

GROMACS - Bug #1693

Jenkins Tests seldomly failing

02/22/2015 01:07 AM - Roland Schulz

Status: Closed	
Priority: Normal	
Assignee:	
Category:	
Target version:	
Affected version - extra info:	Difficulty: uncategorized
Affected version: git master	
Description We have multiple tests seldom failing in Jenkins. Some of those are fp-exception related and those are tracked at #1677 . Others which are not are: - MdrunTests failing with out of memory. Example: http://jenkins.gromacs.org/job/Gromacs_Gerrit_master/8143 , http://jenkins.gromacs.org/job/Gromacs_Gerrit_master/8142/ , http://jenkins.gromacs.org/job/Gromacs_Gerrit_master/8057 . Might be only double precision - cudaMallocHost issue. Example: http://jenkins.gromacs.org/job/Gromacs_Gerrit_master/8139 see also sysadmin bug #1585	
Related issues:	
Related to GROMACS - Bug #1677: floating-point exceptions found	Closed 01/22/2015
Related to GROMACS - Bug #1990: LJ-PME unstable with OpenCL	Closed

Associated revisions

Revision 8c1fedf2 - 06/20/2016 08:52 AM - Mark Abraham

Update pre-submit matrix contents

Converted a config so that we have one that uses neither MPI. Also needed an incidental fix for the build script to make that work.

Refs #1693

Change-Id: leebc939f6c9cf1d3a84681ce212e61059053cf55

History

#1 - 02/22/2015 01:07 AM - Roland Schulz

- Related to Bug #1677: floating-point exceptions found added

#3 - 02/22/2015 01:10 AM - Roland Schulz

- Description updated

#4 - 02/22/2015 01:41 AM - Roland Schulz

MdrunTests pass ASAN and MSAN in double precision. We probably need to wait for this to happen again and look at it before Jenkins removes the core file.

#5 - 05/08/2016 12:17 PM - Erik Lindahl

Have we seen any of this the last few months?

#6 - 05/15/2016 12:27 AM - Erik Lindahl

- Status changed from New to Blocked, need info

Doesn't seem like anybody has found anything; we'll wait until the end of the week, but if no more info has been found or posted by then we'll assume it was fixed over the last year and close it.

#7 - 05/15/2016 10:53 AM - Teemu Murtola

I don't remember seeing those failures in quite some time, but recently I've seen random failures with OpenCL builds (one of the LJ-PME regression tests) and with the essential dynamics regression tests.

#8 - 05/17/2016 03:20 PM - Mark Abraham

Yeah we worked around the OpenCL issue by swapping the debug and release config, but nobody has any theories about why it was happening, and I've been through that logic at least twice.

#9 - 05/17/2016 07:51 PM - Teemu Murtola

Mark Abraham wrote:

Yeah we worked around the OpenCL issue by swapping the debug and release config, but nobody has any theories about why it was happening, and I've been through that logic at least twice.

I did see this with the configurations currently in production, so that swapping did not solve all the issues:

[http://jenkins.gromacs.org/job/Gromacs_Gerrit_master_nrwpo/1011/OPTIONS=gcc-5.2%20openmp%20opencl%20amdappsdk-3.0%20host=bs_nix-amd_gpu,label=bs_nix-amd_gpu/testReport/junit/\(root\)/complex/nbnxn_ljpme_LB_geometric/](http://jenkins.gromacs.org/job/Gromacs_Gerrit_master_nrwpo/1011/OPTIONS=gcc-5.2%20openmp%20opencl%20amdappsdk-3.0%20host=bs_nix-amd_gpu,label=bs_nix-amd_gpu/testReport/junit/(root)/complex/nbnxn_ljpme_LB_geometric/)

The essential dynamics failure is here:

http://jenkins.gromacs.org/job/Gromacs_Gerrit_master_nrwpo/1016/OPTIONS=gcc-4.9%20tsan%20ftpack%20simd=avx2_256%20host=bs_nix1310,label=bs_nix1310/console

#10 - 05/17/2016 10:35 PM - Mark Abraham

Thanks, Teemu.

Do you have any guesses about the ED issue, Carsten?

