SPT-3G Computing

Benedikt Riedel[1], Lincoln Bryant[1], John Carlstrom[1], Thomas Crawford[1], Robert W. Gardner Jr.[1], Nicholas Harrington[2], Nicholas Huang[2], Sasha Rahlin[3], Judith Stephen[1], Nathan Whitehorn[4]

[1]University of Chicago, [2]University of California, Berkeley , [3]Fermi National Accelerator Laboratory, [4]University of California, Los Angeles

Introduction

The South Pole Telescope (SPT) [1] project uses the cosmic microwave background (CMB) to uncover some of the most important features of our Universe and the physics that govern it. The SPT is a 10-meter telescope located at the National Science Foundation (NSF) Amundsen-Scott South Pole station, the best site on Earth for microwave observations, and is optimized for sensitive, high-resolution measurements of the CMB [2]. It is funded jointly by NSF and the Department of Energy (DOE).

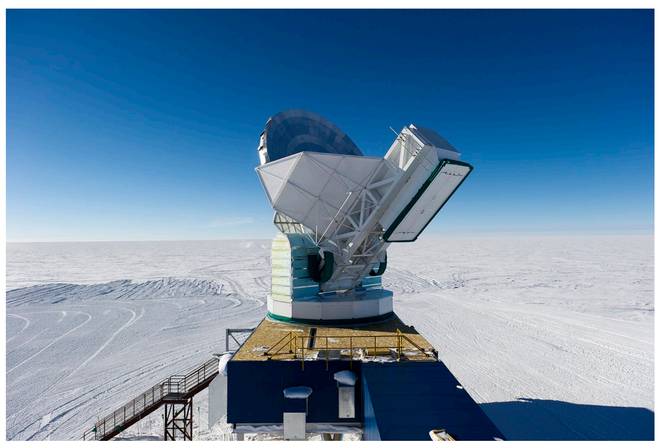

Figure 1: South Pole Telescope. Credit: NSF/SPT

Since the deployment of the SPT in 2007, the SPT team has completed two large surveys:

1. SPT-SZ: 2500-square-degree survey (2007-2011) [3]
2. SPTpol: 500-square-degree survey (2012-2016) [4]

The SPT-SZ and SPTpol observations have led to groundbreaking results that have moved the field of CMB research forward in significant ways. These results include the first galaxy clusters discovered using the Sunyaev-Zel'dovich (SZ) effect and the first detection of the elusive "B-mode" pattern in the polarization of the CMB.

The third-generation camera for SPT, SPT-3G, was deployed during Austral summer 2016-17 (first light January 30, 2017) and delivers a large improvement in sensitivity over the already impressive SPT-SZ and SPTpol surveys [5]. This increase in sensitivity comes from two technological advances:

1. Improved wide-field optical design that allows more than twice as many optical elements in the focal plane, and
2. Pixels that are sensitive to multiple observing bands in a single detector element.

The sensitivity of the SPT-3G receiver will lead to precise constraints on the sum of the neutrino masses and potentially deliver a detection of the primordial B-mode signal from a background of gravitational waves from the epoch of inflation.

Computing Requirements

The significant advances in sensitivity delivered by the SPT-3G receiver come primarily from increasing the number of detectors at the focal plane of the telescope. With this comes a concomitant increase in the requirements for data storage and computing needs. For a 5 year run time, an estimated 1.2 PB of storage and 150M CPU hours are required. The Open Science Grid (OSG) [6] group at University of Chicago maintains data analysis and storage infrastructure at both the South Pole and at the University of Chicago for the SPT-3G collaboration.

South Pole Computing

At the South Pole, OSG staff deployed new computing infrastructure during the Austral summer 2016-17. The new hardware consists of six servers and two storage chassis. The six servers are two Dell R330 and four Dell R730. The storage chassis are Dell MD1280s.

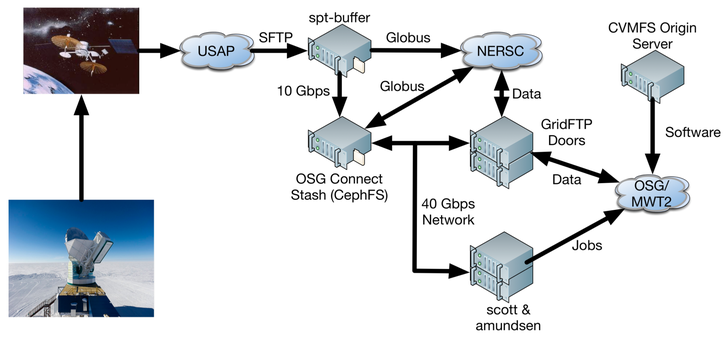

Figure 2: Overview of Data and job flow for SPT. Credit: NSF/SPT/NASA

A single Dell R330 is storage controller nodes for the two Dell MD1280 storage chassis. One of the Dell R730 acts as a hypervisor that hosts central services such as DNS, NFS server, login node, and puppet server. The three remaining Dell R730 are utilized as worker nodes for the HTCondor [7] local cluster. All machines of the same model have the same hardware configuration, such that they can act as a hot spare in case one of the machine cataphatically fails.

