Simulations & Data Challenges: Simulation & Reduction

Reijo Keskitalo

CMB-S4 Collaboration Meeting
April 3-6, 2023

In a nutshell

- We simulate telescope motion and the resulting TOD for
  - Unlensed CMB
  - CMB lensing
  - Foregrounds
  - Instrumental noise and atmosphere
- Noise TOD is written to disk for further study
- The data are filtered using a realistic configuration ground and subscan polynomials and a common mode filter
- Filtered observation TOD are binned into NSide=4096 HEALPix maps stored in chunked HDF5 format (no explicit indexing
- In post processing, maps are co-added to form a variety of data splits

toast_sim_ground.py

- TOAST workflow for simulating and reducing CMB experiments on the ground (not satellite or balloon-borne)
- Made up of 39 data synthesis and reduction operators (and counting)
- Each operator can be individually enabled and configured – workflow enforces consistency
- MPI and OpenMP parallelized. Demonstrated performance over thousands of compute nodes
- World MPI communicator can be split into process groups that each process separate observations but join together in final mapmaking

Inputs - experiment model (see Sara Simon's talk)

- CMB-S4 hardware model translated into a TOAST focalplane file
  - One file for each frequency band, on each telescope
  - Detector position, bandpass, noise and polarization properties
- Observing schedule as a human-readable text file
  - Each entry is a constant elevation scan (CES): time, azimuth range and elevation
  - Broken up into several files if simulating shorter time spans

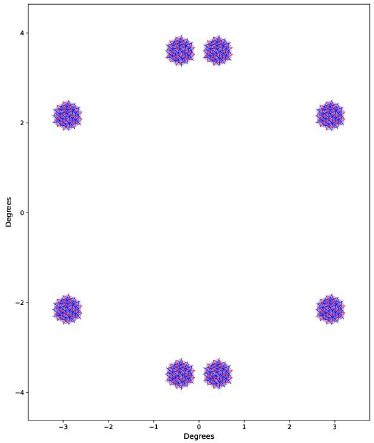

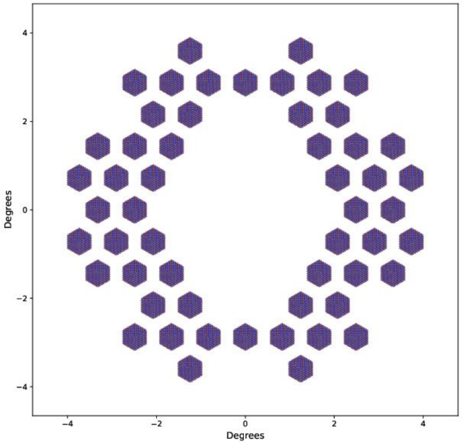

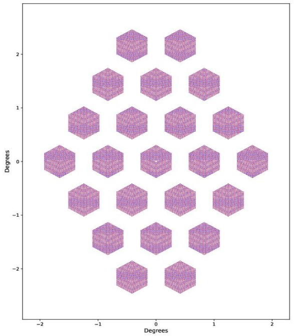

Inputs - sky model

- Sky signal is scanned from HEALPix maps that match the resolution of the output maps – no subpixel structure
- Beam and bandpass are convolved into the map – no anisotropy or per detector variations
- Separate maps for unlensed CMB, CMB lensing and foregrounds

