

UC Irvine

UC Irvine Previously Published Works

Title

CMB-S4 Science Case, Reference Design, and Project Plan

Permalink

<https://escholarship.org/uc/item/9hv6s51j>

Authors

Abazajian, Kevork
Addison, Graeme
Adshead, Peter
[et al.](#)

Publication Date

2019-07-09

DOI

10.48550/arxiv.1907.04473

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

CMB-S4 Science Case, Reference Design, and Project Plan

CMB-S4 Collaboration

July 9, 2019

Preface

This document on the CMB-S4 Science Case, Reference Design, and Project Plan is the product of a global community of scientists who are united in support of advancing CMB-S4 to cross key thresholds in our understanding of the fundamental nature of space and time and the evolution of the Universe. CMB-S4 is planned to be a joint National Science Foundation (NSF) and Department of Energy (DOE) project, with the construction phase to be funded as an NSF Major Research Equipment and Facilities Construction (MREFC) project and a DOE High Energy Physics (HEP) Major Item of Equipment (MIE) project. At the time of this writing, an interim project office has been constituted and tasked with advancing the CMB-S4 project in the NSF MREFC Preliminary Design Phase and toward DOE Critical Decision CD-1. DOE CD-0 is expected imminently.

CMB-S4 has been in development for six years. Through the Snowmass Cosmic Frontier planning process, experimental groups in the cosmic microwave background (CMB) and broader cosmology communities came together to produce two influential CMB planning papers, endorsed by over 90 scientists, that outlined the science case as well as the CMB-S4 instrumental concept [1, 2]. It immediately became clear that an enormous increase in the scale of ground-based CMB experiments would be needed to achieve the exciting threshold-crossing scientific goals, necessitating a phase change in the ground-based CMB experimental program. To realize CMB-S4, a partnership of the university-based CMB groups, the broader cosmology community, and the national laboratories would be needed.

The community proposed CMB-S4 to the 2014 Particle Physics Project Prioritization Process (P5) as a single, community-wide experiment, jointly supported by DOE and NSF. Following P5’s recommendation of CMB-S4 under all budget scenarios, the CMB community started in early 2015 to hold biannual workshops – open to CMB scientists from around the world – to develop and refine the concept. Nine workshops have been held to date, typically with 150 to 200 participants. The workshops have focused on developing the unique and vital role of the future ground-based CMB program.

This growing CMB-S4 community produced a detailed and influential CMB-S4 Science Book [3] and a CMB-S4 Technology Book [4]. Over 200 scientists contributed to these documents. These and numerous other reports, workshop and working group wiki pages, email lists, and much more may be found at the website <http://CMB-S4.org>.

Soon after the CMB-S4 Science Book was completed in August 2016, DOE and NSF requested the Astronomy and Astrophysics Advisory Committee (AAAC) to convene a Concept Definition Taskforce (CDT) to conduct a CMB-S4 concept study. The resulting report was unanimously accepted in late 2017.¹

One recommendation of the CDT report was that the community should organize itself into a formal collaboration. An Interim Collaboration Coordination Committee was elected by the community to coordinate this process. The resulting draft bylaws were refined at the Spring 2018 CMB-S4 community workshop, and overwhelmingly ratified on March 19th 2018, bringing the CMB-S4 Science Collaboration into being, and the first elections for the officers of the collaboration were completed by the end of April 2018. As of summer 2019 the collaboration has 198 members, 71 of whom hold positions within the organizational structure. These members represent 11 countries on 4 continents, and 76 institutions comprising 16 national laboratories and 60 universities. In parallel, the maturation of CMB-S4 as a project, strongly supported by project expertise at the national laboratories, is manifest in the organization of a preProject Development Group and now the Interim Project Office.

¹See https://www.nsf.gov/mps/ast/aaac/cmb_s4/report/CMB_S4_final_report_NL.pdf.

