Supporting information to:

“Parallelization of CPPTRAJ Enables Large Scale Analysis of Molecular Dynamics Trajectory Data”

Submitted to Journal of Computational Chemistry, April 2018.

Table of Contents

1. Detailed Timings for Across-trajectory Parallel Hydrogen Bond Analysis

Run on LoBoS cluster at NIH, Haswell E5-2630 v3 @ 2.4GHz, 2x8 core, Infiniband FDR, NFS mount.

10 runs each PPN.

Command: `hbond HB out nhb.dat avgout hbavg.dat series uuseries UUseries.dat`

375 acceptor-only atoms, 371 acceptor/donor sites, 440 solute hydrogens.

Input Trajectories: 49209 atoms, 9 NetCDF trajectories (2000 frames each), 10.137 GB total.

Data Sets: 1127 integer data sets, 18000 frames (77MB total), hbond data set (12MB), 89MB total.

Data Files: hbavg.dat 116KB, nhb.dat 387KB, UUseries.dat 506MB, 506.318MB total.

PPN	Total	SD	Trajin	SD	Action	SD	Data	SD	Sync	SD
1	22.62	0.86	8.27	1.12	2.35	0.06	9.14	0.06	0.08	0.01
2	23.11	0.36	9.66	0.45	2.41	0.08	9.08	0.14	0.06	0.01
4	33.43	0.24	20.16	0.32	2.56	0.07	9.06	0.10	0.06	0.01
8	56.29	0.26	42.56	0.52	2.68	0.05	9.36	0.05	0.06	0.01
16	102.16	0.14	86.83	1.39	2.70	0.14	9.79	0.10	0.07	0.01

2. Action command lists with brief descriptions

Table S 1: List of all Action commands available in CPPTRAJ 16+ (OpenMP X=enabled).

Action Command	Description	OpenMP
angle	Calculate an angle.	
areapermol	Basic area per molecule calculation in a designated plane.	
atomiccorr	Calculate atomic movement correlations.	X
atomicfluct	Calculate atomic fluctuations.	
atommap	Determine mapping between two reference structures with potentially different atom ordering.	
autoimage	Automatically re-image trajectory.	
average	Calculate average structure.	
bounds	Calculate X Y and Z bounds of selected atoms.	
box	Set box information for input frames.	
center	Center coordinates to a location.	
check	Check input frames for bad overlaps/bond lengths.	X
checkchirality	Check protein residue chirality (D or L).	
checkoverlap	See 'check'.	
checkstructure	See 'check'.	
closest	Retain only specified number of closest solvent molecules to a selection.	X
closestwaters	See 'closest'.	
clusterdihedral	Cluster structures by binning torsions.	
contacts	Count atomic contacts based on distance cutoff.	
createcrd	Create a COORDS data set from input frames.	
createreservoir	Create a NetCDF reservoir for use with reservoir REMD.	
density	Calculate system density.	
diffusion	Calculate diffusion from distance traveled using Einstein relation.	
dihedral	Calculate dihedral angle.	
dihedralscan	Rotate dihedral(s) by a fixed interval or randomly.	
dipole	Calculate dipoles.	
distance	Calculate distances.	
drms	Calculate distance matrix RMS (aka DME).	
drmsd	See 'drms'.	
dssp	Calculate secondary structure using DSSP algorithm.	X
energy	Calculate basic energy terms.	
filter	Filter frames using values from data sets.	
fixatomorder	Correct topologies/frames with non-consecutive atom ordering.	
gist	Grid inhomogenous solvation theory.	X
grid	Bin coordinates in 3D.	
hbond	Calculate hydrogen bonds.	X
image	Image trajectory.	
jcoupling	Calculate J-coupling values using torsions and the Karplus relation.	
lessplit	Split an LES trajectory.	
lie	Calculate linear interaction energy.	

