

## GROMACS - Feature #1182

### improve trajectory writing to support parallel I/O

03/07/2013 04:40 PM - Mark Abraham

<b>Status:</b>	New	
<b>Priority:</b>	Normal	
<b>Assignee:</b>		
<b>Category:</b>		
<b>Target version:</b>	future	
<b>Difficulty:</b>	uncategorized	
<b>Description</b>		
Up to and including 4.6, we collect the whole [xvf] arrays on SIMMASTER. That's poor, and will become limiting on big systems on HPC platforms that might have low memory (e.g. BlueGene/L had 500MB/core when it came along; P and Q have rather more right now). The design for 5.0 should anticipate solving that (e.g. MPI file I/O), even if we don't actually get it done.		
<b>Related issues:</b>		
Related to GROMACS - Task #1170: mdlib reorganization		<b>New</b>

### History

#### #1 - 03/10/2013 08:28 AM - Roland Schulz

A WIP implementation for MPI-IO is available at <https://github.com/rolandschulz/gromacs/commits/CollectiveIO>. It doesn't solve the issue of memory usage, because it still collects a whole frame to one core. But it allows one to use n IO cores. n frames are locally stored on the compute nodes and when all n frames are available, they are collected in parallel to the n IO-nodes. The IO nodes then write those n frames in parallel to the XTC file. Thus it improves performance if the collection or IO takes significant time for very large number of cores.

Another extremely simple improvement is to use Gather instead of Gatherv for the collection of atoms for writing. E.g. on Cray Gatherv isn't optimized and it is much faster to use Gather even if it means that one is communicating a little bit extra (for those nodes having slightly fewer atoms). To include this into the official version we would need to decide whether 1) we always want to use Gather instead of Gatherv or 2) how to switch between Gather/Gatherv.

#### #2 - 05/13/2014 10:42 AM - Mark Abraham

- Target version changed from 5.0 to future