Data Challenge 1

The CMB-S4 Data Management Team

CMB-S4

Goals

Data Challenges help to validate the experiment design and ensure DM readiness to start operations at commissioning. Each Data Challenge has 4 major goals:

1. Validate the point design associated with a particular review gate
2. Verify the simulation, reduction, and analysis software stacks
3. Demonstrate process scaling with data volume
4. Demonstrate process efficiency on that epoch’s computational architecture

CHALLENGE	REVIEW GATE	DATA VOLUME	CORE ARCHITECTURE
DC1	CD-1/PDR	12.5% => 1 year	Cori @ NERSC
DC2	CD-2/FDR	25% => 2 years	Perlmutter @ NERSC
DC3	CD-3	50% => 4 years	Aurora @ ALCF
DC4	Pre-Commissioning	100% => 7 years	NERSC-10 @ NERSC

Process & Schedule

Months Before Deadline	Project								Collaboration
	1.09.01 Subsystem Management	1.09.02 Data Movemement	1.09.03 Software Infrastucture	1.09.04 Data Simulation	1.09.05 Data Reduction	1.09.06 Transients	1.09.07 Site Hardware	L1 + Other L2s	SC + AWGs
...	Design, Develop, and Deploy (DD&D) DM infrastructure, including system-specific porting & optimization. Support intermediate design optimization studies.								
6	0. Point Design Freeze								
	1. Finalize DC scope								
5			2. Freeze infrastructure, simulation & reduction software						
4			3. Execute pipelines	4. Verify outputs					
3	6. Deliver	5. Register & archive							
2	7. Review internally								
		8. Distribute			9. Analyze & provide feedback			9. Analyze & provide feedback	
1									
0	Director's Review								
-1	Agency Review								

Schedule

- DC1 was originally aligned with a CD-1/PDR review expected in spring 2022, with the point design being the Preliminary Baseline Design (PBD).
- Given the short schedule and limited resources, the validation goal was descoped to simply demonstrating that the PBD meets the measurement requirements, rather than going all the way to the science requirements.
- CD-1/PDR was postponed, first until late summer 2022 and then until some time in 2023; the DC1 scope and timeline is being adjusted accordingly.
- Currently we are planning to complete DC1 in 2 stages:
    1. Validation of the PBD against the measurement requirements
    2. Validation of the PBD against the science requirements
- DC1 is the first time we have exercised our full set of processes at scale, so we expect many lessons learned!

1. Definition

Julian Borrill, Sara Simon

1. Definition

- The definition of a Data Challenge includes
  - The experiment
    - Instrument
    - Observation
  - The sky
    - CMB
    - Galactic foregrounds
    - Extragalactic foregrounds
  - The data products

1. Definition: Experiment

- Instrument definition is the preliminary baseline design
  - Detector and telescope information from pBD spreadsheet (frozen version from DC1)
  - Detector noise and noise vs. elevation values from jbolo (LATs +SAT MF/HF, SAT MF)
  - Platescales from SAT/LAT groups
- Observation Definition
  - Observation Efficiency from survey strategy group with time domain breakdown
  - Scan Strategy parameterization from survey strategy group
    - SAT - BICEP/Keck-like scans
    - SPLAT - SPT-like raster scans
    - CHLAT - variable speed, constant elevation scans

1. Definition: Sky

- CMB: realizations from 2 cosmologies with zero and threshold values of r, ΔN_{eff} in order to test analysis pipelines for false positives and negatives respectively.
- Galactic foregrounds - 3 models (optimistic, best-guess, and pessimistic) from the PanEx group*, spanning the range of possibilities and foregrounds and as a common set of foregrounds with other CMB experiments (SO, LiteBIRD, …)
- Extragalactic foregrounds - WebSky clusters + lensing