#11 - 05/17/2016 10:46 PM - Roland Schulz

Should we invest into creating a better way of keeping core dumps? It was disabled because it was a issue with disk space. But for seldom failing tests it should help a lot with debugging it. And if we are smart with storage (compressing, keeping the important ones, ...) it should be possible. Ideally we would not just have a core dump but even a record of execution (such as rr or undoDB). That should make debugging those trivial.

#12 - 05/17/2016 11:00 PM - Mark Abraham

Roland Schulz wrote:

Should we invest into creating a better way of keeping core dumps? It was disabled because it was a issue with disk space. But for seldom failing tests it should help a lot with debugging it. And if we are smart with storage (compressing, keeping the important ones, ...) it should be possible. Ideally we would not just have a core dump but even a record of execution (such as rr or undoDB). That should make debugging those trivial.

It shouldn't be too hard to identify and keep gmx + libgromacs + test binaries, compress those and make them artefacts. Collecting all the core files might be harder?

#13 - 05/17/2016 11:07 PM - Roland Schulz

The core files can also be compressed and made an artifact. Why do you think those would be harder? I think we mainly need a smart retention policy. One which doesn't need huge amount of data but also keeps the artifacts long enough until someone has time to look at it. One option would be to ask people to manual flag those builds which have a crash which doesn't seem related to the current commit. An alternative would be to keep those where a retrigger doesn't crash.

#14 - 05/18/2016 12:36 AM - Mark Abraham

Roland Schulz wrote:

The core files can also be compressed and made an artifact. Why do you think those would be harder? I think we mainly need a smart retention policy. One which doesn't need huge amount of data but also keeps the artifacts long enough until someone has time to look at it. One option would be to ask people to manual flag those builds which have a crash which doesn't seem related to the current commit. An alternative would be to keep those where a retrigger doesn't crash.

There's potentially core files in lots of places (e.g. all the regressiontests paths when someone breaks tpxio) and if ever there's a crash from a binary run from the same place, then AFAIK it's not immediately clear how to connect the crash instance with the core file it produces.

#15 - 05/18/2016 12:41 AM - Roland Schulz

The previous solution took all core files in all folders. We might want to put an upper limit on it, in case all/most regressiontests fail. I don't think we run the same binary in the same folder. For the unit tests it is different binaries and for the regressiontests it is different folders. We could set core_pattern so that the core file name contains the binary name so that for unit tests it is obvious which core file is for which binary. But the core file also contains that information inside.

#16 - 05/18/2016 09:13 AM - Carsten Kutzner

Mark Abraham wrote:

Thanks, Teemu.

Do you have any guesses about the ED issue, Carsten?

I remember once seeing a segfault with ED in one of these tests, however I could not reproduce that on my workstation. After a retrigger it was gone again. Can one at least see what is in the mdrun.out file?

Abnormal return value for ' gmx mdrun -ntmpi 2 -ei sam.edi -eo flooding1.xvg >mdrun.out 2>&1' was -1

#17 - 06/09/2016 11:56 AM - Mark Abraham

Roland Schulz wrote:

The previous solution took all core files in all folders. We might want to put an upper limit on it, in case all/most regressiontests fail. I don't think we run the same binary in the same folder. For the unit tests it is different binaries and for the regressiontests it is different folders. We could set core_pattern so that the core file name contains the binary name so that for unit tests it is obvious which core file is for which binary. But the core file also contains that information inside.

regressiontests runs gmx with grompp, check, mdrun, check, check

#18 - 06/09/2016 12:06 PM - Mark Abraham

One OpenCL case continues to fail at complex/nbnxn-ljpme-LB-geometric, so I will disable it. I had an idea that it might be the multiple FFTs causing some assumption about timing to be flawed, but nbnxn-ljpme-LB is the one with the extra FFTs.

Note that both the OpenCL configs have had problems in debug mode (see <https://gerrit.gromacs.org/#/c/5461/>) so for now we will test only release mode

#19 - 06/09/2016 02:12 PM - Gerrit Code Review Bot

Gerrit received a related patchset '1' for Issue [#1693](#).

