

GROMACS - Task #1963

collect examples of systems where users tried to do something that performed badly

05/17/2016 11:10 PM - Mark Abraham

Status:	New
Priority:	Low
Assignee:	Mark Abraham
Category:	
Target version:	
Difficulty:	uncategorized
Description	
This task aims to collect examples of inputs that users thought were reasonable, but whose performance was unexpectedly poor. These may be useful targets for performance optimization (or better documentation/warnings), or benchmark sets. I suspect there are other examples in Redmine (perhaps even from Chris Neale) which I can cross-reference here when I can find them.	

History

#1 - 05/17/2016 11:21 PM - Roland Schulz

For now you mean manual right? I think it would be worth considering whether we want to add automatic data collection. Of course usage tracking has a bad reputation and we would need to make sure we do it with well understood opt-in. But I think both for funding as well as for the goals laid out in your task description it could be very valuable.

#2 - 05/17/2016 11:27 PM - Mark Abraham

Yeah, manual for now. I don't like the idea of automatic collection, and frankly it's not all that clear how we'd decide what to collect. But if we knew how to decide what to collect, and how to get an explicit "opt in," I'm open to alternatives.

#3 - 05/18/2016 02:57 AM - Szilárd Páll

A separate tool that when manually invoked can collect data (log, input files) and submit it to us could work (something like some Linux distros' bug reporting tool). There would only need to be a simple collection server on the other end. I guess sorting things out might end up being the most tedious work.

#4 - 05/18/2016 09:13 PM - Chris Neale

- File *manyPullGroupsSlowsMdrun.tgz* added

Many pull groups slow down mdrun in gromacs5.1.2 (with or without GPU) and 2016-beta1 (with GPU not tested)

More information can be found here:

https://mailman-1.sys.kth.se/pipermail/gromacs.org_gmx-users/2016-May/105764.html

tarball includes all files needed to run the testing and includes the output from a short run with 128 separate pull restraints. In this case, I get 8.6 ns/day vs. 55.1 ns/day without using any pull groups (whereas adding a single pull restraint only slows the sampling down to 51.3 ns/day)

#5 - 08/09/2016 07:30 AM - Chris Neale

Large free energy decoupling groups slow down gromacs enormously. Even groups of modest size compared with the system size really slow things down too much to be useful for things like REST. Files to test this can be found in issue 742 (<https://redmine.gromacs.org/issues/742>) named as RESTandREMDtps.tgz . The plumed/gromacs solution for REST works well, so it is unclear that the REST use-case is strong enough to support modifications of gromacs, but I am cross-posting for completeness since I think this thread is a fantastic idea.

Files

manyPullGroupsSlowsMdrun.tgz	2.92 MB	05/18/2016	Chris Neale
------------------------------	---------	------------	-------------