HPC / HTC Software Infrastructure for the Synthesis and Analysis of CMB Datasets

Julian Borrill (PI) [1,2] Colin Bischoff [3] Thomas Crawford [4] Matthew Hasselfield [5] Reijo Keskitalo [1,2] Theodore Kisner [1,2] Akito Kusaka [1,2,6] Nathan Whitehorn [7]

[1]University of California, Berkeley [2]LBNL [3]University of Cincinnati [4]University of Chicago [5]Flatiron Institute [6]University of Tokyo [7]University of California, Los Angeles

TOAST (Existing HPC Framework)

Time Ordered Astrophysics Scalable Tools (TOAST) is a public software framework originally developed for parallel simulation and analysis of CMB satellite data using MPI. It provides support for simulating instrumental effects and applying timestream operations and map-making to data that is distributed in memory across many processes. The framework uses communication between processes to efficiently support algorithms that account for correlations across the data set. TOAST has been used for data processing for the ESA/NASA Planck Satellite[1] and for simulation and analysis of other satellite, ground, and balloon-borne telescopes.

TOAST development and deployment has benefited from years of use on DOE supercomputing facilities and has been ported across multiple generations of computing architectures. The current version of TOAST is written in a combination of C++ and Python and has been successfully run at scale on up to 150,000 processes on the Cori supercomputer at LBNL. Current development plans include the porting of some kernals to CUDA to leverage current and upcoming heterogenous compute architectures.

Related Work

As part of this project, efforts are underway to simplify the internal TOAST data model which will make it easier to translate data between the two frameworks.

Acknowledgements

This work is supported by NSF Award 1835865.

References

[1] Planck Collaboration.
`NPIPE`: Joint *Planck* LFI and HFI data processing.
In Preparation, 2020.

[2] T. DeYoung.
Icetray: A software framework for icecube.
Computing in high energy physics and nuclear physics. Proceedings, Conference, CHEP'04, 78:463–466, 2005.

[3] Ruth Pordes, Don Petravick, Bill Kramer, Doug Olson, Miron Livny, Alain Roy, Paul Avery, Kent Blackburn, Torre Wenaus, Frank Würthwein, Ian Foster, Rob Gardner, Mike Wilde, Alan Blatecky, John McGee, and Rob Quick.
The open science grid.
Journal of Physics: Conference Series, 78:012057, jul 2007.

New Development

We are developing tools and training to enable CMB experiments to seamlessly use both HPC and HTC systems for their simulation and processing needs.

Compatible Data Models

The local data on each MPI process in a TOAST workflow can be conceptually mapped to a sequence of frames in the SPT3G data model. However, each experiment uses their own "schema" for specifying the detailed contents of their time-ordered data. We are making these data models consistent by simplifying the TOAST data model and implementing an API for specifying which data and metadata should be translated during the conversion.

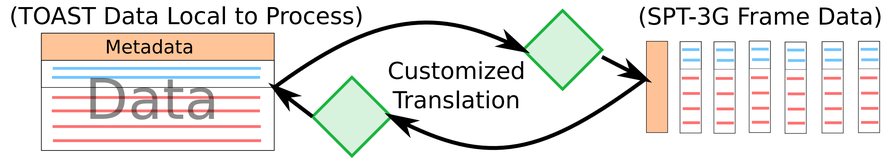

Pipeline Interoperability

Having compatible data models in memory allows bidirectional flow of data between pipeline components from both frameworks running in a single workflow. The three cases we consider are SPT-3G code running in a TOAST workflow on an HPC system:

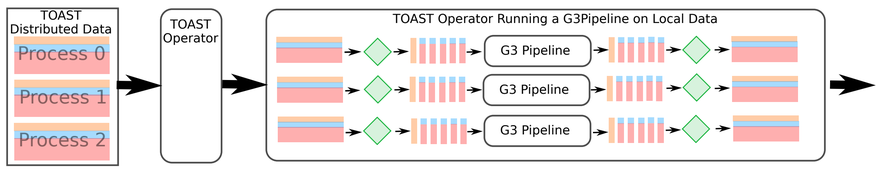

Some TOAST code running within an SPT-3G workflow on grid resources:

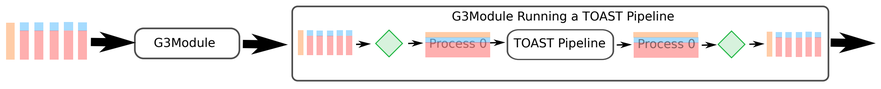

And a hybrid situation running on an HPC system:

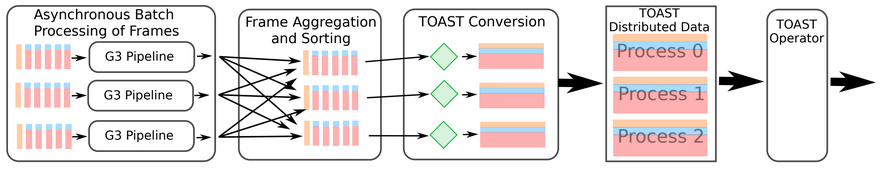

Community Outreach

Initial concepts and prototypes of this work are being tested within the Simons Observatory, SPT, and CMB-S4 collaborations. This effort will expand to dedicated workshops / tutorials as the software tools mature.

SPT-3G (Existing HTC Framework)

The public SPT-3G software framework was developed and released by the SPT collaboration. It provides facilities for data acquisition, storage (including compression), and handling. Because the map-making technique generally used in SPT analyses does not require tight coupling between nodes, the collaboration adopted a distributed HTC strategy for computing using the Open Science Grid[3] as a primary resource.

Because the computational needs of SPT-3G are similar to those of high-energy physics analysis, the software stack developed by SPT shares a great deal of commonality with the architecture of typical HEP data processing, in particular the IceTray framework[2] developed by the IceCube Neutrino Observatory and successfully in use on large scales since 2005.

Deployment on the grid has also benefited strongly from leveraging the HEP middleware ecosystem, including GridFTP, CVMFS, and OSG. Several other collaborations (POLARBEAR2, Simons Observatory) also use portions of the software.

SPT-3G is written in a combination of C++ and Python. This software stack is currently being used in production by the SPT-3G collaboration, handling the highest data rates in the field and running in parallel on many thousands of distributed cores.

Related Work

As part of this project, new features were added to support batch processing of frames across a group of MPI processes. This will ease the construction of hybrid workflows on HPC systems.

Relevant Links

TOAST Software:

https://github.com/hpc4cmb/toast

SPT-3G Software:

https://github.com/CMB-S4/spt3g_software