List of Contributors and Endorsers

Kevork Abazajian,¹ Graeme Addison,² Peter Adshead,³ Zeeshan Ahmed,⁴ Steven W. Allen,⁵ David Alonso,⁶ Marcelo Alvarez,^{7,8} Adam Anderson,⁹ Kam S. Arnold,¹⁰ Carlo Baccigalupi,¹¹ Kathy Bailey,¹² Denis Barkats,¹³ Darcy Barron,¹⁴ Peter S. Barry,¹⁵ James G. Bartlett,¹⁶ Ritoban Basu Thakur,¹⁷ Nicholas Battaglia,¹⁸ Eric Baxter,¹⁹ Rachel Bean,¹⁸ Chris Bebek,⁸ Amy N. Bender,¹² Bradford A. Benson,⁹ Edo Berger,²⁰ Sanah Bhimani,²¹ Colin A. Bischoff,²² Lindsey Bleem,¹² Sebastian Bocquet,²³ Kimberly Boddy,² Matteo Bonato,^{24,25} J. Richard Bond,²⁶ Julian Borrill,^{8,7} François R. Bouchet,²⁷ Michael L. Brown,²⁸ Sean Bryan,²⁹ Blakesley Burkhardt,^{30,31} Victor Buza,¹³ Karen Byrum,¹² Erminia Calabrese,³² Victoria Calafut,¹⁸ Robert Caldwell,³³ John E. Carlstrom,¹⁵ Julien Carron,³⁴ Thomas Cecil,¹² Anthony Challinor,³⁵ Clarence L. Chang,¹² Yuji Chinone,⁷ Hsiao-Mei Sherry Cho,⁴ Asantha Cooray,¹ Thomas M. Crawford,¹⁵ Abigail Crites,¹⁷ Ari Cukierman,⁴ Francis-Yan Cyr-Racine,¹³ Tijmen de Haan,⁸ Gianfranco de Zotti,²⁴ Jacques Delabrouille,¹⁶ Marcel Demarteau,¹² Mark Devlin,¹⁹ Eleonora Di Valentino,²⁸ Matt Dobbs,³⁶ Shannon Duff,³⁷ Adriaan Duivenvoorden,³⁸ Cora Dvorkin,¹³ William Edwards,⁸ Joseph Eimer,² Josquin Errard,¹⁶ Thomas Essinger-Hileman,³⁹ Giulio Fabbian,³⁴ Chang Feng,³ Simone Ferraro,⁸ Jeffrey P. Filippini,³ Raphael Flauger,¹⁰ Brenna Flaugher,⁹ Aurelien A. Fraisse,⁴⁰ Andrei Frolov,⁴¹ Nicholas Galitzki,¹⁰ Silvia Galli,²⁷ Ken Ganga,¹⁶ Martina Gerbino,¹² Murdock Gilchriese,⁸ Vera Gluscevic,⁴² Daniel Green,¹⁰ Daniel Grin,⁴³ Evan Grohs,⁷ Riccardo Gualtieri,³ Victor Guarino,¹² Jon E. Gudmundsson,³⁸ Salman Habib,¹² Gunther Haller,⁴ Mark Halpern,⁴⁴ Nils W. Halverson,⁴⁵ Shaul Hanany,⁴⁶ Kathleen Harrington,⁴⁷ Masaya Hasegawa,⁴⁸ Matthew Hasselfield,⁴⁹ Masashi Hazumi,⁴⁸ Katrin Heitmann,¹² Shawn Henderson,⁴ Jason W. Henning,¹⁵ J. Colin Hill,⁵⁰ Renée Hložek,⁵¹ Gil Holder,³ William Holzzapfel,⁷ Johannes Hubmayr,³⁷ Kevin M. Huffenberger,⁵² Michael Huffer,⁴ Howard Hui,¹⁷ Kent Irwin,⁵ Bradley R. Johnson,⁵³ Doug Johnstone,^{54,55} William C. Jones,⁴⁰ Kirit Karkare,¹⁵ Nobuhiko Katayama,⁵⁶ James Kerby,¹² Sarah Kernovsky,⁵⁷ Reijo Keskitalo,^{8,7} Theodore Kisner,^{8,7} Lloyd Knox,⁵⁸ Arthur Kosowsky,⁵⁹ John Kovac,¹³ Ely D. Kovetz,² Steve Kuhlmann,¹² Chao-lin Kuo,⁵ Nadine Kurita,⁴ Akito Kusaka,⁸ Anne Lahteenmaki,⁶⁰ Charles R. Lawrence,⁶¹ Adrian T. Lee,^{7,8} Antony Lewis,³⁴ Dale Li,⁴ Eric Linder,⁸ Marilena Loverde,⁶² Amy Lowitz,¹⁵ Mathew S. Madhavacheril,⁴⁰ Adam Mantz,⁵ Frederick Matsuda,⁶³ Philip Mauskopf,²⁹ Jeff McMahon,⁴⁷ Matthew McQuinn,⁶⁴ P. Daniel Meerburg,⁶⁵ Jean-Baptiste Melin,⁶⁶ Joel Meyers,⁶⁷ Marius Millea,⁶⁸ Joseph Mohr,²³ Lorenzo Moncelsi,¹⁷ Tony Mroczkowski,⁶⁹ Suvodip Mukherjee,²⁷ Moritz Münchmeyer,⁷⁰ Daisuke Nagai,²¹ Johanna Nagy,^{71,51} Toshiya Namikawa,⁷² Federico Nati,⁷³ Tyler Natoli,⁷¹ Mattia Negrello,³² Laura Newburgh,²¹ Michael D. Niemack,¹⁸ Haruki Nishino,⁴⁸ Martin Nordby,⁴ Valentine Novosad,¹² Paul O'Connor,⁷⁴ Georges Obied,¹³ Stephen Padin,¹⁵ Shivam Pandey,¹⁹ Bruce Partridge,⁴³ Elena Pierpaoli,⁴² Levon Pogosian,⁴¹ Clement Pryke,⁴⁶ Giuseppe Puglisi,⁵ Benjamin Racine,¹³ Srinivasan Raghunathan,⁷⁵ Alexandra Rahlin,⁹ Srinu Rajagopalan,⁷⁴ Marco Raveri,¹⁵ Mark Reichenadter,⁴ Christian L. Reichardt,⁷⁶ Mathieu Remazeilles,²⁸ Graca Rocha,⁶¹ Natalie A. Roe,⁸ Anirban Roy,¹¹ John Ruhl,⁷⁷ Maria Salatino,¹⁶ Benjamin Saliwanchik,²¹ Emmanuel Schaan,⁸ Alessandro Schillaci,¹⁷ Marcel M. Schmittfull,⁵⁰ Douglas Scott,⁴⁴ Neelima Sehgal,⁶² Sarah Shandera,⁴⁹ Christopher Sheehy,⁷⁴ Blake D. Sherwin,⁷⁸ Erik Shirokoff,¹⁵ Sara M. Simon,⁴⁷ Anže Slosar,⁷⁴ Rachel Somerville,^{30,31} David Spergel,^{40,30} Suzanne T. Staggs,⁴⁰ Antony Stark,¹³ Radek Stompor,¹⁶ Kyle T. Story,⁷⁹ Chris Stoughton,⁹ Aritoki Suzuki,⁸ Osamu Tajima,⁸⁰ Grant P. Teply,¹⁰ Keith Thompson,⁵ Peter Timbie,⁸¹ Maurizio Tomasi,⁸² Jesse I. Treu,⁴⁰ Matthieu Tristram,⁸³ Gregory Tucker,⁸⁴ Caterina Umiltà,²² Alexander van Engelen,²⁶ Joaquin D. Vieira,³ Abigail G. Vieregge,¹⁵ Mark Vogelsberger,⁸⁵ Gensheng Wang,¹² Scott Watson,⁸⁶ Martin White,^{8,7} Nathan Whitehorn,⁷⁵ Edward J. Wollack,³⁹ W. L. Kimmy Wu,¹⁵ Zhilei Xu,¹⁹ Siavash Yasini,⁴² James Yeck,⁸¹ Ki Won Yoon,⁵ Edward Young,⁴ Andrea Zonca¹⁰

