DATA REDUCTION IN FLASH-X

RAJEEV JAIN
Principal Specialist, Research Software Engineering
Mathematics and Computer Science Division
Argonne National Laboratory, Chicago, IL

Joint work with Houjun Tang,
Computer Research Scientist,
Scientific Data Management Group,
Berkeley Lab, Berkeley, CA

8/17/23
2023 HDF5 Users Group Meeting, The Ohio State University

FLASH-X

- First released as FLASH in 2000
- Primarily written in Fortran
- Base discretization is Eulerian, using finite volume methods
- Fundamental abstraction is a block with surrounding halo of ghost cells
- Supports both uniform and adaptive meshes
- Flash-X released in February 2022 https://flash-x.org/
  - Fundamentally altered infrastructure
  - Some architectural features are unchanged
  - Modernized quality and sustainability processes
- Highly scalable component based multiphysics simulation code for heterogeneous compute architecture
- AMR is used to:
  - Reduce memory footprint
  - Reduce computation

R&D 100 award winner for 2022

COMPUTATIONAL EXPERIMENTS WITH FLASH AND FLASH-X

Cosmological cluster formation

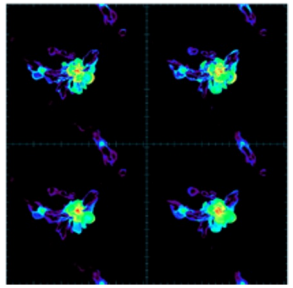

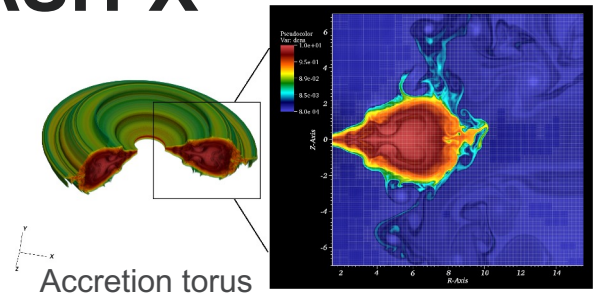

Accretion torus

Rayleigh-Taylor instability

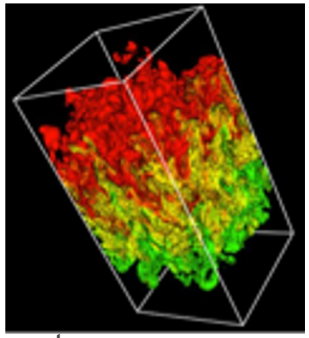

Core Collapse Supernova

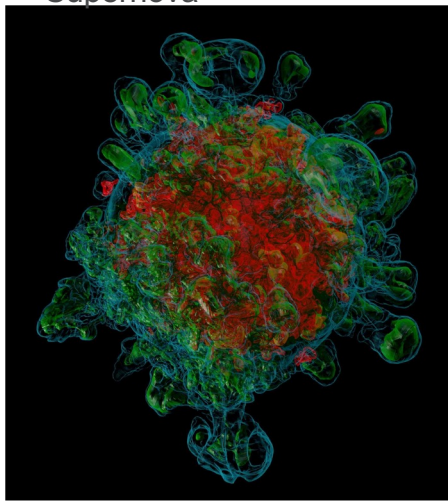

Type Ia SN

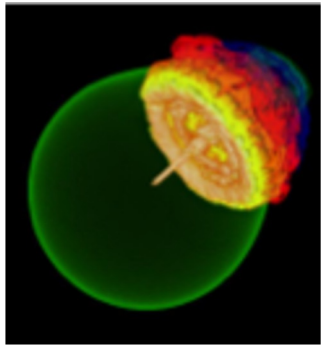

Rigid body structure

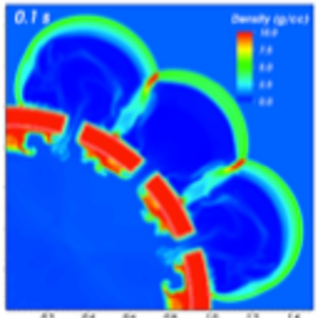

Gravity impact on boiling

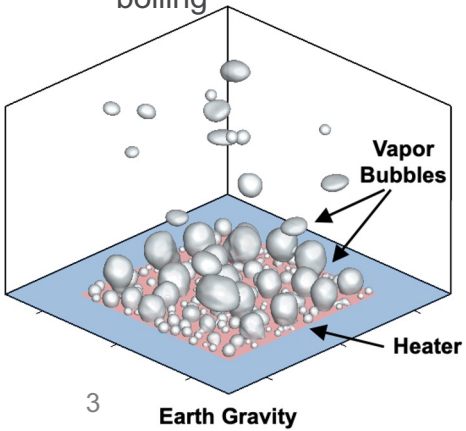