* see Friday's talk by Susan & Brandon

1. Definition: Data Products

- Stage 1: validation against measurement requirements
    - Atmosphere+noise
        - Per-observation detector timestreams for data cuts
        - Per-observation frequency maps for data quality & transients
        - Full- and half-mission frequency maps for analyses
    - CMB
        - Per-observation frequency maps for data quality & transients
        - Full- and half-mission frequency maps for analyses, including transfer function
- Stage 2: validation against science requirements
    - Foregrounds
        - Per-observation frequency maps for data quality & transients
        - Full- and half-mission frequency maps for analyses

2. Code Freeze

Reijo Keskitalo, Ted Kisner, Sara Simon

2. Code Freeze: Internal Process

Achieving a final release of all software packages is an iterative process:

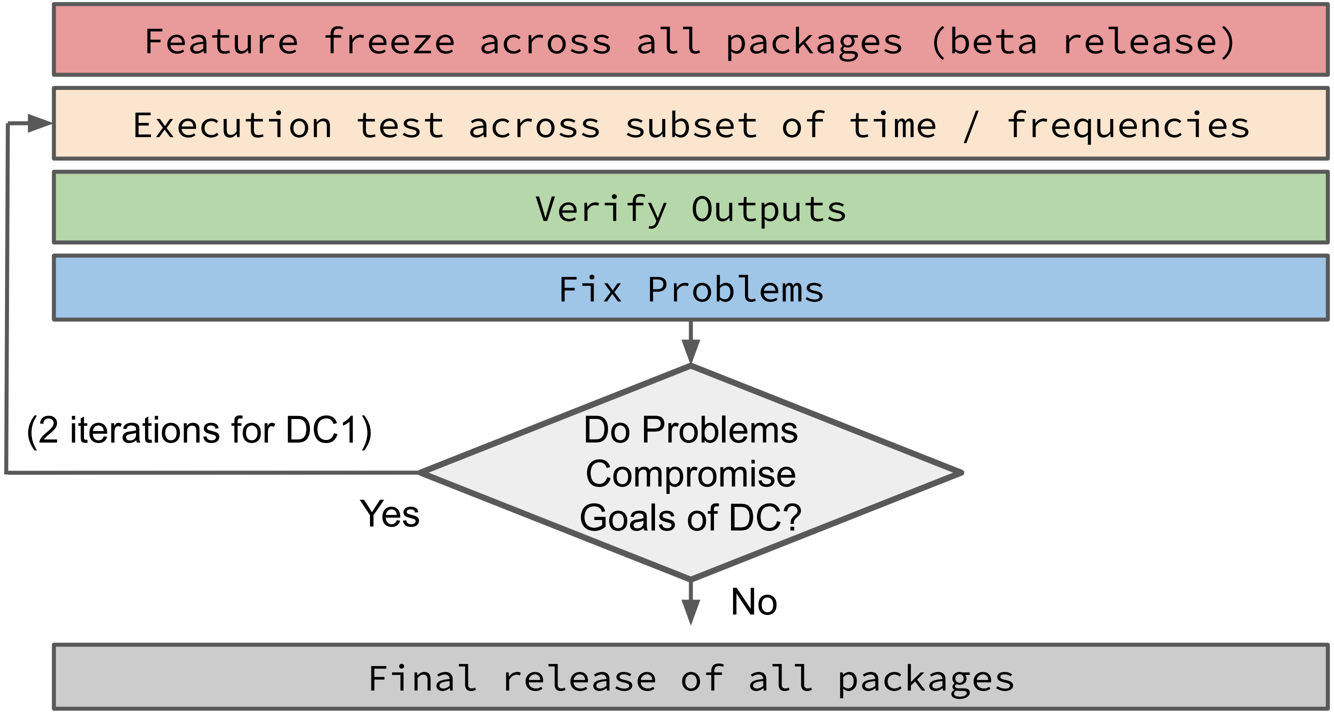

2. Code Freeze: Infrastructure

- Initial iteration of code freeze process identified several areas of improvement and problems:
  - Scripting of job batching and data transfer from scratch to community filesystem
  - Data dumping performance
  - Sporadic data corruption of simulated noise model data
- Fixes implemented and some data regenerated
- Other data files being fixed in post-processing rather than being regenerated
- Lessons learned: More effort needed in data verification during this iterative process