lipidorder	Compute lipid order parameters.	
makestructure	Manipulate structure by rotating dihedrals to specified values or to match a reference structure.	
mask	Select atoms, intended for use with distance-based masks (updates every frame).	X
matrix	Calculate various matrices.	X (covariance)
minimage	Calculate distance to closest adjacent unit cell.	X
molsurf	Calculate Connolly surface area.	
multidihedral	Calculate multiple dihedral angles.	
multivector	Calculate multiple vectors.	
nastruct	Nucleic acid structure analysis following 3DNA procedure.	
nativecontacts	Calculate native and non-native contacts based on distance cutoff and reference structure.	
outtraj	Write out trajectory in line with other Actions (as opposed to the end of all Actions like 'trajout').	
pairdist	Pair distribution function.	
pairwise	Calculate electrostatic and van der Waals energy, potentially comparing to a reference structure.	
principal	Calculate principal axes of a system and optionally align along them.	
projection	Calculate projection of coordinates along eigenvectors.	
pucker	Calculate ring pucker.	
radgyr	Calculate radius of gyration.	
radial	Calculate radial distribution function	X
randomizeions	Randomize ion position by swapping with solvent.	
replicatecell	Replicate unit cell in specified direction(s).	X
rms	Calculate coordinate root-mean-square deviation to a reference structure.	
rmsd	See 'rms'	
rog	See 'radgyr'	
rotate	Rotate system around X, Y, and/or Z axes.	
runavg	Calculate running average of coordinates.	
runningaverage	See 'runavg'	
scale	Scale X, Y, and or Z coordinates by a factor.	
secstruct	See 'dssp'	
setvelocity	Set coordinate velocities.	
spam	Simple approach for profiling bound water molecules.	X
stfcdiffusion	Alternate diffusion calculation from distance travelled.	
strip	Remove atoms from system.	
surf	Calculate LCPO surface area.	X
symmrmsd	Calculate symmetry-corrected coordinate root-mean-square deviation to a reference structure.	
temperature	Calculate temperature from velocity information.	
trans	Translate atoms in X, Y, and/or Z direction.	
translate	See 'trans'.	
unstrip	Reset any modifications to topology/coordinates (from e.g. strip, closest, etc).	
unwrap	Remove imaging from trajectory.	

vector	Calculate various vectors.	
velocityautocorr	Calculate velocity autocorrelation.	X
volmap	Create volumetric map (3D grid).	X
volume	Calculate unit cell volume.	
watershell	Calculate number of waters in first and second solvation shells using distance cutoffs.	X

3. Illustration of Single Replica Processing vs Ensemble Processing

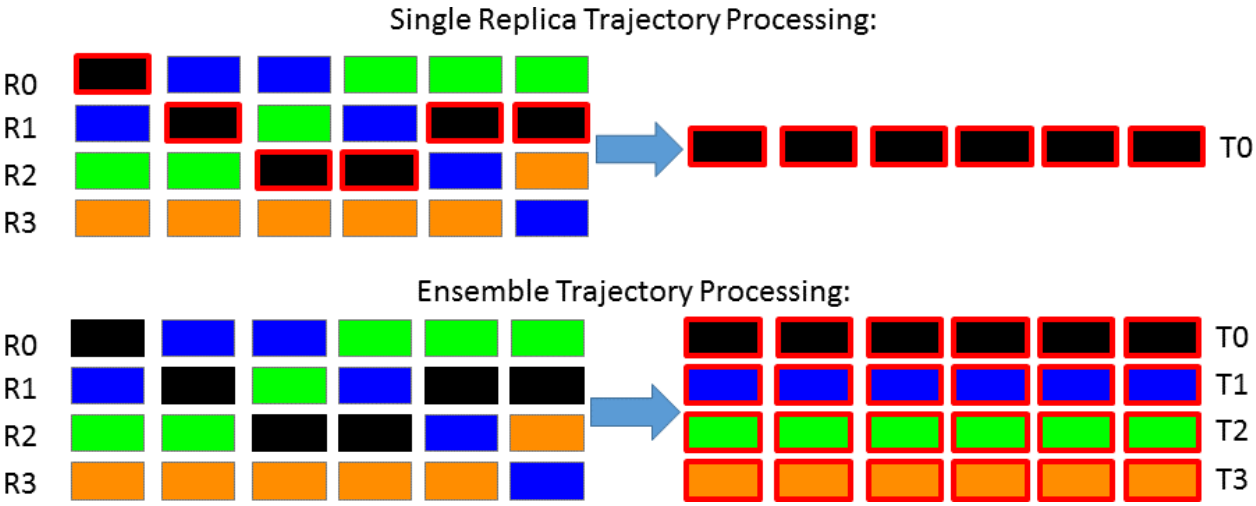

Figure 1. Top) Unsorted "replica" trajectories RX containing multiple temperatures (represented by different colors) are read in, but only frames at the temperature of interest (T0, black boxes, outlined in red) are actually used. Bottom) Unsorted "replica" trajectories RX are read in and sorted, after which frames at all temperatures TX are used.