Ram pressure stripping

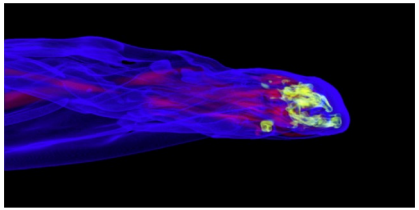

COMPRESSION KEYWORDS FOR PAR FILE

Parameter	Type	Default Value	Description
zfp_accuracy	REAL	0	ZFP accuracy mode, lossy compression
zfp_rate	REAL	0	ZFP rate mode, lossy compression
zfp_precision	REAL	0	ZFP precision mode, lossy compression
zfp_reversible	REAL	0	ZFP reversible mode, lossless compression
sz_abs_error	REAL	0	SZ/SZ3 absolute error mode, lossy compression
sz_rel_error	REAL	0	SZ/SZ3 relative error mode, lossy compression
sz_pw_rel_error	REAL	0	SZ point-wise relative error mode, lossy compression
sz_psnr	REAL	0	SZ/SZ3 peak signal-to-noise ratio mode, lossy compression
sz_norm	REAL	0	SZ/SZ3 mean square error mode, lossy compression
sz_abs_and_rel_error	BOOLEAN	FALSE	SZ3 absolute and/or relative mode, only active when both sz_abs_error and sz_rel_error are specified, lossy compression

For each of the keywords, two options are available

1. checkpoint: ***zfp_chk_accuracy***

and

2. plot files: ***zfp_plt_accuracy***

SEDOV PROBLEM SETUP

- Steps to setup problem:
- 1. clone the repo
- 2. run setup

3 types of compressions available

+hdf5zfp +hdf5sz3 +hdf5sz

./setup Sedov -auto -3d -nxb=32 -nyb=32 -nzb=32 +uhd +pm4dev -gridinterpolation=native DoAnalytical=True +parallelIO ***-objdir=sedov.323232.zfp*** *+***hdf5zfp*** *-site=****frontier***

- 3. *cd sedov.323232.zfp*
- 4. invoke *make –j*
- 5. check and specify simulation parameters in flash.par
- 6. launch the *flashx* executable (script that runs the executable)