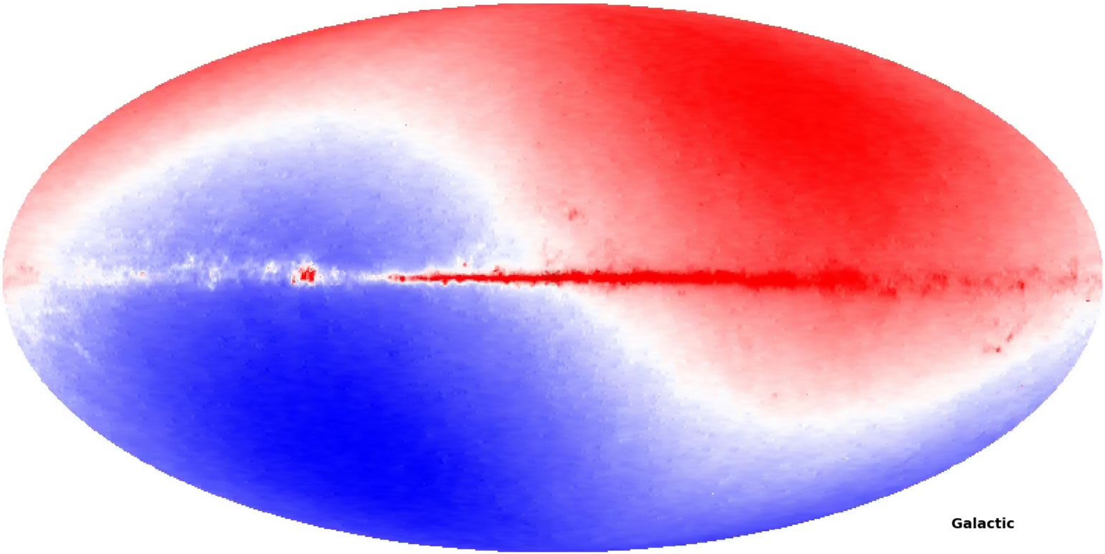

Inputs - atmospheric model

- Uses historic distributions of PWV, wind, surface temperature and surface pressure derived from MERRA-2 data
- Instantiates 3D realizations of atmospheric water vapor that translate into correlated detector noise when observed
- Calibrated on real experimental data from POLARBEAR, ACT, SPT and BICEP/Keck – work on-going
- Visualization of the simulated signal: https://crd.lbl.gov/assets/Uploads/atm-10-30-10v1.gif