4. Examples of Ensemble Processing

Example 1 (with sorting)

Given one ensemble of 16 temperature REMD replica trajectories of which the first is named 'rem1.nc.000', and a continuation of that ensemble of which the first is named 'rem2.nc.000', the following CPPTRAJ input will sort the input trajectories by temperature, remove residues named 'WAT' (i.e. remove water), calculate RMSD to a reference structure at each temperature, then write out 16 combined and sorted temperature trajectories. The **`ensemblesize`** command is used to improve ensemble setup performance in parallel:

```
# Load Topology
parm topology.parm7
# Load Reference Structure
reference myref.rst7
# Indicate ensemble size; improves parallel setup efficiency
ensemblesize 4
# Load Ensembles of Trajectories, Lowest Member Specified
ensemble rem1.nc.000
ensemble rem2.nc.000
```

```
# Action 1: Strip Residues Named WAT from Ensemble
strip :WAT
# Action 2: Calculate RMSD to Reference for Entire Ensemble
rms reference out rmsd.dat
# Write Out Sorted Ensemble Trajectories to sorted.nc.X
trajout sorted.nc netcdf
```

The output trajectories will have a numerical suffix appended to the name starting from 0 to indicate the position in the sorted ensemble (e.g. 'sorted.nc.0', 'sorted.nc.1', etc.). The file 'rmsd.dat' will have RMSDs for each member of the sorted ensemble in separate columns.

Example 2 (no sorting)

The following input could be used to process the results from several separate MD runs (md.300.nc, md.325.nc, md.350.nc, and md.375.nc) in parallel:

```
# Load Topology
parm topology.parm7
# Indicate ensemble size; improves parallel setup efficiency
ensemblesize 4
# Load Ensemble of 4 Trajectories, All Names Specified;
# No sorting will be performed.
ensemble md.300.nc trajnames md.325.nc,md.350.nc,md.375.nc nosort
# Calculate Distance Between Residues 1 and 20 for Ensemble
distance end_to_end :1 :20 out end_to_end.agr
```

In this case the additional trajectory names are specified with the **`trajnames`** keyword since they do not end with numerical suffices (preventing name auto-detection) and the **`nosort`** keyword indicates the trajectories should be processed as they are. The resulting output file 'end_to_end.agr' will contain distances from residue 1 to residue 20 from each of the four trajectories in order (i.e. results from 'md.300.nc' will be first, etc.).

5. Additional Considerations when Parallelizing Actions

Certain Actions require additional setup and/or consolidation steps in order to function properly in parallel. For example, any Action that may require using the first frame to be read in as a reference structure needs to have that structure broadcast from the master rank (which will by definition read in the very first frame of all frames to be read) to all other ranks; examples of this are the RMSD, distance RMSD, and symmetric RMSD calculations (**rms**, **drms**, **symmrmsd**), determination of native contacts (**nativecontacts)**, calculation of nucleic acid structure analysis (**nastruct**), etc. Actions that require calculation of a sum or an average over all frames need ranks to sync separately after all frames have been read; examples are in calculating atomic fluctuations (**atomicfluct**), average structure (**average**), matrices (**matrix**), grids, (**grid**), etc.

Syncing Data Sets to the master can be complicated in parallel for certain Actions that calculate data elements which may be present only in certain parts of the input trajectory frames, such as hydrogen bonds (**hbond**) or native contacts (**nativecontacts**) calculations. For example, if a certain trajectory has a hydrogen bond that is present for the first half of the trajectory but not the second half, some ranks will record the hydrogen bond as present while others will never have seen the hydrogen bond. As a result, only the ranks that have seen the hydrogen bond will have a Data Set for it. Because of this, certain Actions have an additional phase that occurs after the Trajectory Run called **SyncAction**, which is responsible for any additional data consolidation needed. For example, in **SyncAction** for the hydrogen bond Action, each non-master rank first communicates how many hydrogen bonds it has found, then sends the hydrogen bond data to the master. The master will then either create new Data Sets for hydrogen bonds that are not already present or combine the rank's hydrogen bond data with the existing one on master.

Certain Actions require a "read ahead" of a certain number of frames to work properly; currently these are the correlation of atomic motion vectors (**atomiccorr**), calculation of running average structure (**runavg**), and velocity autocorrelation function (**velocityautocorr**) when velocities are not present. These Actions function in parallel by requesting that each rank read ahead a certain number of frames prior to trajectory processing (reported by the **ParallelPreviousFramesRequired** function and processed by the **ParallelPreloadFrames** function prior to trajectory processing). The downside to this is that the frames that are read in this fashion are not processed by the Action queue; therefore, in parallel these commands should not be used with any command that may modify input frames (such as the RMSD calculation). The code will provide printed warnings to this effect for these commands.