COMPRESSION RESULTS FOR SEDOV TEST PROBLEM

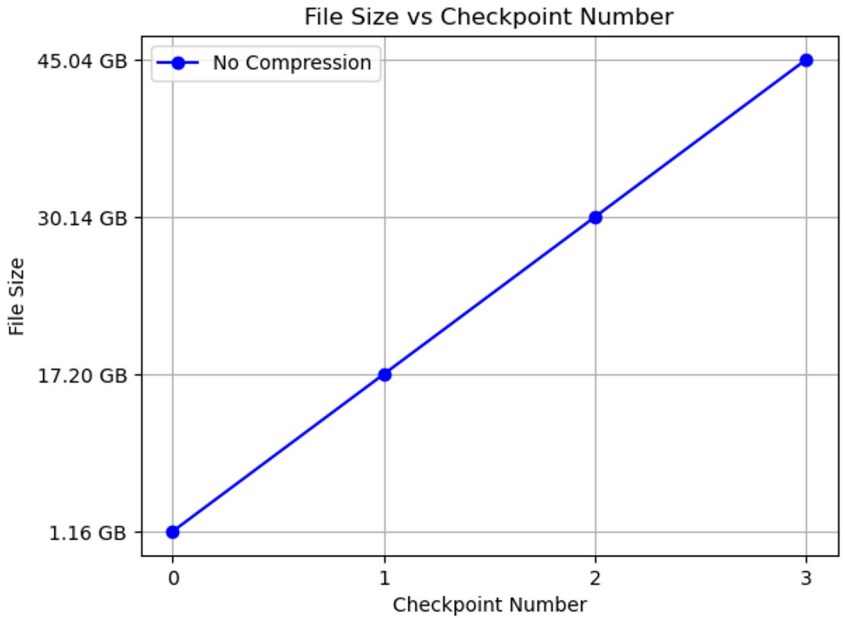

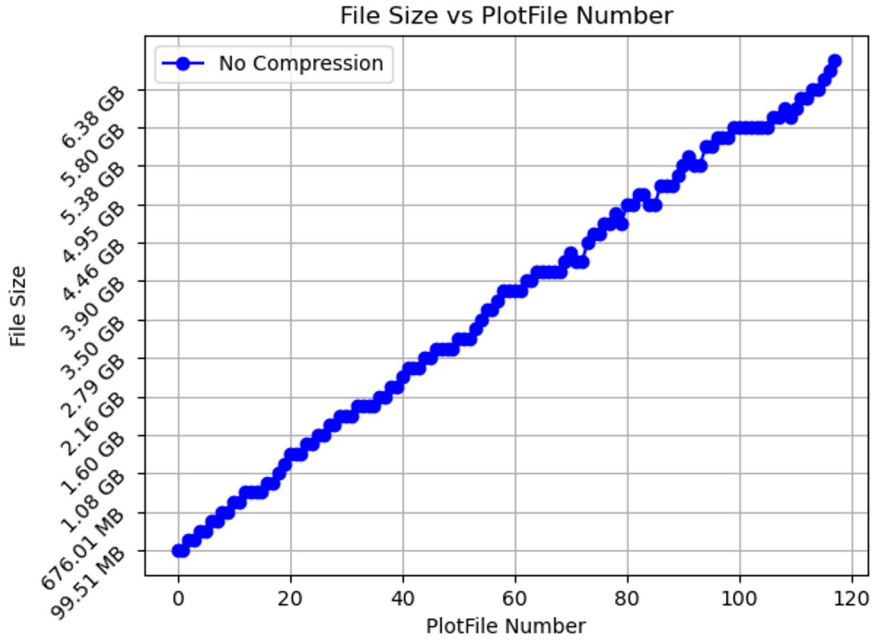

The Sedov problem is a well-known test case in computational astrophysics used to model the point explosion of a strong blast wave in a homogeneous medium. Sedov problem tests how well astrophysical hydrodynamics codes can model shock propagation, compression, rarefaction, and resulting fluid flows under these controlled conditions.

