

GROMACS - Bug #1777

Teach mdrun about explicit -append

07/14/2015 11:20 AM - Mark Abraham

Status:	Closed	
Priority:	Low	
Assignee:		
Category:	mdrun	
Target version:	2016	
Affected version - extra info:	probably all versions since 4.0	Difficulty: uncategorized
Affected version:	5.0.4	

Description

In all cases below, I run

```
gmX mdrun -deffnm first -nsteps 100
gmX convert-tpr -s first -extend 1 -o second
```

as the sole preparation for mdrun in an empty directory (but I repeat those commands for clarity):

1)

```
gmX mdrun -deffnm first -nsteps 100
gmX convert-tpr -s first -extend 1 -o second
gmX mdrun -s second -deffnm second
```

writes second.* output files starting from step 0, which is fine

2)

```
gmX mdrun -deffnm first -nsteps 100
gmX convert-tpr -s first -extend 1 -o second
gmX mdrun -s second -deffnm second -append
```

does the same as 1), which is not necessarily fine. Presumably this is because -append is the default and we haven't taught mdrun how to know that the user explicitly asked for appending. mdrun can't tell whether the use of -deffnm or -append (or missing files) is the error, so we should kick that decision back to the user.

3)

```
gmX mdrun -deffnm first -nsteps 100
gmX convert-tpr -s first -extend 1 -o second
gmX mdrun -s second -deffnm second -cpi first
```

writes second.* output files, starting from step 100, which is fine

4)

```
gmX mdrun -deffnm first -nsteps 100
gmX convert-tpr -s first -extend 1 -o second
gmX mdrun -s second -deffnm second -cpi first -append
```

does the same as 3), which is not necessarily fine. Similarly, mdrun can't tell whether the use of -deffnm, -cpi or -append (or missing files) is the error, so we should kick that decision back to the user.

5)

```
gmx mdrun -deffnm first -nsteps 100
gmx convert-tpr -s first -extend 1 -o second
gmx mdrun -s second -append
```

writes output files with default names, starting from step 0, again not fine because mdrun doesn't know where the error happened

6)

```
gmx mdrun -deffnm first -nsteps 100
gmx convert-tpr -s first -extend 1 -o second
gmx mdrun -s second -cpi first -append
```

writes output files with default names, starting from step 100, again not fine because mdrun doesn't know where the error happened

At https://mailman-1.sys.kth.se/pipermail/gromacs.org_gmx-users/2015-July/099153.html, a new user tried to do 6), and expected the output to be appended to first.*.

7)

```
gmx mdrun -deffnm first -nsteps 100
gmx convert-tpr -s first -extend 1 -o second
gmx mdrun -s second -deffnm second -cpi first -nsteps 100
gmx mdrun -s second -deffnm second -cpi second -append
```

starts from step 200 and appends to the existing second.* files, which is fine

8)

```
gmx mdrun -deffnm first -nsteps 100
gmx convert-tpr -s first -extend 1 -o second
gmx mdrun -s second -deffnm second -cpi first -nsteps 100
gmx mdrun -s second -deffnm second -append
```

starts from step 0 and replaces the existing second.* files, but should give an error. That error should be different depending whether second.cpt exists

9)

```
gmx mdrun -deffnm first -nsteps 100
gmx convert-tpr -s first -extend 1 -o second
gmx mdrun -s second -deffnm second -cpi first -nsteps 100
gmx mdrun -s second -deffnm second -cpi first -append
```

starts from step 100 and replaces the existing second.* files, but should give an error

10)

```
gmx mdrun -deffnm first -nsteps 100
gmx convert-tpr -s first -extend 1 -o second
gmx mdrun -s second -deffnm second -cpi first -nsteps 100
gmx mdrun -s second -deffnm second -cpi second -nsteps 100
gmx mdrun -s second -deffnm second -cpi first -append
```

starts from step 100 and replaces the existing second.* files, but should give an error

<https://gerrit.gromacs.org/#/c/4439/> is related, but IIRC doesn't fix any of this

Related issues:

Related to GROMACS - Bug #1889: mdrun -cpi file presence dilemma

Rejected

Related to GROMACS - Task #1781: re-design benchmarking functionality

Accepted

Related to GROMACS - Task #2169: remove 'continuation' mdp option

New

Associated revisions

Revision 07d12095 - 07/09/2016 05:40 PM - Erik Lindahl

Prevent fragile use cases of checkpoint appending

There are way too many ways we allow runs to be continued and extended. We still allow the checkpoint file to be missing (so `-cpi` can be used for all command lines), but we warn if it is not found. To avoid mistakes with file appending when restarting from checkpoints, we now require that all previous output files must be present (unless `-noappend` is used), and that the file names must match the ones used in the previous run.