2. Code Freeze: Experiment Model Structure

- Python dictionary structure that outputs a human-readable toml file
  - bands: tophat frequency bands and their noise properties
  - wafers: wafer slots and properties of wafers including layout, frequencies, wafer-to-wafer spacing, clocking
  - tubes: tube slotting
  - telescopes: telescope/platform slots including beam size (FWHM) and platescale (mm to degrees on sky)
  - cards: readout card slotting (dummy readout mapping)
  - crates: readout crate slotting (dummy readout mapping)
  - detectors: individual detector properties
- These are positional slots tied to physical locations
- More details can be found in the instrument model documentation, documentation folder, and github repo (s4sim)

2. Code Freeze: Noise Properties

- NET and correlation factors from jbolo→ correlation factor multiplied by per detector NET to get effective per detector NET
- Instrumental 1/f noise (in mHz):

$$\mathrm{PSD} = \mathrm{NET}^2 \frac{f^\alpha + f_{knee}^\alpha}{f^\alpha + f_{min}^\alpha}$$

$\alpha = 1,\ f_{min} = 0.01,$ and $f_{knee} = 50.0$

- NET normalized by the NET at an elevation of 50° is scaled in elevation as:

$$\mathrm{NET} = \frac{A}{\sin\theta_{el}} + C$$

*A and C fit from NET vs. elevation output from jbolo

2. Code Freeze: Simulation

- We use generic, optimized and unit-tested TOAST-3 operators for all simulation steps:
  - Ground simulation operator for simulating telescope motion and drawing weather instances
  - Map scanning operator
  - Noise simulation operator
  - Atmospheric simulation operator
  - TOD dumping to HDF5
- Several optimizations for large focalplanes had to be implemented to make the simulations fit our computing budget

2. Code Freeze: Data Reduction

- We use generic, optimized and unit-tested TOAST-3 operators for all data reduction steps:
  - TOD statistics operator
  - Pointing expansion operators
  - Planet flagging operator
  - Ground filter
  - 1D polynomial filter
  - Common mode filter
  - Map binning operator
- We developed an ad hoc data format to speed up map writing and reduce disk space requirements: a compressed HDF5 format where only the relevant submaps of the full HEALPix map are written.

3. Execution

Reijo Keskitalo

3. Execution (1/2)

- Running on NERSC Cori KNL : 9200+ Knights Landing compute nodes, 68 physical cores, 96GB of memory
- Two execution stages by allocation year (AY):
  - Late AY21: backfill with small CMB jobs (32 nodes, <30 min each)
  - Early AY22: 2 system-wide reservations + large aggregate noise+atmosphere jobs (1024 nodes, 30-60 minutes each)

- All outputs are first written to a fast Lustre filesystem (scratch). NERSC increased our allocation there from 20TB to 1PB to facilitate the simulations.

CMB-S4

3. Execution (2/2)

- Even the 1PB disk space fills up - the data are moved to CMB-S4 project space on the community file system (CFS) as soon as they are written.
  - DOE increased our CFS allocation to 4TB to support DC1
- The transfer rates vary, we have seen rates as high as 8GB/s and as low as 100MB/s. We set up a Globus endpoint that allows transferring files directly to project ownership.
- Current status: CHLAT CMB & noise+atmosphere done for all 3,685 PWV < 2mm observations. About half of the remaining 2,026 high PWV observations also done.