6. Results of Darshan analysis on NCSA Blue Waters

1 process per replica

jobid: 3637698	uid: 30029	nprocs: 192	runtime: 6542 seconds

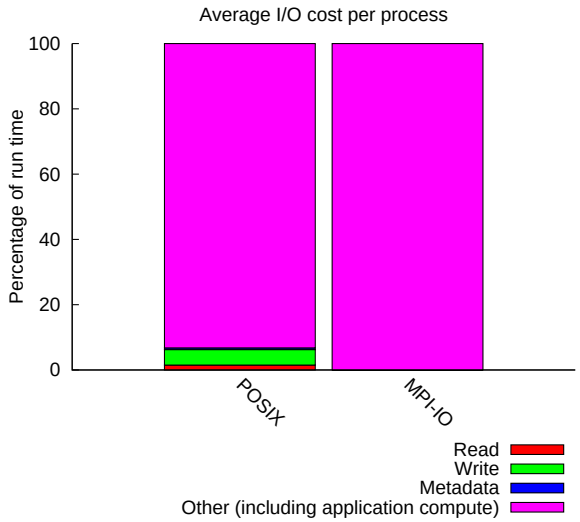

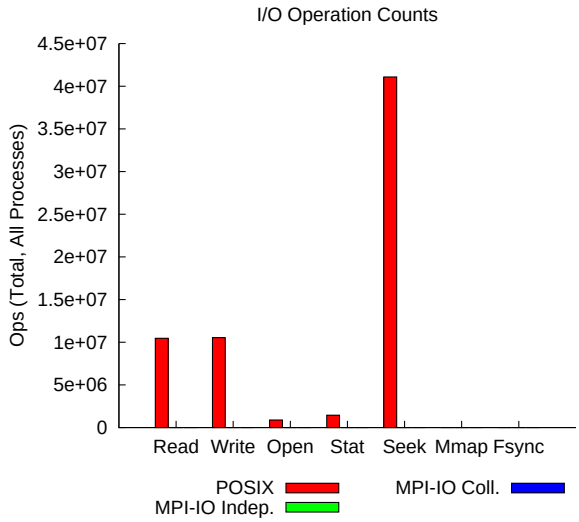

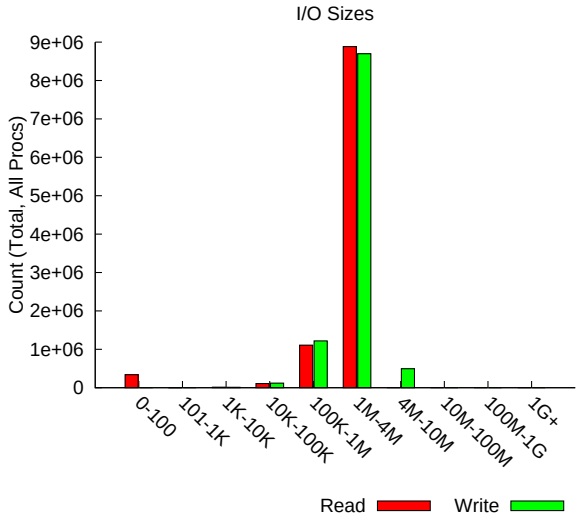

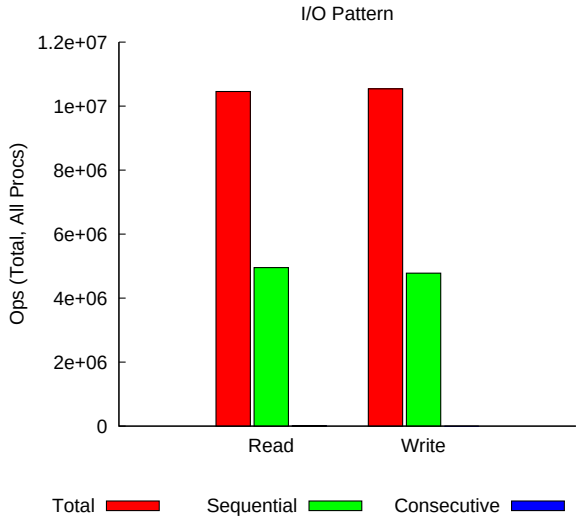

Most Common Access Sizes

access size	count
4194304	555264
4	223104
3	111360
81964	1344

File Count Summary
(estimated by I/O access offsets)

type	number of files	avg. size	max size
total opened	770	879M	1.2G
read-only files	578	1.2G	1.2G
write-only files	0	0	0
read/write files	192	47M	47M
created files	192	47M	47M