¹UC Irvine²Johns Hopkins University³University of Illinois at Urbana-Champaign⁴SLAC⁵Stanford University⁶Oxford University⁷UC Berkeley⁸Lawrence Berkeley National Laboratory⁹Fermilab¹⁰UC San Diego

-
- ¹¹SISSA
¹²Argonne National Laboratory
¹³Harvard University
¹⁴University of New Mexico
¹⁵University of Chicago
¹⁶AstroParticle & Cosmology Laboratory
¹⁷Caltech
¹⁸Cornell University
¹⁹University of Pennsylvania
²⁰Center for Astrophysics, Harvard & Smithsonian
²¹Yale University
²²University of Cincinnati
²³LMU Munich
²⁴INAF
²⁵Italian ALMA Regional Centre
²⁶CITA
²⁷Institut d'Astrophysique de Paris
²⁸University of Manchester
²⁹Arizona State University
³⁰Center for Computational Astrophysics, Flatiron Institute
³¹Rutgers University
³²Cardiff University
³³Dartmouth College
³⁴University of Sussex
³⁵Institute of Astronomy and DAMTP, University of Cambridge
³⁶McGill University
³⁷NIST
³⁸Stockholm University
⁴⁰Princeton University
³⁹NASA Goddard Space Flight Center
⁴¹Simon Fraser University
⁴²University of Southern California
⁴³Haverford College
⁴⁴University of British Columbia
⁴⁵University of Colorado Boulder
⁴⁶University of Minnesota
⁴⁷University of Michigan
⁴⁸KEK
⁴⁹Pennsylvania State University
⁵⁰Institute for Advanced Study
⁵¹University of Toronto
⁵²Florida State University
⁵³Columbia University
⁵⁴National Research Council Canada
⁵⁵University of Victoria
⁵⁶Kavli IPMU
⁵⁷FitBit
⁵⁸UC Davis
⁵⁹University of Pittsburgh
⁶⁰Aalto University
⁶¹JPL
⁶²Stony Brook University
⁶³University of Tokyo
⁶⁴University of Washington
⁶⁵University of Groningen
⁶⁶CEA Saclay
⁶⁷Southern Methodist University
⁶⁸Institut Lagrange de Paris
⁶⁹European Southern Observatory
⁷⁰Perimeter Institute
⁷¹Dunlap Institute
⁷²National Taiwan University
⁷³University of Milano-Bicocca
⁷⁴Brookhaven National Laboratory
⁷⁵UCLA
⁷⁶University of Melbourne
⁷⁷Case Western Reserve University
⁷⁸University of Cambridge
⁷⁹Descartes Lab
⁸⁰Kyoto University
⁸¹University of Wisconsin–Madison
⁸²Università degli Studi di Milan
⁸³LAL
⁸⁴Brown University
⁸⁵Massachusetts Institute of Technology
⁸⁶Syracuse University

Executive Summary

This document presents the science case, Reference Design, and project plan for the Stage-4 ground-based cosmic microwave background experiment CMB-S4. CMB-S4 was conceived by the community during the 2013 Snowmass physics planning activity, as the path forward to realizing the enormous potential of CMB measurements for understanding the origin and evolution of the Universe, from the highest energies at the dawn of time through the growth of structure to the present day. The science case is spectacular, including the search for primordial gravitational waves as predicted from inflation, constraints on relic particles, setting the neutrino mass scale, unique and complementary insights into dark energy and tests of gravity on large scales, elucidating the role of baryonic feedback on galaxy formation and evolution, opening up a window on the transient Universe at millimeter wavelengths, and even the exploration of the outer solar system. The CMB-S4 Legacy Survey covering over half the sky with unprecedented sensitivity through the millimeter-wave band will have profound and lasting impact on Astronomy and Astrophysics and provide a powerful complement to surveys at other wavelengths, such as the Large Synoptic Survey Telescope (LSST) and the *Wide Field Infrared Survey Telescope (WFIRST)*, and others yet to be imagined.