4. Verification

Colin Biscoff, Reijo Keskitalo

4. CHLAT Pipeline Verification Plans

Check that output maps and auxiliary data products match expectations based on input noise levels and simulated sky signals

- Instantaneous Sensitivity: For the co-added maps derive effective instantaneous sensitivity→ check that it is between the numbers estimated for high elevation + good weather and low elevation + poor weather
- Observing Efficiency: Check that net observing efficiency from hit maps matches the defined efficiencies
- Map Depth: Check that the power spectrum of the noise+atmosphere map deconvolved with the beam and pipeline transfer function meets the PBDR measurement requirements
- Signal Transfer Function: Estimate signal transfer function from a single CMB realization that is processed into a filtered CMB map and compare to the input CMB spectrum

4. CHLAT PDB validation status

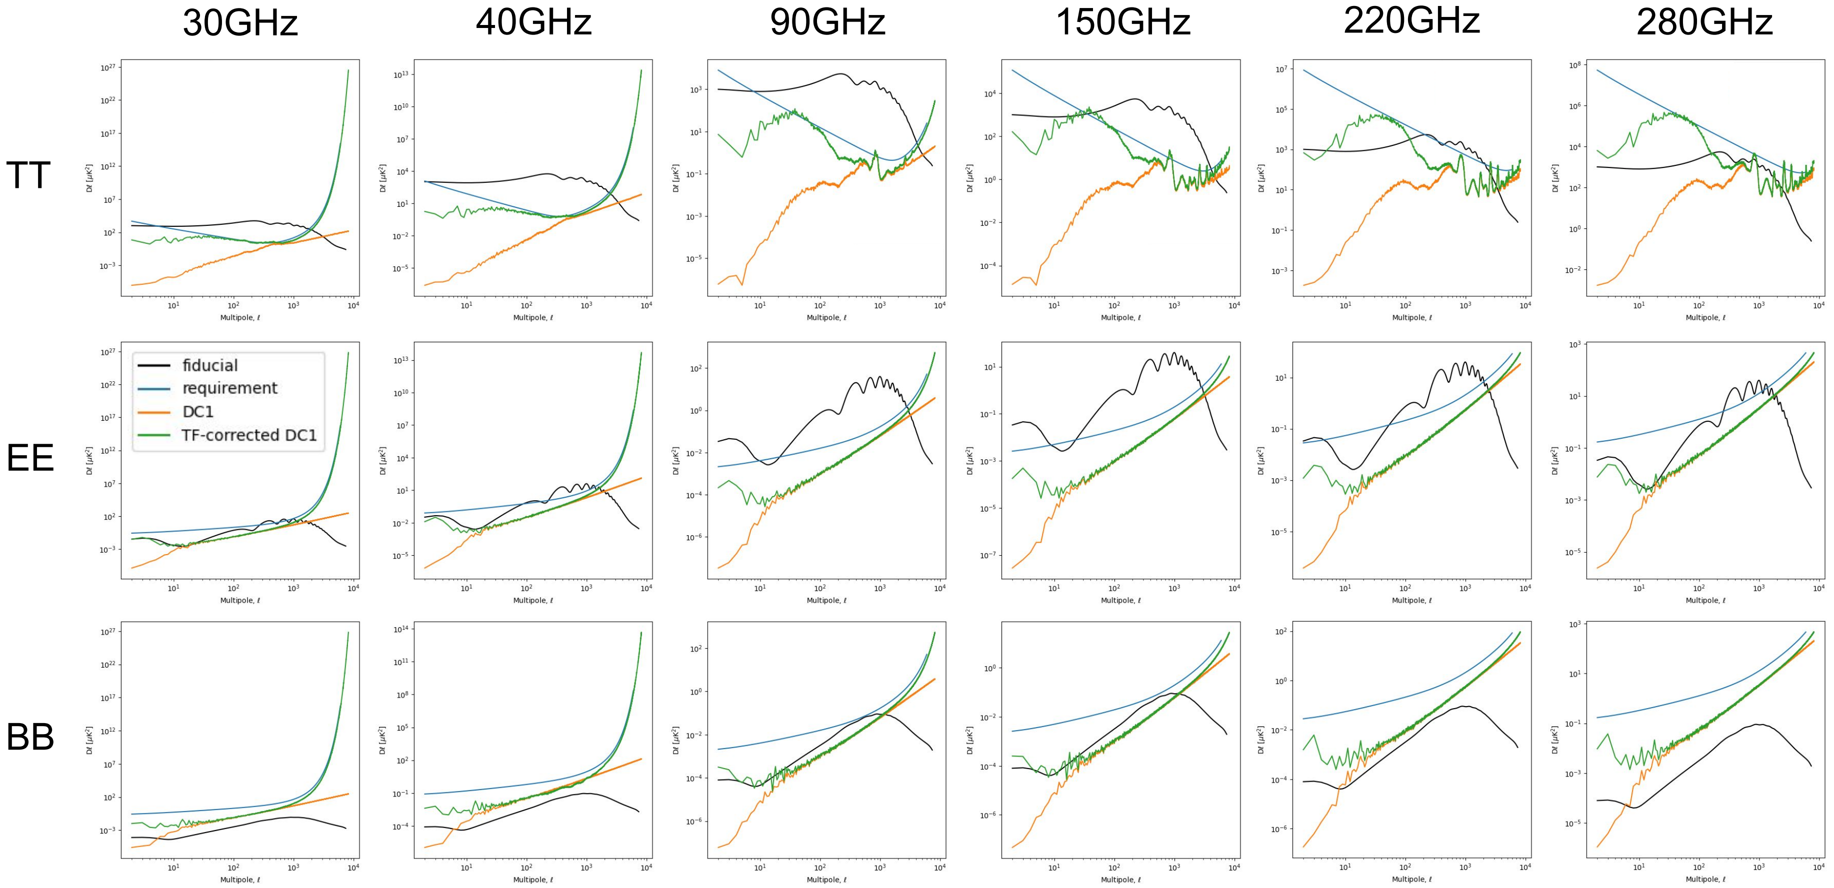

DC1 CHLAT meets the measurement requirement in TT, EE and BB.

4. SAT validation plans

Follow procedure for previous validation of Design Tool SAT sims (analysis / plots by Clem). BB noise spectra are compared to expectation from scaling BICEP/Keck noise and PBDR measurement requirements.