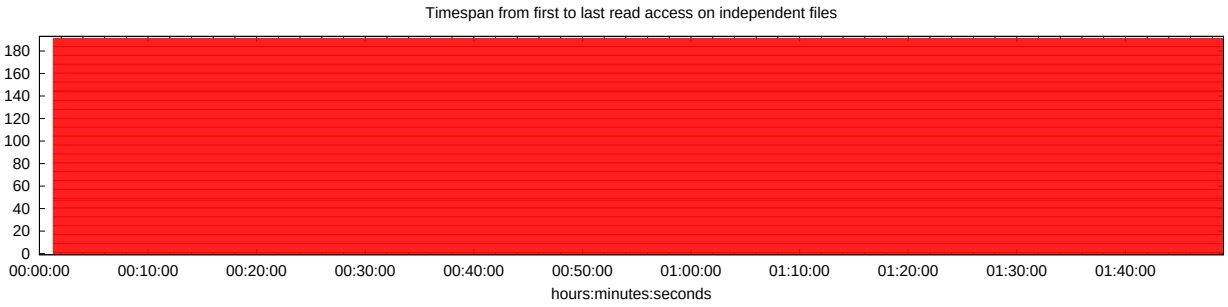

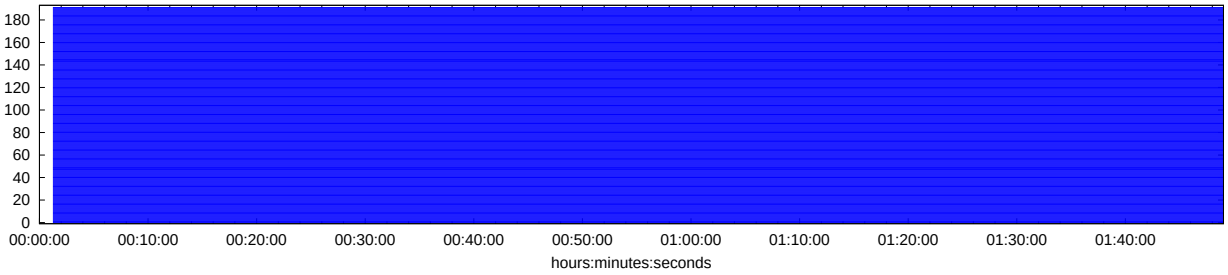

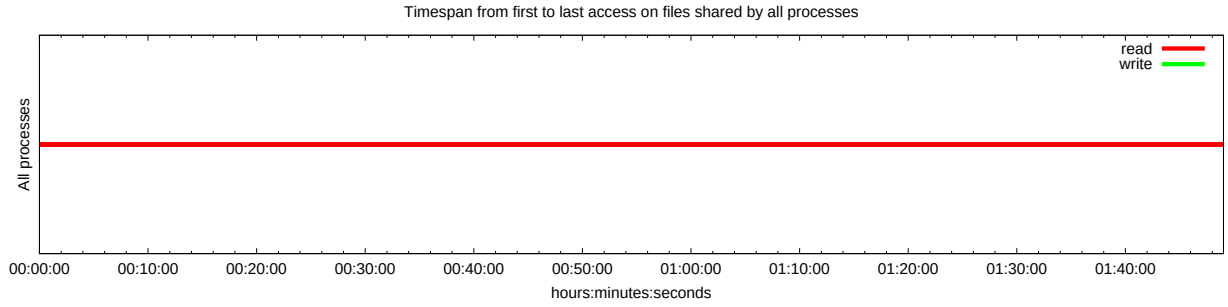

Average I/O per process

	Cumulative time spent in I/O functions (seconds)	Amount of I/O (MB)
Independent reads	54.937552	148300.697231
Independent writes	309.531370	168808.633598
Independent metadata	0.244743	N/A
Shared reads	44.410771	11577.651382
Shared writes	0.000000	0.000000
Shared metadata	32.748305	N/A

Data Transfer Per Filesystem

File System	Write		Read	
	MiB	Ratio	MiB	Ratio
/	0.00000	0.00000	265.83215	0.00001
/mnt/c	32411257.65088	0.00000	30696377.10168	0.99999

