

Benchmark suite - Feature #1105

produce a benchmark suite

01/07/2013 02:57 PM - Mark Abraham

Status:	New	
Priority:	Normal	
Assignee:		
Category:		
Target version:		
Description		
<p>For 4.6 we definitely need a set of benchmark systems. We can't anticipate all kinds of simulations, hardware and compilers, so users have to participate in the configuring and running of GROMACS if they want near-maximum performance. We need to give them some idea what they might be able to achieve.</p> <p>At least initially, we just need a few systems:</p> <ol style="list-style-type: none">1) something with membrane, protein and water2) something that can contrast the performance of group vs verlet3) something people can use for scaling tests4) something suitable for GPUs5) some small peptide system for quick tests on little hardware of whether some code change that doesn't affect parallelization does affect performance6) something that has been used as a benchmark in the past for GROMACS 4 or other MD codes7) something with Martini? <p>I'm hoping Szilard's 134k atom system can address points 1-4. I can do point 5. Berk, Erik, what's lying around for point 6? Or what's been/being used for benchmarks for the papers people are writing? Who can we hit up for point 7? Anything else we should provide?</p> <p>I think we should provide all necessary input intended for grompp in 4.6. That way, if a user needs to tweak something to test a particular context, they can do so.</p> <p>I guess we should create a benchmark git repo, with a subfolder for each test. Each subfolder has all the files it needs to do grompp and mdrun, a few sentences of README, and two one-line shell scripts "run-grompp.sh" and "run-mdrun.sh". Users are going to have to write their own glue for their queuing systems, MPI config, whatever.</p> <p>I dunno if it's easy/possible to set up some kind of dashboard that we/users can submit results to. It'd be very easy for people to submit crap that we wouldn't want to look "official". Probably best avoided, given that we're going to want to eyeball all the benchmark numbers we publish.</p> <p>Somebody please make me a manager of this Redmine project!</p>		
Related issues:		
Related to Benchmark suite - Feature #1106: publish benchmark numbers	New	01/07/2013

History

#1 - 01/07/2013 04:34 PM - Erik Lindahl

- Assignee set to Mark Abraham

- Target version changed from 4.6 to future

Changing target to "future". This doesn't mean it will wait for 5.0, but it's not something we'll have to hold the 4.6 release for; we already have a set of different benchmarks systems separated into small/medium/large/verylarge cases that I think we could build this on.

#2 - 01/08/2013 12:48 PM - Rossen Apostolov

Mark Abraham wrote:

Somebody please make me a manager of this Redmine project!

Done. I also added you to the other sub-projects.

#3 - 01/09/2013 01:33 PM - Mark Abraham

Erik Lindahl wrote:

Changing target to "future". This doesn't mean it will wait for 5.0, but it's not something we'll have to hold the 4.6 release for; we already have a set of different benchmarks systems separated into small/medium/large/verylarge cases that I think we could build this on.

OK, but at some point soon after 4.6 we need to get serious here. Configuring and running 4.6 is no longer easy, and all our hard work on performance is wasted if people who want to take advantage of that work can't find out reliably whether they are doing so.

#4 - 05/22/2013 05:52 AM - Mark Abraham

- Target version changed from future to 4.6.x

#5 - 06/19/2014 12:07 PM - Rossen Apostolov

- Assignee deleted (Mark Abraham)

- Target version deleted (4.6.x)

#6 - 06/19/2014 12:08 PM - Rossen Apostolov

- Related to Feature #1106: publish benchmark numbers added

#7 - 09/30/2014 06:47 PM - Mark Abraham

NB there is a large set of input files we already have. Curating it into something that can run automatically needs someone to dedicate some hardware (doable), decide on a test harness (needs someone to volunteer - not me any time soon), and work out how to present results.

#8 - 10/01/2014 07:47 AM - Alexey Shvetsov

Also it will be good to have some system containing Protein and {R,D}NA.

PS also for dedicated hw for running such tests it seems to me that it may be possible to use few nodes from systems @SpbSTU SCC after it start working with new system (Intel Haswell nodes, and Xeon Phi nodes)

#9 - 12/11/2014 10:51 PM - Roland Schulz

Any reason why the dppc which was in the old gmxbench isn't in the new benchmarks? Currently the largest "large" input is ~140k (kv12) and the smallest "verylarge" is ~1M (virus_capsid). Having the dppc (~490k) would fill in this rather large gap.

#10 - 12/11/2014 11:30 PM - Roland Schulz

For the virus_capsid it would be nice if a version without virtual sites existed.

#11 - 12/12/2014 05:46 AM - Alexey Shvetsov

I can contribute full nucleosome system to benchmark (~700k atoms with large number of charges).

#12 - 02/16/2015 05:23 AM - Roland Schulz

Ideally the benchmark could write the results in JMeter xml format. Then we could use the Jenkins [Performance Plugin](#) to visualize the performance trend and one could quickly see performance regression. It might also be nice if it contained an xml for running it as a [OpenBenchmarking](#) test suite. This way Gromacs could quickly test the performance of different hardware without having to know anything about Gromacs.

#13 - 05/27/2016 06:54 PM - Roland Schulz

It would be very nice to have this released. What are the required TODOs before it can be released?

#14 - 09/13/2016 10:31 PM - Roland Schulz

Can we make the benchmarks git repo public as is? If not why not?
And put it on gerrit so that we can use gerrit for suggested changes?

I submitted a patch to git which converts a few benchmarks from group to verlet. I assume that we don't need/want any group benchmarks anymore. For the lysozyme I also changed the cut-off because the previous one (rcoulomb=0.9,rvdw=1.4) isn't supported with verlet (new rcoulomb=rvdw=1.2). To keep the same PME accuracy I changed fourier_spacing from 0.12 to 0.15. Please let me know if any of those changes are not desired.