Subject: Re: Large TIFF file question
Posted by bowman on Tue, 15 Jan 2002 22:13:33 GMT
View Forum Message <> Reply to Message

In article <a228o1$4n6$1@news.kodak.com>, "Neil Talsania"
<talsania@kodak.com> wrote:

> So, to make a long story short, my question is how do I open a large TIFF
> file, without exhausting memory?

Ah, the irony. Asking us, the poor downtrodden End Users to provide help
to RSI's corporate overlord.

Sorry ... you'll just have to get used to a certain amount of abuse in
this newsgroup.

If you provide info on the OS and IDL versions you are using, the various
users of said versions can likely tell you about process memory limits,
etc.

Ken Bowman

Subject: Re: Large TIFF file question
Posted by Dick Jackson on Tue, 15 Jan 2002 23:06:25 GMT
View Forum Message <> Reply to Message

"Neil Talsania" <talsania@kodak.com> wrote in message
news:a228o1$4n6$1@news.kodak.com...
> Hi,
> I have what should be a simple question (I hope!). I am trying to run an
> IDL routine that was given to me. The routine has run successfully on
small
> images, but when I try to run it on my 1.5 Gig image it fails on the
memory
> allocation.
>
> Looking at the code, it does the following:
>
> a = float(read_tiff(filename).

I might guess that if you did it in two stages, you'd see something
interesting:

1.
aTemp = read_tiff(filename)

- this should use roughly 1.5 GB if it's an ordinary TIFF file with three

bytes per pixel (RGB)

2.
a = float(a)

- this would convert every byte to a 4-byte float, using roughly 6.0 GB!
(The aTemp can be deleted, of course, and your original wouldn't end up
with this 1.5 GB hanging around.)

Perhaps this is the problem, and you may need to get creative to find a
solution. (subsampling the array for further use?)

Cheers,
--
-Dick

Dick Jackson / dick@d-jackson.com
D-Jackson Software Consulting / http://www.d-jackson.com
Calgary, Alberta, Canada / +1-403-242-7398 / Fax: 241-7392

Subject: Re: Large TIFF file question
Posted by Craig Markwardt on Wed, 16 Jan 2002 04:39:49 GMT
View Forum Message <> Reply to Message

"Dick Jackson" <dick@d-jackson.com> writes:

> "Neil Talsania" <talsania@kodak.com> wrote in message
> news:a228o1$4n6$1@news.kodak.com...
>> Hi,
>> I have what should be a simple question (I hope!). I am trying to run an
>> IDL routine that was given to me. The routine has run successfully on
> small
>> images, but when I try to run it on my 1.5 Gig image it fails on the
> memory
>> allocation.
>>
>> Looking at the code, it does the following:
>>
>> a = float(read_tiff(filename).
>
...
> Perhaps this is the problem, and you may need to get creative to find a
> solution. (subsampling the array for further use?)

Or, how about reading only a portion of the image at a time using the
SUB_RECT keyword. This is a technique known as tiling, and of course
the slightly more difficult part is the logic to stitch together

the tiles at the end.

Craig

--
 ---------- ----------
Craig B. Markwardt, Ph.D. EMAIL: craigmnet@cow.physics.wisc.edu
Astrophysics, IDL, Finance, Derivatives | Remove "net" for better response
 ---------- ----------

Subject: Re: Large TIFF file question
Posted by Martin Downing on Wed, 16 Jan 2002 09:17:04 GMT
View Forum Message <> Reply to Message

"Craig Markwardt" <craigmnet@cow.physics.wisc.edu> wrote in message
news:ond70alxii.fsf@cow.physics.wisc.edu...
> "Dick Jackson" <dick@d-jackson.com> writes:
>
>> "Neil Talsania" <talsania@kodak.com> wrote in message
>> news:a228o1$4n6$1@news.kodak.com...
>>> Hi,
>>> I have what should be a simple question (I hope!). I am trying to
run an
>>> IDL routine that was given to me. The routine has run successfully on
>> small
>>> images, but when I try to run it on my 1.5 Gig image it fails on the
>> memory
>>> allocation.
>>>
>>> Looking at the code, it does the following:
>>>
>>> a = float(read_tiff(filename).
>>
> ...
>> Perhaps this is the problem, and you may need to get creative to find a
>> solution. (subsampling the array for further use?)
>
> Or, how about reading only a portion of the image at a time using the
> SUB_RECT keyword. This is a technique known as tiling, and of course
> the slightly more difficult part is the logic to stitch together
> the tiles at the end.
>
> Craig
>
> --
> ---------- ----------
> Craig B. Markwardt, Ph.D. EMAIL: craigmnet@cow.physics.wisc.edu

