Computational challenges with SPT-3G and CMB-S4

Nathan Whitehorn

Michigan State University

March 17, 2022

MICHIGAN STATE UNIVERSITY

NSF

U.S. DEPARTMENT OF ENERGY | Office of Science

The South Pole Telescope

- Telescope designed for precision imaging of the cosmic microwave background
- Arcminute angular resolution
- Current receiver is SPT-3G: **16000 detectors**
- Began science operations in 2018
- $\sim$ 1 TB/day data rate
- NSF and DOE Support

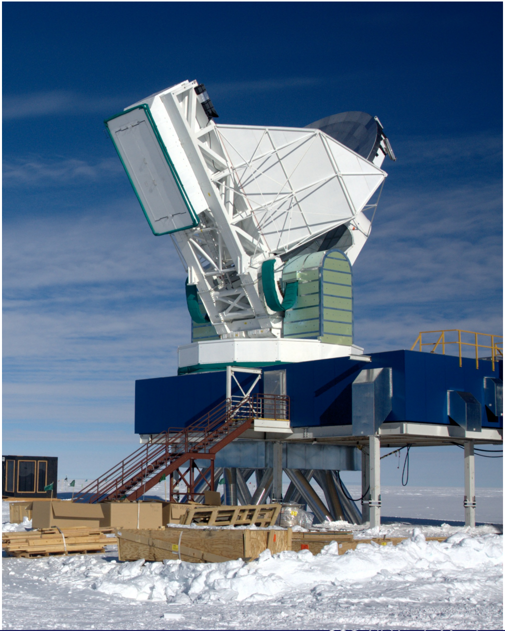

CMB-S4

- CMB data rates rise by a factor of 10 every 5–10 years
- Next-generation project with 30x the number of channels as SPT-3G
- 70 TB/day (800 MB/s) data rate – LHC-scale computing challenge
- Turns on late 2020s
- South Pole and Chilean Sites
- Planning computing architecture now

Science

- CMB temperature fluctuations to arcminute scales
- Measurement of E and B polarization modes
- Gravitational lensing measurements
- Galaxy clusters through Sunyaev-Zel'dovich effect
- Neutrino mass and number of species
- Properties of dark matter
- Inflation (through delensing)
- High-z star-forming galaxies
- Transient mm-wave sky
- ...

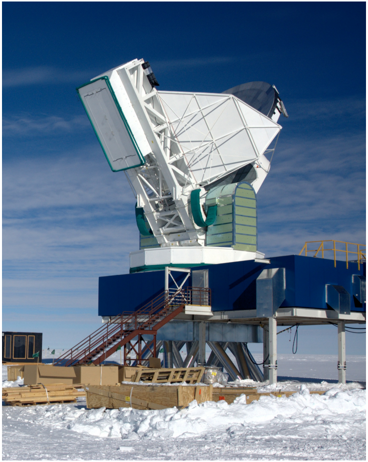

Raw Data from a CMB Telescope

Measure the current through resistors – CMB + many other things:

- Atmospheric emission
- pA-scale electronic noise
- pV-scale instabilities in electronics
- Galactic foregrounds
- Thermal fluctuations of focal plane
- Crosstalk
- Ground leakage
- Solid-state physics in wiring
- ...

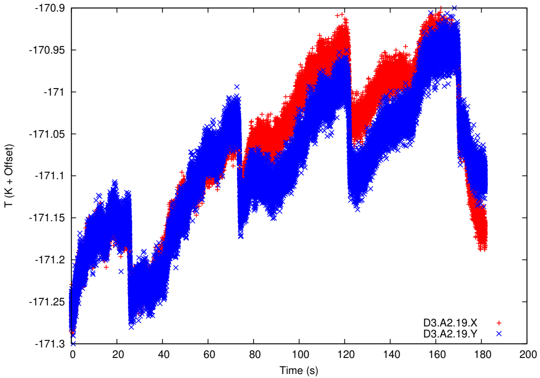

Making a Map of the Sky

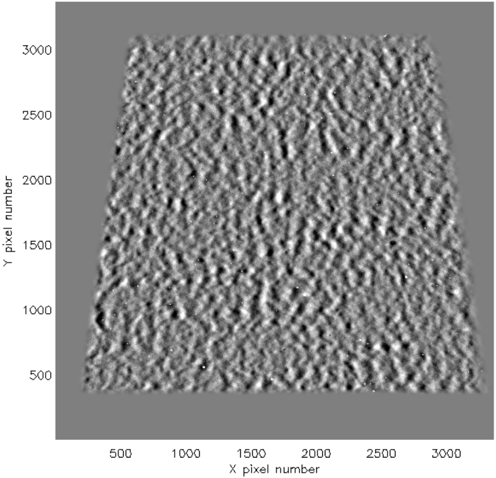

- Filter data for glitches, spectral lines, etc.
- Decorrelate detector modes and filter contaminants
- Compute detector polarization couplings
- Bin measured detector samples by sky position
- Average together all detectors
- **Every sample for every detector (12000 at 152 Hz) for all of time**
- Lots of computing, but few cycles/byte

Key Analysis Problems

- Process measurements from thousands of detectors as telescope is scanning into map of the sky
- Filter out instrumental noise in the process
- Measure, through Monte Carlo, effect of data processing on true sky
- **SPT Collecting** > 1 **TB of data per day** – largest in the field by factor of 4
- **No cuts**
- $\sim 5 - 10$ million CPU-hours per paper for SPT – 30x that for CMB-S4?