Fixes #1777.

Change-Id: Id9e89773a4a9214be6dbb76676c526e98e12bd37

History

#1 - 07/14/2015 11:38 AM - Mark Abraham

Some ambiguity here is resolved if we remove the ability for the user to give arbitrary names to `mdrun` output files. `-deffnm` is useful and has been around for a while, but is not a great name if that's the only functionality available. I would like something like `gmx mdrun -name`, perhaps with `-deffnm` also supported for backward compatibility. Similarly, I think we should drop `-multi` in favour of `-multidir`, not least because I expect that the latter is easier to implement and test in combination with `gmx mdrun -name`

#2 - 04/03/2016 10:47 PM - Erik Lindahl

I would strongly recommend that we do NOT allow appending in any other case than continuing from a checkpoint.

There is a history of stupid errors just because we have had umpteen different ways to continue a run, which leads to way too many combinations to test... and then there are bugs.

So, in short: The `-append` flag is only meant to continue from checkpoint, not a convenience flag to concatenate completed trajectories on-the-fly.

The reason we originally added it was originally that we wanted to play it safe, and we might still want it for the extreme corner case of file systems that don't support appending, but nothing else.

#3 - 05/23/2016 12:45 AM - Erik Lindahl

- Status changed from New to Accepted

Based on discussions with Mark last Friday, we both decided it's better to limit append to checkpoint continuation, as it was originally designed for. By default it is always on, but we do allow users to force separate files instead.

However, it was never meant as a convenience replacement for being able to concatenate any independent trajectories already at runtime, and should not be used as such.

#4 - 05/23/2016 06:01 PM - Mark Abraham

Agreed. That decision might break some existing workflows, but we should apologise for past errors and focus on building tools that work provably well for the important mainstream cases.

#5 - 05/23/2016 06:08 PM - Mark Abraham

- Related to Bug #1889: mdrun -cpi file presence dilemma added

#6 - 05/25/2016 12:03 AM - Erik Lindahl

Issuing a fatal error when `-append` is used without checkpointing is easy, but not sufficient.

It seems like most of these cases are caused by using `convert-tptr` to alter runs, and/or the `-nsteps` option to `mdrun`. The whole idea of arbitrarily altering input files, renaming them, and overriding settings on the command line - but still being able to continue the same trajectory - opens a horrible can of worms that is quite clearly WAY too complex to handle correctly.

I would vote for

- 1) Removing the `-extend` option to `convert-tptr`, and limit it to a tool used for altering `tptr` files for analysis.
- 2) Replacing the `-nsteps` option with a `-maxsteps` option. We will NOT write more data than what was requested in the input, but it is possible to stop earlier.

When a user want a particular number of steps we have a clearly prescribed way to achieve that: The "nsteps" option in the `mdp` file. Let's stick to that and avoid screwing things up.

#7 - 05/25/2016 12:18 AM - Erik Lindahl

Here's a proposed solution for the more immediate problem - in combination these should solve all the issues listed:

- 1) Refuse an explicit -append without checkpoint.
- 2) Refuse to run when we cannot find the files to append to (unless -noappend was used)
- 3) Store the names of the original files in the checkpoint, and refuse to append to anything else.

#8 - 05/25/2016 01:08 AM - Gerrit Code Review Bot

Gerrit received a related patchset '1' for Issue [#1777](#).

Uploader: Erik Lindahl (erik.lindahl@gmail.com)

Change-Id: Id9e89773a4a9214be6dbb76676c526e98e12bd37

Gerrit URL: <https://gerrit.gromacs.org/5890>

#9 - 05/25/2016 01:09 AM - Erik Lindahl

- Status changed from Accepted to Fix uploaded

#10 - 07/09/2016 09:51 PM - Erik Lindahl

- Status changed from Fix uploaded to Resolved

Applied in changeset [07d120957728dd271752128113b58b45f6b8e194](#).

#11 - 07/09/2016 09:55 PM - Mark Abraham

- Target version changed from future to 2016

#12 - 07/11/2016 05:09 PM - Mark Abraham

- Status changed from Resolved to Closed

#13 - 03/21/2017 09:48 AM - Mark Abraham

- Related to Task #1781: re-design benchmarking functionality added

#14 - 05/05/2017 04:11 PM - Mark Abraham

- Related to Task #2169: remove 'continuation' mdp option added