> Astrophysics, IDL, Finance, Derivatives | Remove "net" for better response
> ---------- ----------

If you are crazy/unfortunate enough to be doing this on a windows OS, you'll
be facing the 1/2Gb limit on process memory, and anyway no matter how much
memory you have the chances are you will be watching the grass grow as page
faulting takes up most of the time :([I'd be happy for someone to prove
me wrong!]. Craig's method is undoubtedly the way to go.

Martin

Subject: Re: Large TIFF file question
Posted by David Fanning on Wed, 16 Jan 2002 13:24:11 GMT
View Forum Message <> Reply to Message

Martin Downing (martin.downing@ntlworld.com) writes:

> If you are crazy/unfortunate enough to be doing this on a windows OS, you'll
> be facing the 1/2Gb limit on process memory, and anyway no matter how much
> memory you have the chances are you will be watching the grass grow as page
> faulting takes up most of the time :([I'd be happy for someone to prove
> me wrong!]. Craig's method is undoubtedly the way to go.

I thought one of the features of IDL 5.4 or 5.5 (I
can't recall, since I just woke up and I'm sitting
here scratching myself and waiting for the coffee
to boil) was an RSI hack that allowed the PCs to
exceed these memory limits. I remember this as being
one of the most significant, but completely unheralded,
items of that release.

Cheers,

David

--
David W. Fanning, Ph.D.
Fanning Software Consulting
Phone: 970-221-0438, E-mail: david@dfanning.com
Coyote's Guide to IDL Programming: http://www.dfanning.com/
Toll-Free IDL Book Orders: 1-888-461-0155

Subject: Re: Large TIFF file question
Posted by Mark Rivers on Wed, 16 Jan 2002 14:54:28 GMT
View Forum Message <> Reply to Message

David Fanning <david@dfanning.com> wrote in message
news:MPG.16af2ce16a9304e39897d3@news.frii.com...
> Martin Downing (martin.downing@ntlworld.com) writes:
>
>> If you are crazy/unfortunate enough to be doing this on a windows OS,
you'll
>> be facing the 1/2Gb limit on process memory, and anyway no matter how
much
>> memory you have the chances are you will be watching the grass grow as
page
>> faulting takes up most of the time :([I'd be happy for someone to
prove
>> me wrong!]. Craig's method is undoubtedly the way to go.
>
> I thought one of the features of IDL 5.4 or 5.5 (I
> can't recall, since I just woke up and I'm sitting
> here scratching myself and waiting for the coffee
> to boil) was an RSI hack that allowed the PCs to
> exceed these memory limits. I remember this as being
> one of the most significant, but completely unheralded,
> items of that release.

If you find anything documenting that I'd be most interested to hear about it. I routinely bump into this limit on Windows machines with 1GB of RAM, reading 3-D tomography data sets that are 400-600 MB. .RESET_SESSION_ALL sometimes helps, but I have to exit/restart IDL very frequently because the memory gets fragmented.

Question for the group: IDL runs on a number of 64-bit operating systems, e.g. Solaris 8, etc. But my understanding was that IDL on such platforms was still 32 bits, so that, for example, the largest array element IDL could access was still a 32-bit pointer and a 4GB array would be an absolute limit, with 1-2 GB being more typical system-specific limits. Is this true?

The new 64-bit Itanium processors have arrived, and there is a 64-bit version (beta) of Windows to support them. I hope IDL releases a version SOON that can take advantage of the additional memory. Hardware has caught up to software sooner than we all expected.

Mark