Variance in Shared Files

File Suffix	Processes	Fastest			Slowest			σ	
		Rank	Time	Bytes	Rank	Time	Bytes	Time	Bytes
...RAJ/rem.crd.114	192	33	0.124623	8.1M	113	3.785623	2.3G	0.52	1.75e+08
...RAJ/rem.crd.055	192	71	0.101840	8.1M	54	3.653352	2.3G	0.523	1.75e+08
...RAJ/rem.crd.035	192	176	0.168210	8.1M	34	3.306591	2.3G	0.276	1.75e+08
...RAJ/rem.crd.016	192	172	0.099367	8.1M	15	3.057339	2.3G	0.262	1.75e+08
...RAJ/rem.crd.136	192	128	0.083264	8.1M	135	2.988848	2.3G	0.23	1.75e+08
...RAJ/rem.crd.044	192	109	0.169516	8.1M	43	2.986023	2.3G	0.206	1.75e+08
...RAJ/rem.crd.129	192	49	0.386875	8.1M	128	2.980287	2.3G	0.187	1.75e+08
...RAJ/rem.crd.105	192	64	0.079357	8.1M	104	2.948687	2.3G	0.22	1.75e+08
...RAJ/rem.crd.129	192	118	0.091200	8.1M	128	2.940905	2.3G	0.209	1.75e+08
...RAJ/rem.crd.005	192	144	0.126275	8.1M	4	2.922351	2.3G	0.192	1.75e+08
...RAJ/rem.crd.069	192	76	0.080420	8.1M	68	2.910364	2.3G	0.203	1.75e+08
...RAJ/rem.crd.134	192	160	0.099607	8.1M	133	2.905410	2.3G	0.201	1.75e+08
...RAJ/rem.crd.001	192	25	0.201269	13M	0	2.892505	2.3G	0.189	1.75e+08
...RAJ/rem.crd.129	192	169	0.094766	8.1M	128	2.886342	2.3G	0.228	1.75e+08
...RAJ/rem.crd.087	192	94	0.076258	8.1M	86	2.880964	2.3G	0.204	1.75e+08
...RAJ/rem.crd.017	192	176	0.174198	8.1M	16	2.868618	2.3G	0.192	1.75e+08
...RAJ/rem.crd.001	192	118	0.224117	13M	0	2.868463	2.3G	0.188	1.75e+08
...RAJ/rem.crd.016	192	165	0.077711	8.1M	15	2.863224	2.3G	0.2	1.75e+08
...RAJ/rem.crd.001	192	165	0.220775	13M	0	2.861363	2.3G	0.185	1.75e+08
...RAJ/rem.crd.188	192	144	0.075851	8.1M	187	2.860075	2.3G	0.21	1.75e+08

