

GROMACS - Bug #995

Better parallel random number generator

08/18/2012 06:09 AM - Roland Schulz

Status: Closed	
Priority: Normal	
Assignee:	
Category:	
Target version: 5.0	
Affected version - extra info:	Difficulty: uncategorized
Affected version: 5.0-beta1	
Description	
<p>The random number generators currently used for have correlation between parallel processes (as seed the input seed plus the rank is used), and are not reproducible (which makes testing and debugging unnecessary hard). The code doesn't say which algorithm is used, so I'm not sure whether the algorithm itself doesn't have any correlations.</p> <p>http://www.deshawresearch.com/resources_random123.html is small API licensed under BSD with a couple of PRNG which are fast, don't fail any statics tests, and are reproducible. We should consider replacing the current PRNG with one of those. The paper http://dl.acm.org/citation.cfm?doi=2063405 has the details including performance numbers, claiming it is the fastest PRNG which doesn't fail any statistics tests of TestU01 (as far as I know the gold standard for RNG) and additional tests for parallel RNG.</p>	
Related issues:	
Related to GROMACS - Task #1545: test Random123 on unsupported platforms	Closed 07/01/2014

Associated revisions

Revision ef2d19b3 - 02/26/2014 10:45 AM - Roland Schulz

Replace all mdrun rngs with cycle based rng

The stateful random number generator (rng) used previously doesn't produce reproducible results in parallel for sd/bd and doesn't produce reproducible results for continuation for replica exchange. The rng state has been removed from the checkpoint file.

Fixes #995

Change-Id: Id2a5d064cf363c54db3c16a0675cfeba553feaaa

History

#1 - 09/11/2012 09:11 PM - Szilárd Páll

+1 for what you are suggesting!

Have browsed through the paper a few months ago and it seemed to be impressive work.

#2 - 09/11/2012 10:08 PM - Roland Schulz

The question is in which version should it go? Is it a bugfix or feature? The reason we might want to categorize it as a bugfix is that the current approach to seed with seed+mpi_rank doesn't give statically independent random numbers. Whether that is a problem for the algorithms we use, I don't know.

#3 - 10/24/2012 04:02 PM - Alexey Shvetsov

Another idea may be use following approach

1. use fist rng to generate seeds for mpi/openmp processes
2. init mpi/openmp process with that seed

i use that for mc in nsfactor.c <http://redmine.gromacs.org/projects/gromacs/repository/revisions/master/entry/src/tools/nsfactor.c> lines 200-235

#4 - 10/24/2012 05:17 PM - Roland Schulz

As is explained in the paper, this doesn't guarantee to have uncorrelated random numbers.

#5 - 11/10/2012 07:26 PM - Roland Schulz

- Target version set to 5.0

#6 - 01/20/2014 12:39 AM - Mark Abraham

- Affected version set to 5.0-beta1

Sounds good to me. Is someone willing to act on this in the near future?

#7 - 01/20/2014 08:24 PM - Roland Schulz

I suggest we use the Threefry-2×64-20 algorithm.

Alternatives:

- ARS-7: 2.2x faster. Requires hardware support (should be available on all modern CPU) and no support for GPU (not sure this is important). If we choose this we need a fall-back and the regression-tests wouldn't match. Not sure how to deal with that. Otherwise I would recommend ARS-7.
- Threefry-2×64-13: 1.24x faster. Smaller number of rounds. Still crush resistant so it should be fine but the paper recommends to have a few extra rounds of safety margin. And given that the performance difference is small that's probably not a bad idea.
- Threefry-4×64-20: 1.25x higher throughput. Returns 4 8byte random numbers. Given that we usually don't have use for 4 the effective speed is lower.

Further, I suggest to use as the 2 counter inputs, the atomic number and the step number, and as the 2 keys, the seed from the mdp file and a hardcoded one in the code. Any better suggestion?

#8 - 01/20/2014 10:37 PM - Gerrit Code Review Bot

Gerrit received a related patchset '1' for Issue [#995](#).

Uploader: Roland Schulz (roland@rschulz.eu)

Change-Id: Id2a5d064cf363c54db3c16a0675cfeba553feea

Gerrit URL: <https://gerrit.gromacs.org/3010>

#9 - 02/26/2014 11:00 AM - Roland Schulz

- Status changed from New to Resolved

- % Done changed from 0 to 100

Applied in changeset [ef2d19b37cb9978512f3ce8f0a6385cc44244d4d](#).

#10 - 02/26/2014 11:46 AM - Mark Abraham

Lulz @ <http://dilbert.com/strips/comic/2001-10-25/>

#11 - 04/03/2014 07:47 PM - Roland Schulz

- Status changed from Resolved to Closed

#12 - 07/01/2014 03:00 PM - Szilárd Páll

- Related to Task #1545: test Random123 on unsupported platforms added