The Dell MD1280s are used to provide a large storage pool for online data analysis at the South Pole and as a data store for data that cannot be transferred via satellite. The large storage pool is managed using ZFS and exported to the R730 using NFS. The other storage chassis is configured as a JBOD. This is done to allow for data retrieval at the end of every Austral summer, i.e. disk filled with data are replaced with new disks for the upcoming season of data-taking.

During the Austral summer of 2017-18 performed we regular maintenance and software upgrade tasks. All machines were upgraded to Scientific Linux 7 and a general software update was performed. The hardware did not require any maintenance beyond replacing failed hard drives. We retrieved the raw data from the previous season and it has been shipped back to University of Chicago for ingestion into OSG Stash.

Northern Hemisphere Computing

The computing infrastructure in the northern hemisphere consists of two pieces: transferring, managing, and archiving data from the South Pole and data analysis. The data transfer from the South Pole is handled through the infrastructure provided b the United States Antarctic Program (USAP). SPT has an 125 GB/day data allocation on the South Pole Tracking and Data Relay Satellite System (SPTR) [8]. The daily transfer is then retrieved from USAP's servers in Denver, CO to a dedicated server at the University of Chicago. From this server the data is added to the OSG Stash and replicated to the High Performance Storage System (HPSS) tape archival system at DOE's National Energy Research Scientific Computing Center (NERSC) at the Lawrence Berkeley National Laboratory [9].

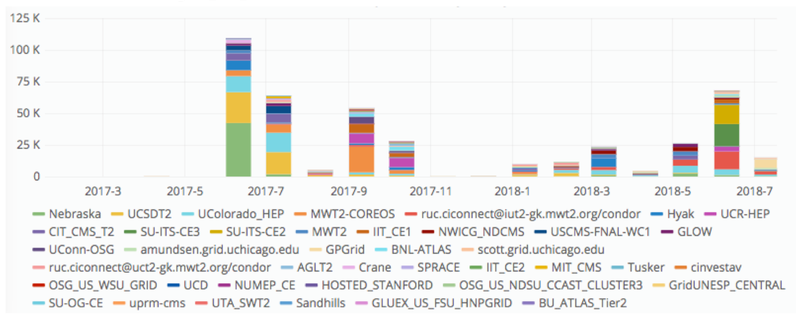

Figure 3: Summary of Wall Hours used by SPT across the OSG

For data analysis, we setup two Dell R630 servers to allow collaboration to perform interactive data analysis and submit large data reduction pipelines to the Open Science Grid. In addition to the servers, we deployed a copy of the SPT-3G software dependencies across the OSG and on the two dedicated nodes using the CERN-VM Filesystem (CVMFS) [10]. To allow for interactive data analysis, we have deployed a JupyterHub [9] instance on each serve The users can access the SPT data on OSG workers nodes using GridFTP [11] from OSG Stash.

Conclusion

We have described the computing workflow for the SPT 3G telescope both on site and at the University of Chicago. SPT using OSG as its pre-dominate computing resource is the first instance in which a CMB telescope has done so.

References

[1] The South Pole Telescope, arXiv: astro-ph/0411122
[2] The 10 Meter South Pole Telescope, arXiv: astro-ph/0907.4445
[3] SPT-SZ: a Sunyaev-ZePdovich survey for galaxy clusters, AIP Conference Proceedings, Volume 1185, pp. 475-477 (2009)
[4] SPTpol: an instrument for CMB polarization measurements with the South Pole Telescope, arXiv: astro-ph/1210.4970
[5] SPT-3G: A Next-Generation Cosmic Microwave Background Polarization Experiment on the South Pole Telescope
[6] The open science grid, 2007 J. Phys.: Conf. Ser., 78, 012057
[7] Distributed Computing in Practice: The Condor Experience, Concurrency and Computation: Practice and Experience, Vol. 17, No. 2-4, pages 323-356, February-April, 2005
[8] https://www.usap.gov/technology/contentHandler.cfm?id=1971
[9] http://www.nersc.gov/users/storage-and-file-systems/hpss/
[10] CernVM – a virtual software appliance for LHC applications, 2010 J. Phys.: Conf. Ser. 219 042003
[11] Globus GridFTP: what's new in 2007, Proceedings of the First International Conference on Networks for Grid Applications, GridNets '07