4 processes per replica

jobid: 3637697	uid: 30029	nprocs: 768	runtime: 75 seconds

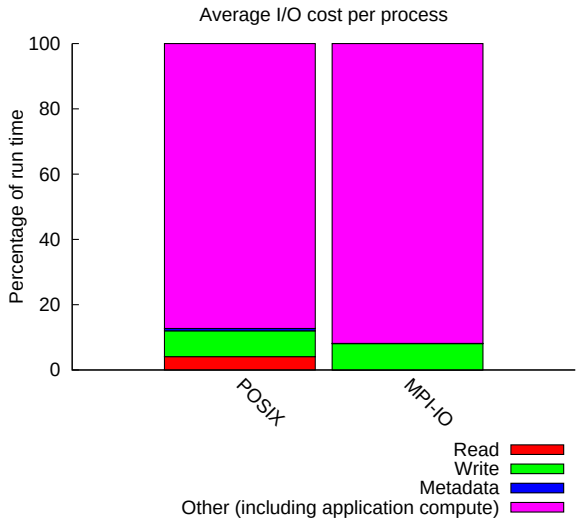

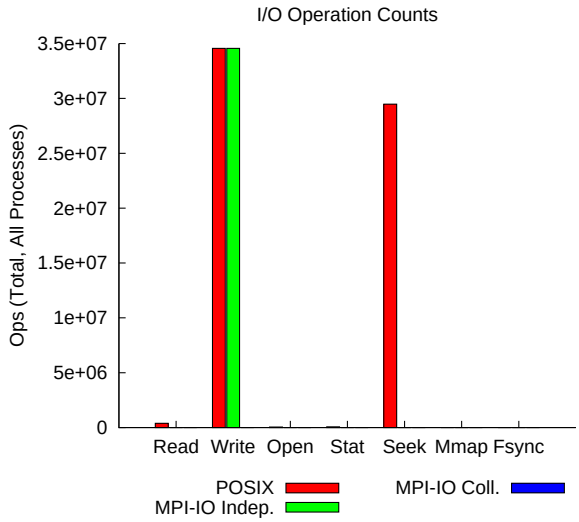

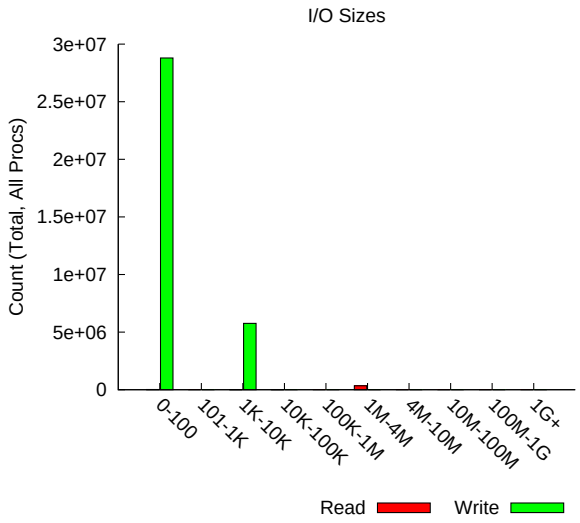

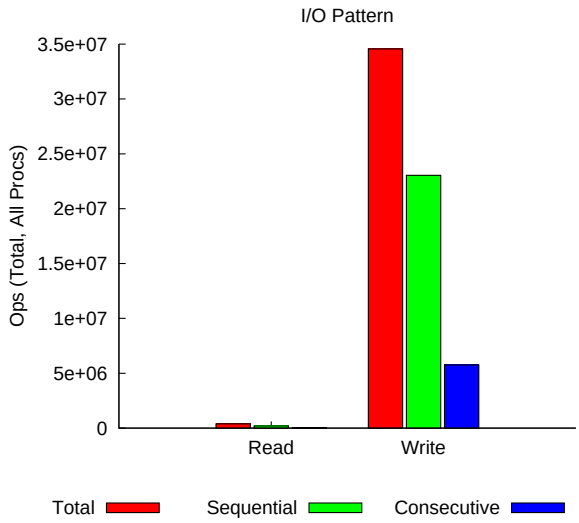

Most Common Access Sizes

access size	count
8	11520192
24	11520000
4	5769024
1572	5760576

File Count Summary
(estimated by I/O access offsets)

type	number of files	avg. size	max size
total opened	770	879M	1.2G
read-only files	578	1.2G	1.2G
write-only files	0	0	0
read/write files	192	47M	47M
created files	192	47M	47M