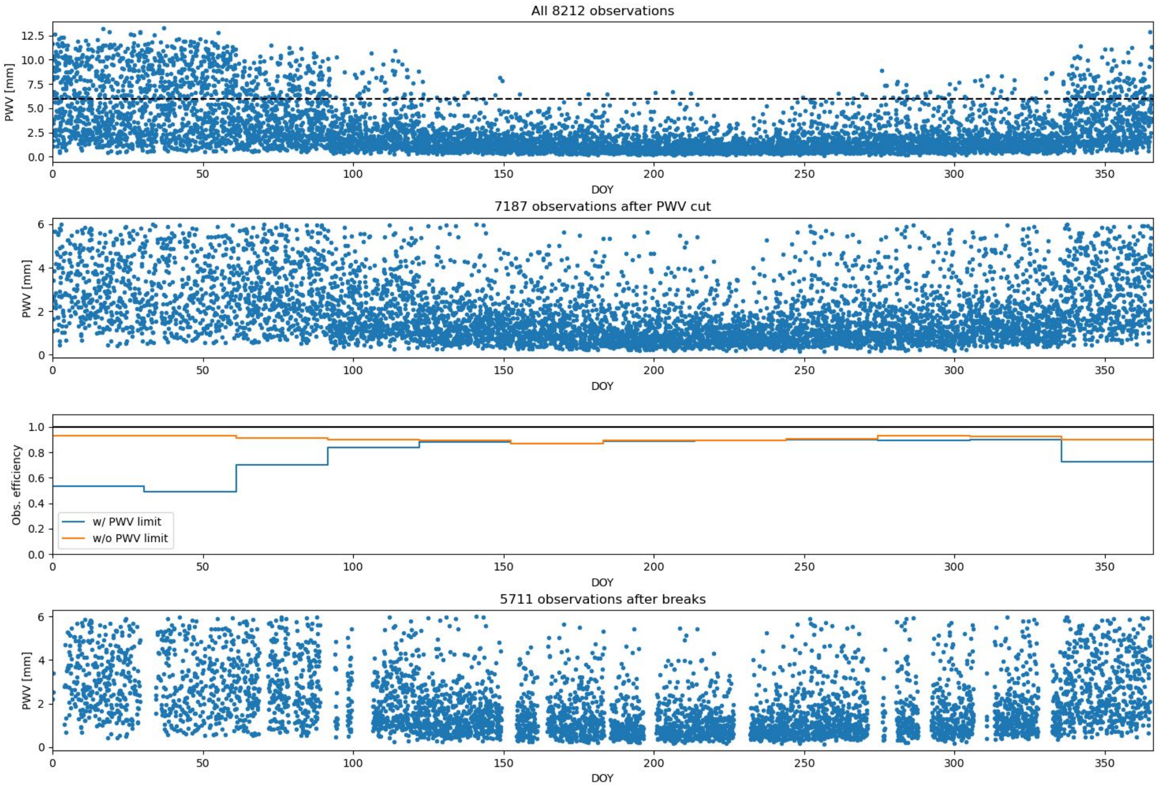

Inputs - noise model

- Instrumental noise based on an uncorrelated 1/f model
- BoloCalc used to predict response to changes in observing elevation and PWV

Other simulation elements

- Gain calibration errors drawn from a Gaussian distribution with a 1% width
- No ground pick-up, just the mode loss from filtering
- Perfect pointing and polarization parameters
- Focalplane yield is simulated, 20% of the detectors are flagged as bad
- Samples with 5° of the Moon are flagged as bad

The simulations in action

- Link to a month-long scanning movie on the Google drive: https://drive.google.com/file/d/1hQjhjmRma13Qoddprk8wXQ9rLbdal2Sa/view?usp=share_link

Hardware - Perlmutter CPU partition

- 3,072 compute nodes
- 2 AMD Epyc Milan CPUs per node, 64 cores each
- 512GB of RAM per node
- All-flash scratch file system
- Node-for-node, we measure ~10X speed-up over the Cori KNL system (previous flagship system)

Computing problem - CHLAT

- Simulate and reduce 5,711 hour-long observations in 6 frequency bands and 4 flavors, total of **137,064** observation maps
- Detector counts and sampling rates
  - LF 768 detectors@220Hz
  - MF 46,656 detectors@220Hz
  - HF 21,574 detectors@440Hz
- Each map can be run separately but requires 4 nodes for memory (MF/HF)
- TOAST supports large aggregate runs that produce multiple outputs but these are harder to fit in the queue and are subject to load imbalance
- Submitting each observation as a separate job would choke the scheduler, massive jobs have significant drawbacks too

  => allocate 16 or 24 nodes per job for 12 hours at a time, break them into 4-node gangs (8 nodes for noise) and have each of them loop over observations until the time runs out
  - MF/HF jobs with CMB or foreground complete in ~200 seconds
  - MF/HF jobs with noise take 400-800 seconds (simulates atmosphere, outputs TOD)
- ~55k node hours, equal to all of the 3,072 nodes for 18 hours

Data reduction – filtering and data quality

The filter stack was tuned on the design tool simulations and comprises

- 5th order azimuth-locked ground polynomial
- 1st order subscan polynomial
- Common mode filter operated at the optics tube (~0.5° across) level at every time sample

Application of the filter stack is quick:

- 35 seconds for MF
- 30 seconds for HF

Samples drawn over the Galaxy or point sources are not included in template fits.

TOD statistics (mean, standard deviation, skewness and kurtosis) are collected and recorded before and after filtering for further study.

Data reduction – mapmaking

Filtered TOD are binned into noise-weighted observation maps for rapid co-addition. We use nested NSide=4096 HEALPix pixelization (0.86’ across, 201M pixels)

Mapmaking cost:

- 16 seconds for LF
- 71 seconds for MF
- 68 seconds for HF

Single observation maps are sparse but explicit indexing adds extra overhead, especially when co-adding observation maps. We use chunked HDF5 files and break up the sky into 3,072-pixel submaps. Only the observed submaps are built, recorded or loaded.

TOAST deploys with command lines tools for co-adding the observation maps and converting them into conventional FITS format.

Next steps

- SPLAT simulations are mostly similar to CHLAT. Varying observation time may pose challenges to job scripting.
- SPSAT simulation and mapmaking is cheap due to lower sampling rate and map resolution. Computational challenge comes from the evaluation of the observation matrix.