CMB-S4 was recommended by the 2014 Particle Physics Project Prioritization Panel (P5) report *Building for Discovery: Strategic Plan for U.S. Particle Physics in the Global Context* and by the 2015 National Academies report *A Strategic Vision for NSF Investments in Antarctic and Southern Ocean Research*. The community further developed the science case in the 2016 *CMB-S4 Science Book* and surveyed the status of the technology in the 2017 *CMB-S4 Technology Book*. This work formed the foundation for the joint NSF-DOE Concept Definition Task Force (CDT), a subpanel of the Astronomy and Astrophysics Advisory Committee (AAAC), an FACA committee advising DOE, NASA, and NSF. The CDT report was enthusiastically accepted by the AAAC in October 2017. The CDT report provides clear guidance on the science goals and measurement requirements for CMB-S4, along with a strawperson instrument design, schedule, and cost.

CMB-S4 will be a joint agency program with roughly comparable support from NSF and DOE. The construction phase is expected to be funded as an NSF MREFC project and a DOE HEP MIE project. The CMB-S4 Collaboration and the DOE laboratory-based CMB-S4 pre-Project Development Group were established shortly after the CDT report was accepted. They are working together as a team to advance the CMB-S4 project. This team brings together the CMB community, as well as the considerable expertise of the national laboratories, and it is this team that has produced this report.

This report builds on the work done by the CDT and the strawperson design concept in the CDT report. We fully support the two key elements of the CDT concept for CMB-S4: (1) *it requires multiple cameras and telescopes distributed across two sites, and (2) the experiment will be undertaken by a single collaboration and run as one project*. The latter point is essential given the magnitude of the increase in science reach and complexity over existing CMB projects.

The CMB-S4 Reference Design presented here uses proven existing technology that has been developed and demonstrated by CMB experimental groups over the last decade, scaled up to unprecedented levels. The design and implementation plan addresses the considerable technical challenges presented by the required scaling-up of the instrumentation and by the scope and complexity of the data analysis and interpretation. Features of the design and plan include: superconducting detector arrays with well-understood and robust material properties and processing techniques; high-throughput mm-wave telescopes and optics with unprecedented precision and rejection of systematic contamination; full internal characterization of astronomical foreground emission; large cosmological simulations and improved theoretical modeling; and computational methods for extracting minute correlations in massive, multi-frequency data sets that include noise and a host of known and unknown signals.

Science Goals and Measurement Capabilities

Following the CMB-S4 Science Book and the CDT report we set three transformative science goals for CMB-S4.

- The first goal is to measure the imprint of primordial gravitational waves on the CMB polarization anisotropy, quantified by the tensor-to-scalar ratio r . The simplest models that naturally explain the observed departure from scale invariance of the density perturbations predict $r > 0.001$, and a particularly well-motivated subclass of these models predicts $r > 0.003$. CMB-S4 will be able to detect primordial gravitational waves for $r > 0.003$ at greater than 5σ . Such a detection would yield the first evidence of the quantization of gravity and point to inflationary physics near the energy scale associated with grand unified theories, probing energy scales far beyond the reach of the LHC or any conceivable collider experiment, and providing additional evidence in favor of the idea of the unification of forces. The measurement of the energy scale of inflation would have broad implications for many other aspects of fundamental physics, including key aspects of string theory. In the absence of a detection, the upper limit of $r < 0.001$ at 95% CL achievable by CMB-S4 would significantly advance our understanding of inflation. It would rule out large classes of inflationary models and dramatically impact how we think about the theory. To some, the remaining class of models would be contrived enough to give up on inflation altogether.
- The second goal for CMB-S4 is to detect or strongly constrain departures from the thermal history of the Universe predicted by the Standard Model of particle physics. Many well-motivated extensions of the Standard Model to higher energies predict low-mass relic particles. Departures from the standard history are conveniently quantified by the contribution of light relic particles to the effective number of relativistic species in the early Universe, N_{eff} . CMB measurements are sensitive to the contribution of relic particles to the energy density in the early Universe and therefore only depend on the interaction cross-sections of the relics with Standard-Model particles through the temperature at which the relics decouple. CMB-S4 will constrain $\Delta N_{\text{eff}} \leq 0.06$ at the 95% confidence level allowing detection of, or constraints on, a wide range of light relic particles even if they are too weakly interacting to be detected by lab-based experiments. CMB-S4 will be the most robust and precise probe of the thermal history of our Universe and will improve bounds on the decoupling temperature compared to Stage-3 CMB experiments or planned large-scale structure surveys by an order of magnitude or more, depending on the spin of the particle.
- The third science goal for CMB-S4 is to provide a unique and powerful survey of a large fraction of the sky at centimeter (cm) to millimeter (mm) wavelengths at unprecedented depth and angular resolution. Such a data set would provide enormous legacy value to the broader Astronomy and Astrophysics communities and would complement and enhance the LSST optical survey of the same region, as well as other planned and yet-to-be-imagined surveys and data from both ground- and space-based instruments. Using the signature of gravitational and electromagnetic interactions between matter and the CMB as it traverses the expanse of the Universe, mm-wave maps of sufficient depth and resolution would provide highly complementary data for investigations of dark energy, modifications to general relativity, and neutrino properties. For example, a sufficiently deep, wide, and high-resolution data set would enable two independent and competitive determinations of the sum of the neutrino masses, using weak gravitational lensing and the evolution of the number density of galaxy clusters. These data would also provide a unique and powerful probe of the influence of baryonic feedback on the formation of galaxies and clusters of galaxies. With sufficient depth and observing cadence, such a cm-to-mm-wave survey would also open an entirely new window on the transient and dynamic Universe, including mm-wave searches for orphan GRB afterglows and dwarf planets.