Uploader: Mark Abraham (mark.j.abraham@gmail.com)

Change-Id: Ieebc939f6c9cf1d3a84681ce212e61059053cf55

Gerrit URL: <https://gerrit.gromacs.org/5941>

#20 - 06/13/2016 01:13 PM - Mark Abraham

- Related to Bug #1990: LJ-PME unstable with OpenCL added

#21 - 09/26/2017 01:06 PM - Mark Abraham

That LJPME test continues to have an issue (

[http://jenkins.gromacs.org/view/Gerrit%20pre-submit/job/Matrix_PreSubmit_master/1150/OPTIONS=gcc-5%20openmp%20simd=avx_128_fma%20openc1%20amdappsdk-3.0%20host=bs_nix-amd_gpu.label=bs_nix-amd_gpu/testReport/\(root\)/complex/nbnxn_ljpme_LB_geometric/](http://jenkins.gromacs.org/view/Gerrit%20pre-submit/job/Matrix_PreSubmit_master/1150/OPTIONS=gcc-5%20openmp%20simd=avx_128_fma%20openc1%20amdappsdk-3.0%20host=bs_nix-amd_gpu.label=bs_nix-amd_gpu/testReport/(root)/complex/nbnxn_ljpme_LB_geometric/)), even though that build is RelWithDebug

#22 - 09/28/2017 04:07 PM - Aleksei lupinov

[http://jenkins.gromacs.org/job/Matrix_PreSubmit_master/1219/OPTIONS=gcc-5%20openmp%20simd=avx_128_fma%20openc1%20amdappsdk-3.0%20host=bs_nix-amd_gpu.label=bs_nix-amd_gpu/testReport/junit/\(root\)/complex/nbnxn_ljpme_LB_geometric/](http://jenkins.gromacs.org/job/Matrix_PreSubmit_master/1219/OPTIONS=gcc-5%20openmp%20simd=avx_128_fma%20openc1%20amdappsdk-3.0%20host=bs_nix-amd_gpu.label=bs_nix-amd_gpu/testReport/junit/(root)/complex/nbnxn_ljpme_LB_geometric/)

It likes my GPU timing changes, definitely. They introduce some asserts, too. I guess it's a rare synchronization issue.

#23 - 10/02/2017 08:43 PM - Aleksei lupinov

[http://jenkins.gromacs.org/job/Matrix_PreSubmit_master/1275/OPTIONS=gcc-5%20openmp%20simd=avx_128_fma%20openc1%20amdappsdk-3.0%20host=bs_nix-amd_gpu.label=bs_nix-amd_gpu/testReport/junit/\(root\)/complex/nbnxn_ljpme_geometric/](http://jenkins.gromacs.org/job/Matrix_PreSubmit_master/1275/OPTIONS=gcc-5%20openmp%20simd=avx_128_fma%20openc1%20amdappsdk-3.0%20host=bs_nix-amd_gpu.label=bs_nix-amd_gpu/testReport/junit/(root)/complex/nbnxn_ljpme_geometric/)

#24 - 10/03/2017 11:57 AM - Mark Abraham

Also found with my HostAllocator change at

[http://jenkins.gromacs.org/job/Matrix_PreSubmit_master/1278/OPTIONS=gcc-5%20openmp%20simd=avx_128_fma%20openc1%20amdappsdk-3.0%20host=bs_nix-amd_gpu.label=bs_nix-amd_gpu/testReport/junit/\(root\)/complex/nbnxn_ljpme_LB_geometric/](http://jenkins.gromacs.org/job/Matrix_PreSubmit_master/1278/OPTIONS=gcc-5%20openmp%20simd=avx_128_fma%20openc1%20amdappsdk-3.0%20host=bs_nix-amd_gpu.label=bs_nix-amd_gpu/testReport/junit/(root)/complex/nbnxn_ljpme_LB_geometric/)

#25 - 10/03/2017 05:07 PM - Aleksei lupinov

Here's nbnxn_vsite failure for a change.

This really seems to behave like a missing sync point.

