

Martin Moche

Från: Martin Moche
Skickat: den 6 april 2020 09:04
Till: Gerstel, Markus (DLSLtd,RAL,LSCI)
Kopia: dials-support; Sebastian Thorarensen; Filip Polbratt
Ämne: Sv: DIALS or computer issue

Dear Markus,

Thanks for your detailed explanation!

I will test again using perhaps 16 cores out of 32 available on each node at NSC Tetralith.

I will involve staff at the National Supercomputer Centre (NSC) in Linköping and share if they find any alternative solution for avoiding DIALS crashes.

Best regards,
Martin

Från: Gerstel, Markus (DLSLtd,RAL,LSCI) <markus.gerstel@diamond.ac.uk>
Skickat: den 4 april 2020 23:55
Till: Martin Moche <martin.moche@ki.se>
Kopia: dials-support <dials-support@lists.sourceforge.net>
Ämne: Re: DIALS or computer issue

Dear Martin,

thank you for the data. I had a look and you are right that this is a variant of our known running-out-of-memory-issue <https://github.com/dials/dials/issues/659>, however in your case there is a twist. We do have a check in place before we go and try allocating the memory we need, and if we find out that we'd need more memory than we can have we reduce the number of parallel processes to fit into memory, trying to avoid an out-of-memory crash. In your case this logic determined that dials.integrate should use 22 cores, and then it still crashed.

I believe this is because our memory check apparently does the following:

- We calculate the amount of memory required by one process
- We take the available system memory and multiply it by a factor of 0.75 (parameter 'integration.block.max_memory_usage')
- We divide the first second number by the first number, and use the rounded down value of that division to limit the number of processes.

This is probably fine, but it builds on the assumption that *everything else* running on your computer fits inside 'system memory * (1 - integration.block.max_memory_usage) + swap size'.

This is because "available system memory" literally means "theoretically available physical memory in this machine". It does not mean "memory available to allocate". Incidentally it also excludes swap memory.

For your immediate case it is easiest to just reduce the number of processes with the xia2 'nproc' option below 22, depending on the memory situation on your machine.

For us DIALS folks I think we need to do a bit better with the memory estimation. I propose the following logic:

- Find the amount of memory available for allocation ([psutil.virtual_memory\(\).available](#)) and the amount of memory available in swap ([psutil.swap_memory\(\).free](#))
- Subtract the second from the first to obtain the amount of memory available for *immediate* allocation, multiply it by 'integration.block.max_memory_usage'
- Divide this by the amount of memory required by one process, round down and limit the number of processes to this number.
- If this number falls below 1 then
 - Check whether the amount of memory available for immediate allocation *including swap* is greater than the amount of memory required by one process
 - If this is the case, then continue with 1 process
 - Otherwise crash out with the usual "Not enough memory ..." error

Further:

- raise the default of 'integration.block.max_memory_usage' from 0.75 to 0.90 or maybe even higher. Reason for this is that the 'unknown' caused by the rest of the system has mostly gone away.
- write all the obtained numbers into the log. That would be really helpful for debugging...
- and for a bonus: limit amount of 'memory available for immediate allocation' to `resource.getrlimit(resource.RLIMIT_RSS)[0] - psutil.Process().memory_info().rss`, because psutil ignores ulimits.

-Markus

From: Gerstel, Markus (DLSLtd,RAL,LSCI) <markus.gerstel@diamond.ac.uk>
Sent: Friday, April 3, 2020 11:34 AM
To: Martin Moche <martin.moche@ki.se>; xia2.support@gmail.com <xia2.support@gmail.com>
Cc: Eleonore von Castelmur <eleonore.von.castelmur@liu.se>
Subject: RE: DIALS or computer issue

Dear Martin,

Dials failed in the integration step, most likely due to running out of memory.
Could you please send us the file
`/proj/xray/users/x_marmo/targets/eleonore/nat3/dials/DEFAULT/NATIVE/SWEEP1/integrate/13_dials.integrate_INTEGRATE.log` to have a look at what is going on in that step?

-Markus

From: Martin Moche <martin.moche@ki.se>
Sent: 03 April 2020 10:30
To: xia2.support@gmail.com
Cc: Eleonore von Castelmur <eleonore.von.castelmur@liu.se>
Subject: DIALS or computer issue

Dear xia2 and DIALS developers

I fail running xia2 DIALS with latest 2.2.1 version on a high resolution dataset. The same dataset runs ok using XDS.

DIALS 2.2.1 runs ok on lower resolution datasets.

It may be an issue with the computer or perhaps an issue with the software.

Best regards,
Martin

Martin Moche, Ph.D.
Head of Protein Crystallography
Karolinska Institutet
PSF/MBB/Biomedicum
Solnavägen 9
171 65 Stockholm
Sweden
mobile: +46-73-322 93 27
phone: +46-8-524 868 43

--

This e-mail and any attachments may contain confidential, copyright and or privileged material, and are for the use of the intended addressee only. If you are not the intended addressee or an authorised recipient of the addressee please notify us of receipt by returning the e-mail and do not use, copy, retain, distribute or disclose the information in or attached to the e-mail.

Any opinions expressed within this e-mail are those of the individual and not necessarily of Diamond Light Source Ltd.

Diamond Light Source Ltd. cannot guarantee that this e-mail or any attachments are free from viruses and we cannot accept liability for any damage which you may sustain as a result of software viruses which may be transmitted in or with the message.

Diamond Light Source Limited (company no. 4375679). Registered in England and Wales with its registered office at Diamond House, Harwell Science and Innovation Campus, Didcot, Oxfordshire, OX11 0DE, United Kingdom