast new pair of processors and 4GB of RAM, and this would be a cake-walk, as long as you're careful to keep extra copies of the array to a minimum. If you're desperate for speed, MAX is easy to write in C: just read in chunks at a time about the size of your memory and loop through.

JD