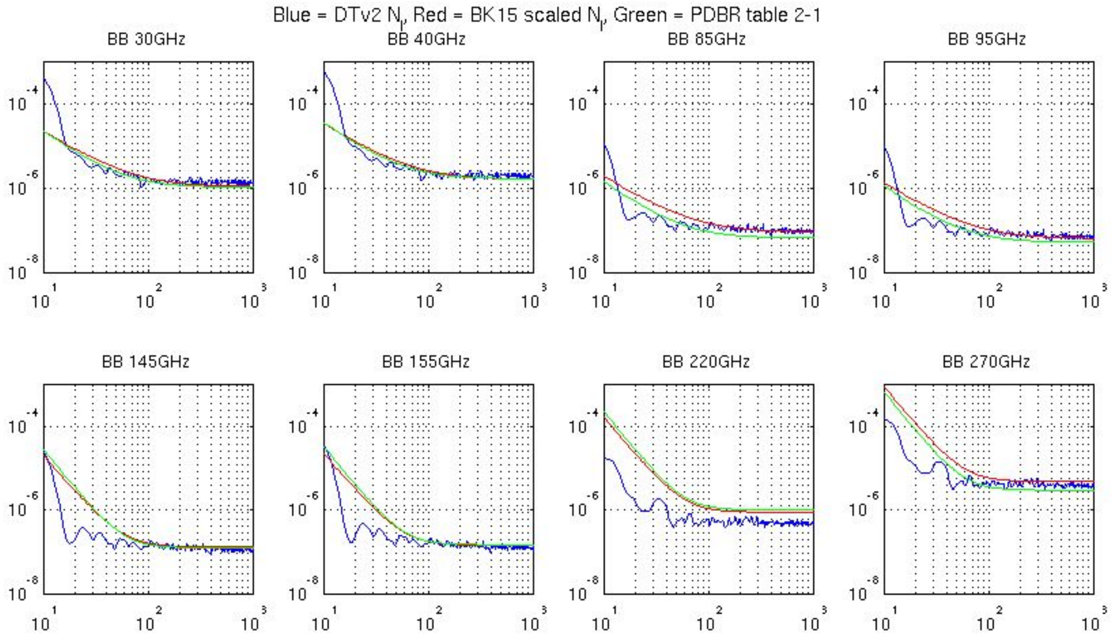

4. Data Quality

In addition to verifying the transfer function and noise properties of the full maps, we have a large quantity of time-ordered data and per-observation maps, as well as statistics that are generated and save during mapmaking. These will be used to start development of data quality tools.

- Framework for calculating and applying cut statistics
- Data visualization for human inspection of TOD and per-observation maps
- Inputs for transient analysis pipeline

5. Registration & Archiving

Eli Dart

5. Registration & Archiving

- Current registration and archiving activities are done by hand
- Data is published in the project portal, accessible via Globus (see distribution section)
- Two systems will be evaluated for future use by the project
  - Librarian - used by Simons Observatory
  - Rucio
    - Originally written by ATLAS experiment at LHC
    - Now adopted by LHC/CMS as well
    - Being evaluated for use by LSST/Rubin
- As Data Movement staff are brought on board, evaluation efforts will commence

6. Delivery

Julian Borrill

CMB-S4

6. Delivery

- Data delivery to the project and collaboration will consist of
  - Documentation
  - Announcement
  - Training session(s)
- These will be in place by the end of DC1

7. Internal Review

Julian Borrill

CMB-S4

7. Internal Review

"Those who do not learn from history are condemned to repeat it" - Churchill

"History repeats itself, first as tragedy, second as farce" - Marx

- Internal reviews allow for more critical self-assessment than external reviews (dirty laundry)
- Internal reviews also provide valuable input for the ensuing external review (clean laundry)
- We will have a mid-DC1 mini-review when the CHLAT runs are completed, re-scope/schedule the remainder of DC1 based on that, and then hold a full internal review at the end of DC1
  - This will include an opportunity for input from the AWGs & other L2s

8. Distribution

Andrea Zonca

8. Distribution

- Data portal based on LSST data portal (itself based on Globus Data Portal)
- Architecture:
  - Data on CMB-S4 Community File System
  - Browsable web interface on Spin (Kubernetes at NERSC)
  - Transfer via Globus
- Hierarchical permissions:
  - Public (everyone with a Globus account)
  - CMB-S4 Collaboration+Project (created from membership database)
  - Data Management
- Deployed at data.cmb-s4.org

Browse Dataset

Displaying files in dataset **DC1: CMB Coadded maps LAT Chile**. Learn more about this data set.

Choose the files you wish to download, then click the **Transfer** button to begin transferring files from NERSC. See here for detailed instructions.

Select all | Unselect all | Transfer

File Name	Size	Select
coadd_LAT0_CHLAT_f090_1of2_cov.fits	9.7 GB	☐
coadd_LAT0_CHLAT_f090_1of2_invcov.fits	4.8 GB	☐
coadd_LAT0_CHLAT_f090_1of2_map.fits	4.8 GB	☐
coadd_LAT0_CHLAT_f090_1of2_rcond.fits	1.6 GB	☐
coadd_LAT0_CHLAT_f090_2of2_cov.fits	9.7 GB	☐
coadd_LAT0_CHLAT_f090_2of2_invcov.fits	4.8 GB	☐
coadd_LAT0_CHLAT_f090_2of2_map.fits	4.8 GB	☐
coadd_LAT0_CHLAT_f090_2of2_rcond.fits	1.6 GB	☐
coadd_LAT0_CHLAT_f090_cov.fits	9.7 GB	☐