COMPRESSION RESULTS FOR SEDOV TEST PROBLEM

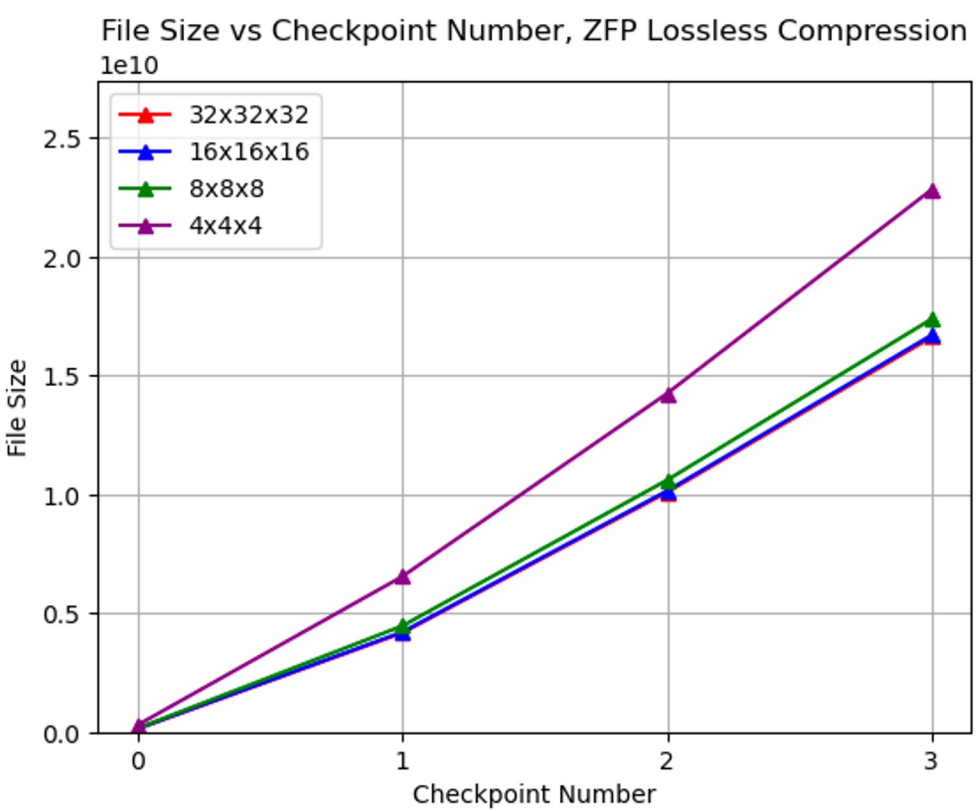

Chunk sizes have an effect on final size of the file for lossless compression!

4x4x4 has the highest size

16x16x16 and 32x32x32 have very similar file sizes

COMPRESSION RESULTS FOR SEDOV TEST PROBLEM

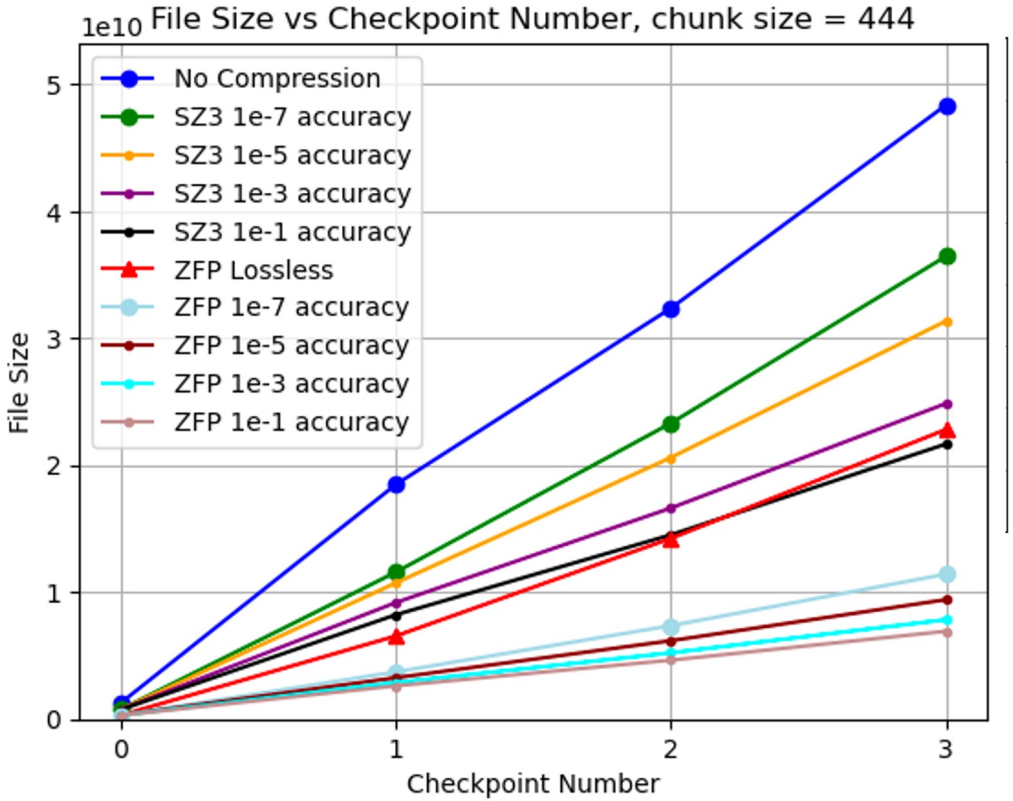

With chunk size of 4x4x4:

SZ3 with high accuracy gives the worst compression ratio. SZ3 is worse than ZFP lossless compression! even for 1e-3 accuracy!

Lower accuracy values of 0.1 are used to understand how errors are propagated.

