
Subject: Processing getting slower and slower?

Posted by [Zoltan Bartalis](#) on Mon, 06 Mar 2006 16:52:12 GMT

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

Hello all

[Regular reader, first time poster]

I am processing satellite data with IDL, involving reading in data, processing it, storing it in several buffer arrays and writing (appending) it to lots of binary files (about one million) as soon as any of the buffers is full. Without going into more details, a more general question: has anyone encountered that while going through several processing cycles involving intensive file opening, writing and closing, the processing gets slower and slower for each cycle, although the operations in each cycle are more or less the same? I am declaring all the buffer arrays and other needed variables `_before_` looping through the read-in of the input data files, so I would expect no change in the memory usage from cycle to cycle. Is this an IDL issue or a platform one? Running under Windows XP.

Thanks in advance,
/Zoltan, Vienna

--

Zoltan Bartalis Institute of Photogrammetry and Remote Sensing
 Vienna University of Technology
 Gusshausstrasse 27-29
 A-1040 Vienna, Austria
email: zb @ ipf . tuwien . ac . at Tel (++43 1) 58801 12240
web: http://www.ipf.tuwien.ac.at/zb Fax (++43 1) 58801 12299

Subject: Re: Processing getting slower and slower?

Posted by [Nigel Wade](#) on Wed, 08 Mar 2006 14:49:30 GMT

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

Zoltan Bartalis wrote:

> Hello all
>
> [Regular reader, first time poster]
>
> I am processing satellite data with IDL, involving reading in data,
> processing it, storing it in several buffer arrays and writing (appending)
> it to lots of binary files (about one million) as soon as any of the buffers
> is full. Without going into more details, a more general question: has

> anyone encountered that while going through several processing cycles
> involving intensive file opening, writing and closing, the processing gets
> slower and slower for each cycle, although the operations in each cycle are
> more or less the same? I am declaring all the buffer arrays and other needed
> variables _before_ looping through the read-in of the input data files, so I
> would expect no change in the memory usage from cycle to cycle. Is this an
> IDL issue or a platform one? Running under Windows XP.
>
> Thanks in advance,
> /Zoltan, Vienna
>

Is there any change in memory usage as the loops progress?

It might be a filesystem issue. As you continually open/append/close the files they grow in size, and each append takes longer as it has to find the end of a progressively longer file. I think this will be made much worse by the fact you are writing many files simultaneously, so each file will become very badly fragmented, and fragmented files take much longer to locate to EOF.

--

Nigel Wade, System Administrator, Space Plasma Physics Group,
University of Leicester, Leicester, LE1 7RH, UK

E-mail : nmw@ion.le.ac.uk

Phone : +44 (0)116 2523548, Fax : +44 (0)116 2523555

Subject: Re: Processing getting slower and slower?
Posted by [Zoltan Bartalis](#) on Thu, 09 Mar 2006 16:59:15 GMT
[View Forum Message](#) <> [Reply to Message](#)

Thanks David, Lajos and Nigel for your answers. The explanation is probably a combination of the issues you mentioned. Unfortunately due to time pressure I couldn't investigate the problem further - I simply adopted a more "unorthodox" solution, I now gave up processing globally and started doing it region-wise. That means more manual interference with the processing, and the total processing time is more or less unchanged, but at least there are now fewer files to handle simultaneously, and as soon as a batch is ready, I can go on working with it.

Thanks for your help again,
-Zoltan-

"Nigel Wade" <nmw@ion.le.ac.uk> wrote in message
news:[dumqsc\\$gsh\\$1@south.jnrs.ja.net](mailto:dumqscgsh1@south.jnrs.ja.net)...
> Zoltan Bartalis wrote:
>

>> Hello all
>>
>> [Regular reader, first time poster]
>>
>> I am processing satellite data with IDL, involving reading in data,
>> processing it, storing it in several buffer arrays and writing
>> (appending)
>> it to lots of binary files (about one million) as soon as any of the
>> buffers
>> is full. Without going into more details, a more general question: has
>> anyone encountered that while going through several processing cycles
>> involving intensive file opening, writing and closing, the processing
>> gets
>> slower and slower for each cycle, although the operations in each cycle
>> are
>> more or less the same? I am declaring all the buffer arrays and other
>> needed
>> variables _before_ looping through the read-in of the input data files,
>> so I
>> would expect no change in the memory usage from cycle to cycle. Is this
>> an
>> IDL issue or a platform one? Running under Windows XP.
>>
>> Thanks in advance,
>> /Zoltan, Vienna
>>
>
> Is there any change in memory usage as the loops progress?
>
> It might be a filesystem issue. As you continually open/append/close the
> files
> they grow in size, and each append takes longer as it has to find the end
> of a
> progressively longer file. I think this will be made much worse by the
> fact you
> are writing many files simultaneously, so each file will become very badly
> fragmented, and fragmented files take much longer to locate to EOF.
>
> --
> Nigel Wade, System Administrator, Space Plasma Physics Group,
> University of Leicester, Leicester, LE1 7RH, UK
> E-mail : nmw@ion.le.ac.uk
> Phone : +44 (0)116 2523548, Fax : +44 (0)116 2523555