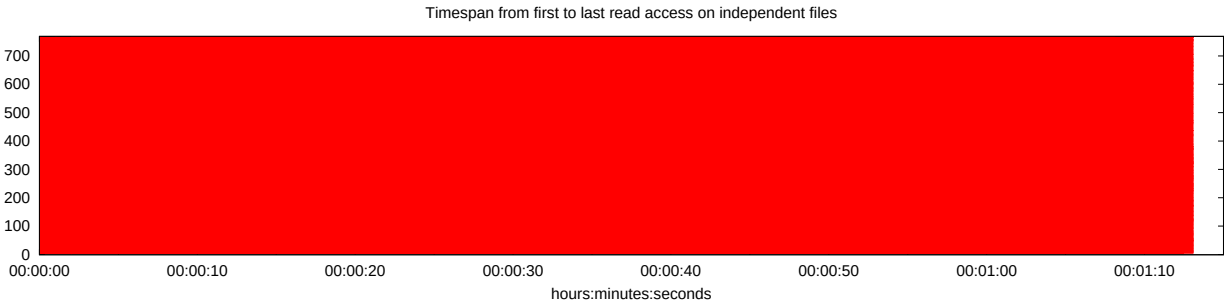

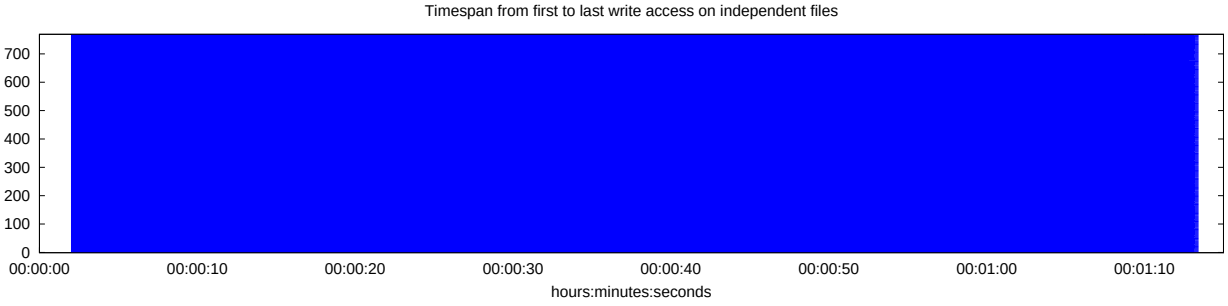

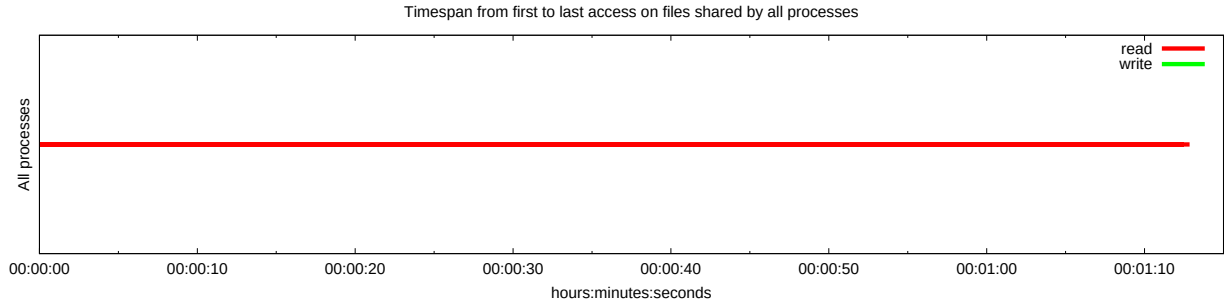

Average I/O per process

	Cumulative time spent in I/O functions (seconds)	Amount of I/O (MB)
Independent reads	2.481183	1759.851776
Independent writes	5.944697	11.730681
Independent metadata	0.256926	N/A
Shared reads	0.575267	22.598728
Shared writes	0.000000	0.000000
Shared metadata	0.272822	N/A

Data Transfer Per Filesystem

File System	Write		Read	
	MiB	Ratio	MiB	Ratio
/	0.00000	0.00000	1063.32861	0.00078
/mnt/c	9009.16333	0.00000	1367858.65796	0.99922

Variance in Shared Files

File Suffix	Processes	Fastest			Slowest			σ	
		Rank	Time	Bytes	Rank	Time	Bytes	Time	Bytes
...0-2.nowat.nc.17	4	71	7.222066	12M	68	9.84986	12M	1.11	1.78e+03
...n0-2.nowat.nc.6	4	27	6.498477	12M	24	8.658202	12M	0.851	1.78e+03
...0-2.nowat.nc.95	4	383	5.897218	12M	381	8.599763	12M	1.09	1.78e+03
...-2.nowat.nc.113	4	455	5.611698	12M	452	8.226118	12M	0.953	1.78e+03
...-2.nowat.nc.171	4	687	5.210494	12M	684	8.031502	12M	1.07	1.78e+03
...-2.nowat.nc.103	4	415	5.389028	12M	412	7.757018	12M	0.852	1.78e+03
...-2.nowat.nc.114	4	459	5.631735	12M	456	7.745795	12M	0.794	1.78e+03
...-2.nowat.nc.179	4	719	5.570011	12M	716	7.712014	12M	0.813	1.78e+03
...-2.nowat.nc.175	4	703	5.212742	12M	700	7.64251	12M	0.863	1.78e+03
...-2.nowat.nc.129	4	519	5.170251	12M	516	7.63082	12M	0.963	1.78e+03
...-2.nowat.nc.131	4	527	5.294287	12M	524	7.620544	12M	0.874	1.78e+03
...-2.nowat.nc.106	4	427	4.999499	12M	424	7.590011	12M	0.968	1.78e+03
...0-2.nowat.nc.27	4	111	6.020544	12M	108	7.583287	12M	0.612	1.78e+03
...0-2.nowat.nc.16	4	67	5.167673	12M	64	7.474077	12M	0.966	1.78e+03
...-2.nowat.nc.139	4	559	4.519389	12M	556	7.448442	12M	1.09	1.78e+03
...-2.nowat.nc.173	4	695	4.886281	12M	692	7.445973	12M	0.967	1.78e+03
...-2.nowat.nc.121	4	487	4.939084	12M	484	7.443133	12M	0.978	1.78e+03
...0-2.nowat.nc.24	4	99	4.958237	12M	96	7.424861	12M	0.903	1.78e+03
...-2.nowat.nc.135	4	543	5.013659	12M	540	7.423809	12M	0.928	1.78e+03
...-2.nowat.nc.160	4	643	4.57623	12M	640	7.413212	12M	1.07	1.78e+03