COMPRESSION RESULTS FOR SEDOV TEST PROBLEM

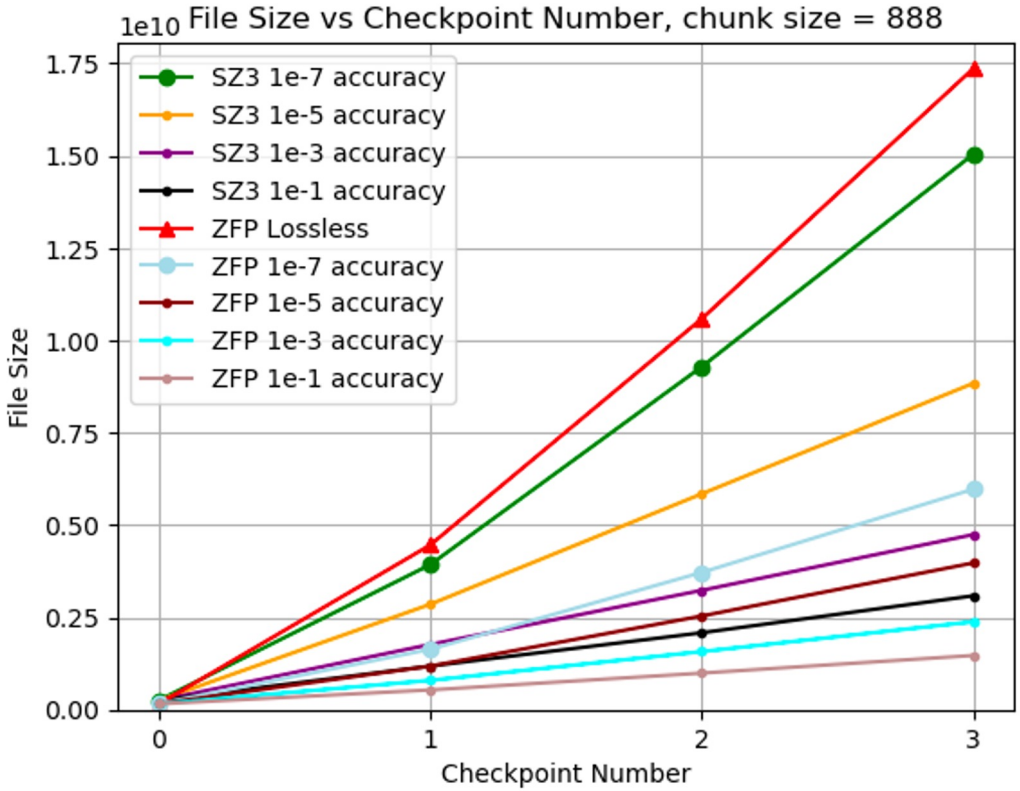

With chunk size of 8x8x8:

SZ3 better than ZFP lossless, but still far behind ZFP size for 1e-7 accuracy.