This third goal—the legacy cm-to-mm-wave dataset—will have the broadest benefit to both the cosmological and astronomical communities, and while it does not provide sharp measurement thresholds that drive the design of CMB-S4, it is crucial that the instrument be designed to deliver the full promise of the legacy science. For instance, the sensitivity and sky coverage required to meet the goal of $\Delta N_{\text{eff}} \leq 0.06$ at the 95% confidence level are a good match to the requirements for the legacy data, but we must also keep in mind parameters unique to this science goal such as observing cadence.

The science goals lead to the following measurement requirements.

- Low-resolution ultra-deep measurements (noise levels $< 1 \mu\text{K-arcmin}$) over an exceptionally low-foreground region covering 3% of the sky are required to meet the primordial gravitational wave goals. These measurements must have high fidelity and low contamination over a wide range of angular scales and frequencies. Large-angular-scale measurements with resolution of around 30 arcminutes and well-determined beam properties and excellent control of systematic contamination are needed to image the B -mode polarization signature of the primordial gravitational waves. Small-angular-scale measurements with resolution of order 1.5 arcminutes are needed for removing the contamination of the degree-scale B modes caused by gravitational lensing of the much stronger CMB E -mode polarization, a process referred to as “delensing.”
- High-resolution (≤ 1.5 arcminutes) deep and wide measurements at a noise level of $1 \mu\text{K-arcmin}$ over approximately 70% of the sky (60% of the sky after applying a Galactic cut) are required to meet the light relic and legacy data goals.
- Multifrequency coverage is required for foreground mitigation. As current measurements have shown, even in the cleanest regions of the sky, the rms fluctuation in Galactic foreground emission is an order of magnitude larger than the predicted B -mode fluctuations for $r = 0.001$. Simulations based on the current best knowledge of the dust and synchrotron foreground emission indicate CMB-S4 can meet its primordial gravitational wave goals using nine frequency bands spanning 20 to 270 GHz for the degree-scale measurements. Fewer bands are needed for the high-angular-resolution delensing and deep/wide surveys.

The CMB-S4 science goals are therefore met with two nested, highly complementary surveys, the deep and wide survey covering 70% of the sky and the ultra-deep survey focused on the detection of degree-angular-scale B -mode polarization generated by primordial gravitational waves, such as those predicted by inflationary models. In the context of their legacy value to the wider community, we will refer to the deep/wide and ultra-deep high-resolution surveys together as the CMB-S4 Legacy Survey.

