Subject: Re: Reading columns of binary data
Posted by JD Smith on Tue, 08 Aug 2006 17:44:38 GMT
View Forum Message <> Reply to Message

On Tue, 08 Aug 2006 16:57:41 +0200, Fï¿½LDY Lajos wrote:

>
> On Mon, 7 Aug 2006, JD Smith wrote:
>
>> I'd regard this as yet another example of why IDL needs a fast
>> compiled "side-loop" primitive for generic looping operations which
>> don't need the full conveniences of the interpreter loop (ability to
>> hit Control-C to interrupt, etc.). I should be able to say:
>>
>> a=fltarr(500000,/NO_ZERO)
>> f=0.0
>> for_noblock i=0L,500000L-1L do begin
>> point_lun,un,i*100L
>> readu,un,f
>> a[i]=f
>> endfor
>>
>> and not have it be 100x slower than the equivalent in C.
>>
>> JD
>>
>
> Hi,
>
> I mentioned assoc yesterday, so let's try again:
>
> arr=assoc(un, [0.])
> for i=0L,500000L-1L do a[i]=arr[25l*i]
>
> it will do the point_lun/readu in one step, so it should be a little bit
> faster. But I think a faster loop here would not help much, as disk I/O
> is the limiting factor.

If that were the case then the IDL loop and a similar C loop (offsetting the file pointer and reading 4 bytes) should perform similarly. I haven't done the test (and would be happy to be proven wrong), but my guess is they wouldn't match by 1-2 orders of magnitude.

JD