

---

Subject: Re: Avoiding memory paging with large data-set  
Posted by [sterner](#) on Thu, 23 May 1996 07:00:00 GMT

[View Forum Message](#) <> [Reply to Message](#)

---

David Foster <[foster@bial1.ucsd.edu](mailto:foster@bial1.ucsd.edu)> writes:

> I have a question for those familiar with UNIX systems stuff. We  
> are using SPARC 10s and 20s under Solaris 2.3, and often have to  
> work with very large data-sets. Sometimes the operations go pretty  
> fast, and sometimes they go REALLY slow! Same machine, even the  
> same data, but the time varies. I can tell that the machine is  
> having to page the memory in and out continuously when things  
> slow down, but I don't know why.

One thing to consider is in what order you access the data.  
If you can access along the first dimension of an array it  
may be quicker than along another dimension. For example,  
access a large 2-d array in x instead of y if possible. I found  
a case where I could speed things up significantly by doing a  
transpose on a 2-d array before access the data. I could extract  
rows then instead of columns and even with the added transpose  
time it was faster. You will have to experiment with this.

...

I just tried an example for you:

```
a = lindgen(500,500)
t0=systime(1) & for i=0,499 do z=a(i,*) & print,systime(1)-t0
; 0.35461497 <- seconds.
t0=systime(1) & for i=0,499 do z=a(*,i) & print,systime(1)-t0
; 0.031723022 <- seconds.
```

That was on a loaded system (HP 7/35, two IDL jobs, plus a tar).  
And I don't think the test run was paging.

Ray Sterner                    [sterner@tesla.jhuapl.edu](mailto:sterner@tesla.jhuapl.edu)  
The Johns Hopkins University    North latitude 39.16 degrees.  
Applied Physics Laboratory      West longitude 76.90 degrees.  
Laurel, MD 20723-6099  
WWW Home page: <http://fermi.jhuapl.edu/s1r/people/res/res.html>

---