
Subject: Re: Speeding up I/O

Posted by [R.G. Stockwell](#) on Tue, 16 Mar 2004 23:53:14 GMT

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

"Ken Knapp" <Ken.Knapp@REMOVEnoaa.gov> wrote in message
news:c37rd7\$9s\$1@news.nems.noaa.gov...

```
> Are there any tricks to speed up I/O?
> Situation: I have 1.2 TB of data to process. The following is a typical
> read statement:
> in = {header:bytarr(128),scan:intarr(1250)}
> image = replicate(in,1250)
> openr,1,file
> readu,1,image
> close,1
> Then I process the image and output something. The slow part is in the
> readu [I used profiler to find that readu IS slowing me down, the
> computations are only a fraction of the read time].
> Any suggestions on how to speed it up?
> Tricks to make the computer find files faster?
> Would Fortran/C be faster?
>
> Any help is greatly appreciated.
> -Ken
> --
> ***** to reply remove the _REMOVE_ *****
>
> Ken Knapp    Ken.Knapp@_REMOVE_noaa.gov
> Remote Sensing and Applications Division
> National Climatic Data Center
> 151 Patton Ave
> Asheville, NC 28806
> 828-271-4339 (voice) 828-271-4328 (fax)
>
```

Just curious, does it take longer than about 0.02 seconds to read in the data?
for instance...

```
t = systime(1)
openr,1,file
readu,1,image
close,1
print,'read time', systime(1)-t
```

Cheers,
bob

Subject: Re: Speeding up I/O

Posted by [marc schellens\[1\]](#) on Wed, 17 Mar 2004 03:09:10 GMT

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

Ken Knapp wrote:

- > Are there any tricks to speed up I/O?
- > Situation: I have 1.2 TB of data to process. The following is a typical
- > read statement:
- > in = {header:bytarr(128),scan:intarr(1250)}
- > image = replicate(in,1250)
- > openr,1,file
- > readu,1,image
- > close,1
- > Then I process the image and output something. The slow part is in the
- > readu [I used profiler to find that readu IS slowing me down, the
- > computations are only a fraction of the read time].
- > Any suggestions on how to speed it up?
- > Tricks to make the computer find files faster?
- > Would Fortran/C be faster?
- >
- > Any help is greatly appreciated.
- > -Ken

The fastest you can get would be to read in one array instead of an array of structures:

```
image = intarr( 128/2 + 1250, 1250)
openr,1,file
readu,1,image
close,1
```

Then you have to separate the header yourself and convert it back to byte.

Anyway, this method is almost as fast as it gets (almost like C). About 80% faster than the one using structures on my machine.

HDH,
marc

Subject: Re: Speeding up I/O

Posted by [schaa](#) on Wed, 17 Mar 2004 08:41:32 GMT

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

Ken Knapp wrote:

- > Are there any tricks to speed up I/O?
- > Situation: I have 1.2 TB of data to process. The following is a typical

```
> read statement:
> in = {header:bytarr(128),scan:intarr(1250)}
> image = replicate(in,1250)
> openr,1,file
> readu,1,image
> close,1
```

Hi Ken,

I could speed up reading binary files using the ASSOC-Command, like this

```
thisHeader = Assoc(Lun, Bytarr(hLen),hOffset)
```

```
thisData = Assoc(Lun,BytArr(dLen),dOffset)
```

maybe it'll do for you as well.

Best Regards

-Ralf

Subject: Re: Speeding up I/O

Posted by [Ken Knapp](#) on Wed, 17 Mar 2004 13:19:07 GMT

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

Marc Schellens wrote:

```
>
> The fastest you can get would be to read in one array instead of
> an array of structures:
>
> image = intarr( 128/2 + 1250, 1250)
> openr,1,file
> readu,1,image
> close,1
>
> Then you have to separate the header yourself and convert it
> back to byte.
```

OK. That sounds good, but is there a good/fast way to get the header information back to byte? How about if the header were a mixture of byte and int's? It seems to me I'd need some sort of function that would allow me to read from one variable to another ... kind of like the ReadS routine, but "binary-style"??

```
>
> Anyway, this method is almost as fast as it gets (almost like C).
> About 80% faster than the one using structures on my machine.
>
> HDH,
> marc
>
```

--

***** to reply remove the _REMOVE_ *****

Ken Knapp Ken.Knapp@_REMOVE_noaa.gov
Remote Sensing and Applications Division
National Climatic Data Center
151 Patton Ave
Asheville, NC 28806
828-271-4339 (voice) 828-271-4328 (fax)

Subject: Re: Speeding up I/O
Posted by [Rick Towler](#) on Wed, 17 Mar 2004 16:49:18 GMT
[View Forum Message](#) <> [Reply to Message](#)

"Ken Knapp" <Ken.Knapp@REMOVEnoaa.gov> wrote in message
news:c39jvi\$1@news.nems.noaa.gov...

> Marc Schellens wrote:

>>

>> The fastest you can get would be to read in one array instead of
>> an array of structures:

>>

>> image = intarr(128/2 + 1250, 1250)

>> openr,1,file

>> readu,1,image

>> close,1

>>

>> Then you have to separate the header yourself and convert it
>> back to byte.

>

> OK. That sounds good, but is there a good/fast way to get the header
> information back to byte? How about if the header were a mixture of byte
> and int's? It seems to me I'd need some sort of function that would
> allow me to read from one variable to another ... kind of like the ReadS
> routine, but "binary-style"??

Use the BYTE() function with the offset parameter.

-Rick