Using High-Throughput Computing

Approaches to making CMB maps

1. "Maximum-likelihood" map making: needs all data (PB!) in memory at once
2. "Filter-and-bin" map making: needs only a few minutes (GB) at once

SPT uses "filter-and-bin", which works great on an HTC architecture:

- Custom software to sip data from disk ~ 1 GB at a time (architecture derived from IceCube analysis software)
- In-band calibration data, no big databases or other choke points
- All open-source (Python + C++), so no licenses

SPT Infrastructure

- Data hosted at MWT2 center, distributed using GridFTP
- Software distribution by `cvmfs`, repository maintained by collaboration
- Custom software (GitHub: CMB-S4/spt3g_software) for analysis using C++ and Python
- Interactive nodes at Chicago double as submission nodes

A *big* thank-you to the staff at MWT2 for all their help.

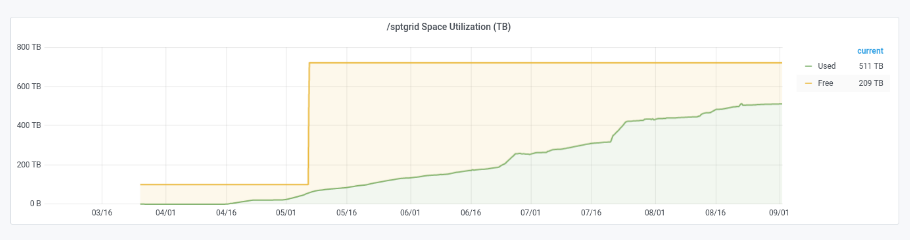

Collaboration Usage

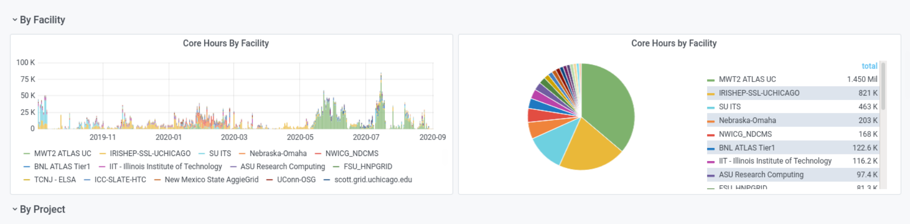

- $\sim$ 5 million CPU-hours in first year, published in 2021
- Widely distributed usage
- OSG users predominantly PhD students and postdocs
- Vast majority of collaboration computing done on OSG – this has made SPT-3G science possible

CMB-S4 Requirements

- Yet another big leap in data volume: $\rightarrow$ 10 PB
- Data too large to transfer from South Pole by satellite – on-site map-making for pole
- Some science goals, especially in Chile, may require fancier map-making algorithms than SPT
- Transient-science requirements may force near-live data analysis (<1 hour latency)

CMB-S4 Requirements

- Yet another big leap in data volume: $\rightarrow$ 10 PB
- Data too large to transfer from South Pole by satellite – on-site map-making for pole
- Some science goals, especially in Chile, may require fancier map-making algorithms than SPT
- Transient-science requirements may force near-live data analysis (<1 hour latency)
- More complex, higher throughput, deeply heterogeneous systems needed

Lessons Learned from SPT

- Think about architecture from the beginning: shoe-horning software into a different computing model is *hard*
- Attention to data transfer scheme is important
- We would love better ways to instantiate small compute environments (e.g. south pole cluster) that look like big ones
- Requirements for opportunistic HTC computing match requirements for small on-site computing clusters at the south pole

Current State-of-the-Art for CMB-S4

- Using a tightly-coupled HPC program for instrument simulations: TOAST [1]
  - Tight coupling needed for atmospheric simulations
- Using SPT-derived tools for initial data processing and archiving
- SPT-like (loosely-coupled, HTC-esque) required for south-pole live processing

[1] `https://github.com/hpc4cmb/toast`

Open Problems for CMB-S4

- Unclear what we will do as yet for final data processing in the North
  - Do we need full maximum-likelihood map-making? This forces us to tightly-coupled systems.
  - If we do, how much of the pipeline can we keep common with the loosely-coupled pipeline we also need?
  - Tight latency requirements on some of it
- Huge dataset with varying degrees of availability
  - How do we integrate late-arriving and early-arriving south pole data without just waiting a year?
  - How do we store and transport it effectively?
- Going to need a *lot* of computing
  - Minimization of effort: does every paper *really* need its own processing?
  - Accelerators: to what extent can we use GPUs? Low cycle-count-to-data-volume ratio makes this hard
  - Where do we get it and how do we connect disparate computing sites?

Conclusions

- Huge amount of science available in this field as precision and channel-count increase
- Rapidly increasing computational demands as field is scaling up
- Successful, OSG/HTC-style analysis pipeline from SPT-3G with published data
- Massive increase in data volume and requirements for CMB-S4
  - New techniques likely required
  - Need to think harder than we have about what we really need
  - We will be scrounging for every single cycle

Conclusions

- Huge amount of science available in this field as precision and channel-count increase
- Rapidly increasing computational demands as field is scaling up
- Successful, OSG/HTC-style analysis pipeline from SPT-3G with published data
- Massive increase in data volume and requirements for CMB-S4
  - New techniques likely required
  - Need to think harder than we have about what we really need
  - We will be scrounging for every single cycle
- Interesting problems ahead!