Design Overview

This document presents the Reference Design for CMB-S4. This design meets the measurement requirements and therefore can deliver the CMB-S4 science goals. This report provides the details of the design, including the technology choices, as well as design or technology options that, if developed in time, could lead to improved performance, lower cost, or reduced risk. Building on the work done for the CDT report strawperson design, the Reference Design is supported by the extensive use of simulations based on our current understanding of the expected level and complexity of foreground emission and on noise levels scaled from existing experiments, as summarized in the appendices.

The major components of the Reference Design are as follows.

- An ultra-deep survey covering 3% of the sky, more if a primordial gravitational wave signal is detected, to be conducted over seven years using: fourteen 0.55-m refractor SATs (at 155 GHz and below) and four 0.44-m SATs (at 220/280 GHz), with dichroic, horn-coupled transition-edge sensor (TES) detectors in each SAT measuring two of the eight targeted frequency bands between 30 and 270 GHz; and one 6-m class crossed-Dragone “delensing” LAT, equipped with detectors distributed over seven bands from 20 to 270 GHz. Measurements at degree angular scales and larger made using refractor telescopes with roughly 0.5-m apertures have been demonstrated to deliver high-fidelity, low-contamination polarization measurements at these scales. The combination of the SATs with the 6-m LAT therefore provides low-resolution B -mode measurements with excellent control of systematic contamination, as well as the high-resolution measurements required for delensing. The ultra-deep survey SATs and 6-m LAT are to be located at the South Pole to allow targeted observations of the single deep narrow field, with provisions to relocate a fraction of the SATs in Chile if, for example, a high level of r is detected or unforeseen systematic issues are encountered.

The total detector count for the eighteen SATs is 153,232 with the majority of the detectors allocated to the 85- to 155-GHz bands. There are four pixel designs. The total number of science-grade 150-mm detector wafers required for eighteen SATs is 204.

The delensing LAT will have a total TES detector count of 114,432, with the majority of the detectors allocated to the 95- to 150-GHz bands. There are four pixel designs. The total number of science grade 150-mm diameter detector wafers required is 76.

- The deep and wide survey covering approximately 70% of the sky to be conducted over seven years using two 6-m crossed-Dragone LATs located in Chile, each equipped with 121,760 TES detectors distributed over eight frequency bands spanning 30 to 270 GHz. The total number of science grade 150-mm diameter detector wafers required is 152.

The total detector count for CMB-S4 is 511,184, and it requires 432 science grade wafers. This is an enormous increase over the detector count of all Stage-3 experiments combined and is required to meet the CMB-S4 science goals.

Project Plan Overview

The CMB-S4 Collaboration and the pre-Project Development Group of experienced project leaders drawn primarily from the national laboratories, jointly contributed to the development of a Work Breakdown Structure (WBS), Organization, Cost Book, Resource Loaded Schedule, and Risk Registry. The Reference Design and project baseline prepared for this document is the basis for subsequent design and project development work that is now being led by the Interim Project Office. A permanent Integrated Project Office will be established in 2020 to manage the construction phase which starts in 2021.

The CMB-S4 Project is a collaborative project, with the scientific Collaboration members serving in technical leadership roles in the project. This is similar to many successful projects including IceCube, ATLAS, and CMS. The project office is responsible for forming partnerships with key stakeholder institutions including DOE National Laboratories, universities, and potential collaborating observatories/projects such as the Simons Observatory, South Pole Observatory, and the CCAT-prime project. Partnerships are also expected to include foreign institutions participating in the CMB-S4 Collaboration and contributing to the CMB-S4 Project.

The CMB-S4 project total estimated cost is currently \$591.6M (fully loaded and escalated to the year of expenditure) including a 35% contingency budget. The cost estimate is the full cost, i.e., it does not take

credit for use of any legacy infrastructure or for contributions from collaborating institutions supported by private and international partners, e.g., large-aperture telescopes currently under construction in Chile as part of the Simons Observatory, and large- and small-aperture telescopes proposed by international collaborators. In-kind contributions delivered by private and international partners are expected and would reduce the total cost to NSF and DOE. It is estimated that the value of in-kind contributions could reduce the total cost of the CMB-S4 project by 20–25%.