COMPRESSION RESULTS FOR SEDOV TEST PROBLEM

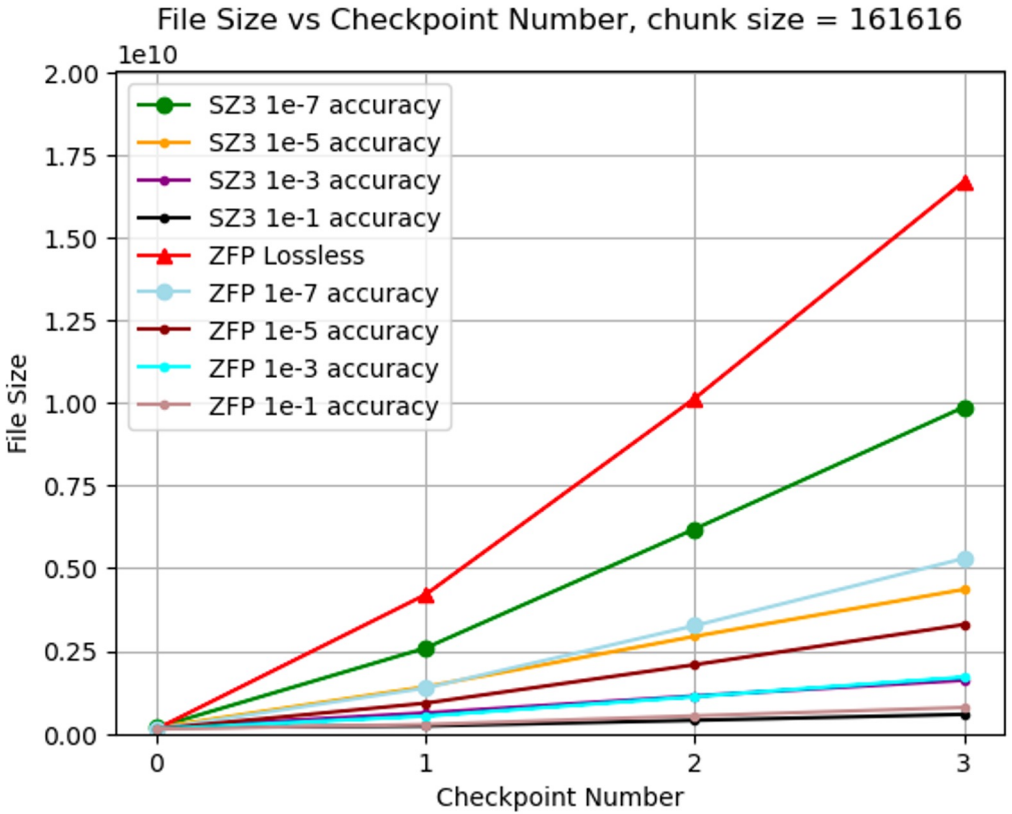

With chunk size of 16x16x16:

SZ3 better than ZFP lossless, but still far behind ZFP size for 1e-7 accuracy.

COMPRESSION RESULTS FOR SEDOV TEST PROBLEM

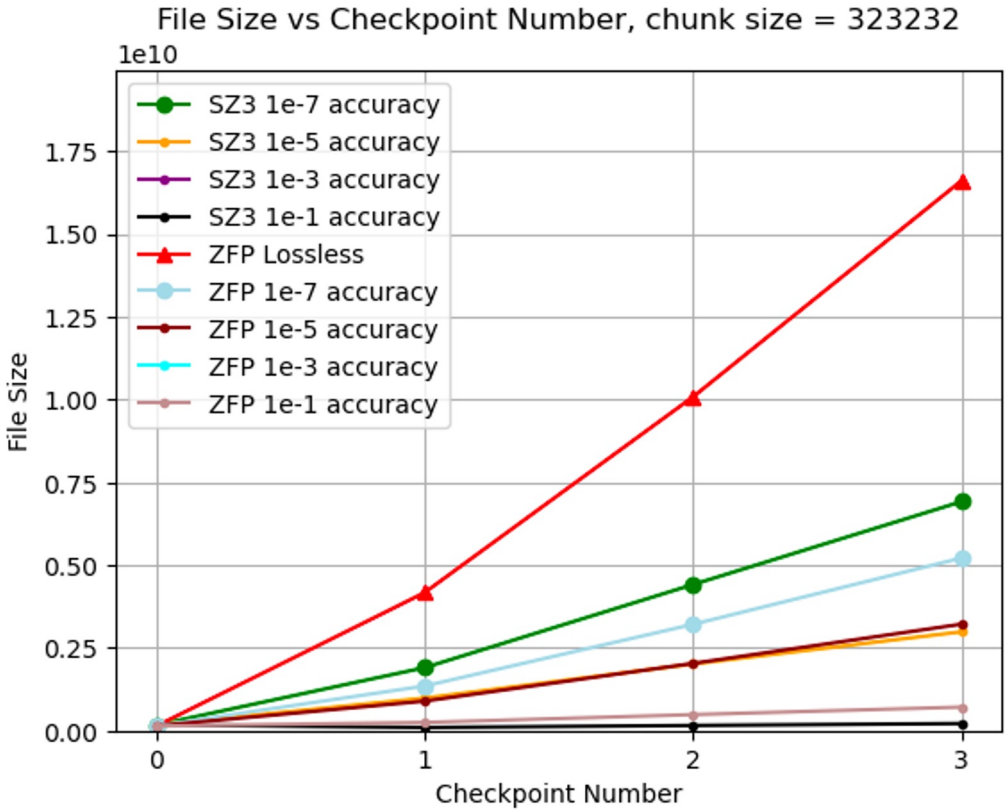

With chunk size of 32x32x32:

SZ3 and ZFP gap for 1e-7 accuracy is the smaller. SZ3 still has higher size for the same accuracy.