[http://jenkins.gromacs.org/job/Matrix_PreSubmit_master/1300/OPTIONS=gcc-5%20openmp%20simd=avx_128_fma%20openc!%20amdappsdk-3.0%20host=bs_nix-amd_gpu,label=bs_nix-amd_gpu/testReport/junit/\(root\)/complex/nbnxn_vsite/](http://jenkins.gromacs.org/job/Matrix_PreSubmit_master/1300/OPTIONS=gcc-5%20openmp%20simd=avx_128_fma%20openc!%20amdappsdk-3.0%20host=bs_nix-amd_gpu,label=bs_nix-amd_gpu/testReport/junit/(root)/complex/nbnxn_vsite/)

#26 - 10/03/2017 05:41 PM - Mark Abraham

OK I'm building gcc 7 on the amd_gpu slave to see if I can get a) TSAN, and then b) TSAN+OpenCL capable of testing for a race

#27 - 10/05/2017 08:06 AM - Aleksei lupinov

[http://jenkins.gromacs.org/job/Matrix_PreSubmit_master/1363/OPTIONS=gcc-5%20openmp%20simd=avx_128_fma%20openc!%20amdappsdk-3.0%20host=bs_nix-amd_gpu,label=bs_nix-amd_gpu/testReport/junit/\(root\)/complex/nbnxn_ljpme_LB_geometric/](http://jenkins.gromacs.org/job/Matrix_PreSubmit_master/1363/OPTIONS=gcc-5%20openmp%20simd=avx_128_fma%20openc!%20amdappsdk-3.0%20host=bs_nix-amd_gpu,label=bs_nix-amd_gpu/testReport/junit/(root)/complex/nbnxn_ljpme_LB_geometric/)

[http://jenkins.gromacs.org/job/Matrix_PreSubmit_master/1362/OPTIONS=gcc-5%20openmp%20simd=avx_128_fma%20openc!%20amdappsdk-3.0%20host=bs_nix-amd_gpu,label=bs_nix-amd_gpu/testReport/junit/\(root\)/complex/nbnxn_ljpme_LB/](http://jenkins.gromacs.org/job/Matrix_PreSubmit_master/1362/OPTIONS=gcc-5%20openmp%20simd=avx_128_fma%20openc!%20amdappsdk-3.0%20host=bs_nix-amd_gpu,label=bs_nix-amd_gpu/testReport/junit/(root)/complex/nbnxn_ljpme_LB/)

Those 2 builds also were overlapping.

#28 - 10/05/2017 03:50 PM - Mark Abraham

Aleksei lupinov wrote:

[http://jenkins.gromacs.org/job/Matrix_PreSubmit_master/1363/OPTIONS=gcc-5%20openmp%20simd=avx_128_fma%20openc!%20amdappsdk-3.0%20host=bs_nix-amd_gpu,label=bs_nix-amd_gpu/testReport/junit/\(root\)/complex/nbnxn_ljpme_LB_geometric/](http://jenkins.gromacs.org/job/Matrix_PreSubmit_master/1363/OPTIONS=gcc-5%20openmp%20simd=avx_128_fma%20openc!%20amdappsdk-3.0%20host=bs_nix-amd_gpu,label=bs_nix-amd_gpu/testReport/junit/(root)/complex/nbnxn_ljpme_LB_geometric/)

[http://jenkins.gromacs.org/job/Matrix_PreSubmit_master/1362/OPTIONS=gcc-5%20openmp%20simd=avx_128_fma%20openc!%20amdappsdk-3.0%20host=bs_nix-amd_gpu,label=bs_nix-amd_gpu/testReport/junit/\(root\)/complex/nbnxn_ljpme_LB/](http://jenkins.gromacs.org/job/Matrix_PreSubmit_master/1362/OPTIONS=gcc-5%20openmp%20simd=avx_128_fma%20openc!%20amdappsdk-3.0%20host=bs_nix-amd_gpu,label=bs_nix-amd_gpu/testReport/junit/(root)/complex/nbnxn_ljpme_LB/)

Those 2 builds also were overlapping.

Interesting, thanks.

Meanwhile, I have single-rank master TSAN+OpenCL reporting a race on buffer clearing. I'll wait for the node to be quiet, rerun and upload some logs.

#29 - 10/05/2017 04:02 PM - Mark Abraham

- File *tsan-logs.tgz* added

Logs and TSAN output of two regressiontests running on recent master HEAD d880251df494bf1948588d916db281f2f8fe110c. I haven't analyzed anything yet, but looks like a real issue, rather than OpenMP false positives.

#30 - 10/06/2017 11:18 AM - Mark Abraham

Mark Abraham wrote:

Logs and TSAN output of two regressiontests running on recent master HEAD d880251df494bf1948588d916db281f2f8fe110c. I haven't analyzed anything yet, but looks like a real issue, rather than OpenMP false positives.

The reported races are all to do with alloc/free/memcpy on a thread maintained by the runtime or driver racing with our usual operations. So I presume there's either a bug with some of that infrastructure, or we are misusing it. Since I was only using one GPU, and no jobs were active, we can rule out issues with **using** two GPUs, or multiple executors on the machine. **Having** two GPUs on the machine could be an issue, e.g. <http://support.amd.com/en-us/kb-articles/Pages/OpenCL2-Driver.aspx> reports issues with OpenCL 2.0 support and multiple GPUs. Device 0 is a 2.0-compatible APU, whereas device 1 is a discrete GPU supporting only 1.2. So I will re-run my tests with -gpu_id 1 and maybe we can try removing the discrete GPU or updating the driver.

#31 - 10/06/2017 12:26 PM - Aleksei lupinov

[http://jenkins.gromacs.org/job/Matrix_PreSubmit_master/1404/OPTIONS=gcc-5%20openmp%20simd=avx_128_fma%20openc!%20amdappsdk-3.0%20host=bs_nix-amd_gpu,label=bs_nix-amd_gpu/testReport/junit/\(root\)/complex/nbnxn_ljpme_LB/](http://jenkins.gromacs.org/job/Matrix_PreSubmit_master/1404/OPTIONS=gcc-5%20openmp%20simd=avx_128_fma%20openc!%20amdappsdk-3.0%20host=bs_nix-amd_gpu,label=bs_nix-amd_gpu/testReport/junit/(root)/complex/nbnxn_ljpme_LB/)

#32 - 10/06/2017 01:44 PM - Mark Abraham

Mark Abraham wrote:

So I will re-run my tests with -gpu_id 1 and maybe we can try removing the discrete GPU or updating the driver.

Using -gpu_id 1 did not help, so we need to try a driver update, and then something like shifting the discrete GPU to e.g. the other AMD build slave (and consequent rework of releng slaves config).

#33 - 10/06/2017 01:49 PM - Mark Abraham

Mark Abraham wrote:

Mark Abraham wrote:

So I will re-run my tests with `-gpu_id 1` and maybe we can try removing the discrete GPU or updating the driver.

Using `-gpu_id 1` did not help, so we need to try a driver update, and then something like shifting the discrete GPU to e.g. the other AMD build slave (and consequent rework of releng slaves config).

<https://redmine.gromacs.org/boards/6/topics/823> opened

#34 - 10/06/2017 02:51 PM - Szilárd Páll

Mark Abraham wrote:

<https://redmine.gromacs.org/boards/6/topics/823> opened

See my comments there. TL; DR: we could attempt and upgrade to the 15.201 (or something like that), but perhaps a better, forward-looking approach is to revamp the AMD OpenCL infrastructure and try to shift to the new stack and less ancient GPUs.

#35 - 10/17/2017 01:22 PM - Aleksei lupinov

Pull_constraint failure for a change

[http://jenkins.gromacs.org/job/Matrix_PreSubmit_master/1615/OPTIONS=gcc-5%20openmp%20simd=avx_128_fma%20opencl%20amdappsdk-3.0%20host=bs_nix-amd_gpu.label=bs_nix-amd_gpu/testReport/junit/\(root\)/complex/pull_constraint/](http://jenkins.gromacs.org/job/Matrix_PreSubmit_master/1615/OPTIONS=gcc-5%20openmp%20simd=avx_128_fma%20opencl%20amdappsdk-3.0%20host=bs_nix-amd_gpu.label=bs_nix-amd_gpu/testReport/junit/(root)/complex/pull_constraint/)

#36 - 10/02/2018 09:02 PM - Mark Abraham

- Status changed from *Blocked*, need info to *Closed*

Aleksei did eventually find and fix a race in GPU buffer clearing, around April 2018, which I presume resolves this issue.

Files

tsan-logs.tgz	20.7 KB	10/05/2017	Mark Abraham
---------------	---------	------------	--------------