For 1e-1 accuracy, SZ3 is better than ZFP!

Ongoing Work

- Accuracy analysis
- Chunk size automation
- Performance measurement
- AsyncIO + Compression
- Setup more problems
  - mesh type
  - computing paradigms (CPU-GPU/GPU-GPU)

Jain, Rajeev, et al. "Accelerating Flash-X Simulations with Asynchronous I/O." 2022 IEEE/ACM International Parallel Data Systems Workshop (PDSW). IEEE, 2022.

ACKNOWLEDGEMENTS

- This work was supported by the U.S. Department of Energy Office of Science, Office of Advanced Scientific Computing Research (ASCR), and by the Exascale Computing Project (17-SC-20-SC), a collaborative effort of the U.S. Department of Energy Office of Science and the National Nuclear Security Administration.

THANK YOU FOR YOUR ATTENTION!

CODE CHANGES FOR SZ3, SZ AND ZFP COMPRESSION

- Include compression library dependencies
- Pass the compression values

```
…int* compression_type,
double* compression_val,
double* compression_val2,...
```

- If chunk size is specificized set it up
- Call appropriate compression routines
- Enabled data chunking in Flash-X
- Largest chunk size is the block size

```
#ifdef FLASH_HDF5_ZFP
#include "H5Zzfp_lib.h"
#include "H5Zzfp_props.h"
#endif

#ifdef FLASH_HDF5_SZ
#include "hdf5_sz/H5Z_SZ.h"
#endif

#ifdef FLASH_HDF5_SZ3
#include "hdf5_sz3/include/H5Z_SZ3.hpp"
#endif
```

```
SZ_errConfigToCdArray(&cd_nelmts, &cd_values, 0, *compression_val, 0, 0, 0);
H5Pset_filter(dataset_plist, H5Z_FILTER_SZ3, H5Z_FLAG_MANDATORY, cd_nelmts, cd_values);
```

ACCELERATING FLASH-X WITH ASYNCHRONOUS I/O

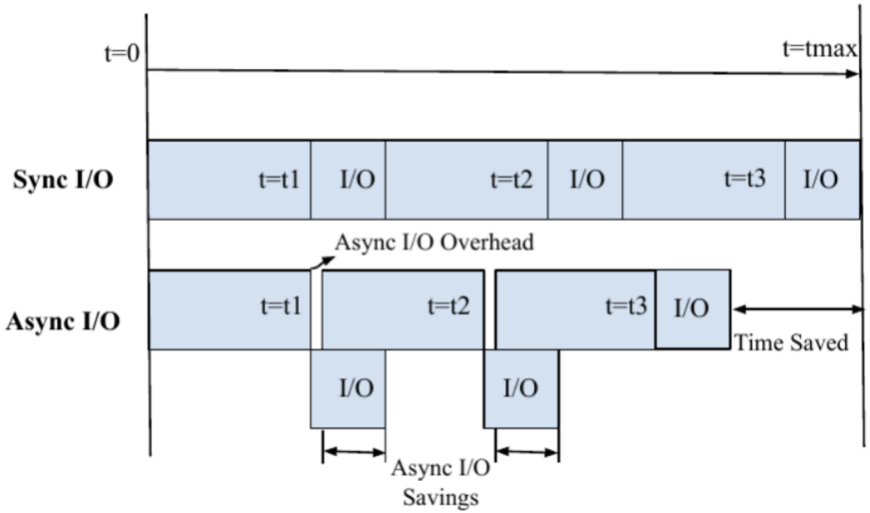

FLASH-X HISTORY

- First released as FLASH in 2000
  - Amalgamation of three independent Fortran 77 style codes
  - Underwent three rounds of architectural refactors over 6 years
  - 2007-now FLASH architecture has remained unchanged
    - Capabilities added, minor tweaks to infrastructure
    - Suitable for distributed memory parallel model only
- Flash-X released in February 2022 https://flash-x.org/
  - Fundamentally altered infrastructure
  - Some architectural features are unchanged
  - Modernized quality and sustainability processes

R&D 100 award winner for 2022