The cost contingency estimate was constructed using input from subject matter experts with previous experience in previous CMB experiments and similar NSF MREFC projects and DOE MIE projects. As the design, cost estimates, and schedules mature the contingency as a percentage of the base cost estimate is expected to decrease to 30% or less.

The project has developed a task based detailed resource loaded schedule which was reviewed by an external panel of experts in December 2018. The schedule has 1110 activities, 1928 relationships, 5 Level 1, 20 Level 2 and 299 Level 3 Milestones. The estimates follow the guidance in the NSF Large Facilities Manual, NSF 17-066 and the Project Management for the Acquisition of Capital Assets, DOE 413.3b.

A CMB-S4 Risk and Opportunity Management Plan describes the continuous risk and opportunity management process implemented by the project, consistent with DOE O413.3B, “Project Management for the Acquisition of Capital Assets,” and the NSF 17-066, “NSF Large Facilities Manual.” The plan establishes the methods of assessing CMB-S4 project risk and opportunities for all subsystems as well as the system as a whole. The project is working on mitigations to ensure that the highest risks identified in the risk registry are lowered to reasonable levels on a time scale consistent with the overall project timeline.

The basic operations model for CMB-S4 will be observations with multiple telescopes and cameras distributed across two sites, with observing priorities and specifications optimized for the CMB-S4 science goals, and data from all instruments shared throughout the entire CMB-S4 collaboration. The operations cost is based on a preliminary bottom-up estimate that includes management, site staff, utilities, instrument maintenance, data transmission, data products, pipeline upgrades, collaboration management, and science analysis. The annual operations cost is \$32M in 2019 dollars, excluding 20 FTE/year of scientist effort supported by DOE research funds.

Data Analysis and Data Release Plan Overview

Many of the key data analysis tools that will be needed for the scientific exploitation of the CMB-S4 data are in fact needed to optimize the final design of the instrument—simulations are particularly important. These tools are included in the Data Management section of the Reference Design and are therefore included as part of the construction phase. We also outline an operations plan that includes developing the science analysis tools that will be needed for the scientific exploitation of the data.

The production and release of the CMB-S4 program data deliverables is an integral part of operations, and the plan is for this effort to be supported by NSF and DOE. These deliverables include the data products from the powerful CMB-S4 Legacy Survey that covers 70% of the sky. The data products will include maps of the temperature and polarization in each of the nine CMB-S4 observing bands, the projected mass map reconstructed from CMB lensing, and source catalogs (e.g., galaxy clusters, active and dusty galaxies, and transient events). These products will be made available through a series of releases to the entire Astronomy and Astrophysics community and will provide a powerful complement to the surveys that will be available or planned at other wavelengths, such as LSST and *WFIRST*. The CMB-S4 maps and catalogs will also

provide exciting targets for detailed follow-up study by the Atacama Large Millimeter/submillimeter Array (ALMA), the *James Webb Space Telescope (JWST)*, and other facilities.

The key science analyses, such as the searches for primordial gravitational waves and light relics, will be pursued by the CMB-S4 collaboration with support shared by NSF and DOE. Non-key science analyses, for example studies of the Legacy Survey maps and catalogs, will be carried out by the wider community of laboratory and university scientists with support expected to be provided by individual NSF and DOE awards.

Contents

1	Science Case	1
1.1	Summary of science goals	2
1.2	Primordial gravitational waves and inflation	10
1.3	The dark Universe	25
1.4	Mapping matter in the Cosmos	50
1.5	The time-variable millimeter-wave sky	65
2	Science and Measurement Requirements	71
2.1	Primary science goals	71
2.2	Flowdown to measurement requirements	72
3	Instrument Overview	81
3.1	Requirements and design drivers	81
3.2	Instrument configuration	81
3.3	Reference design summary	83
3.4	Manufacturability and reliability	83
3.5	Design approach	85
3.6	Options	85
4	Reference Design	87
4.1	Sites	87
4.2	Large telescopes	93
4.3	Cameras for large telescopes	104
4.4	Small telescopes and cameras	114
4.5	Detectors and readout	127
4.6	Data acquisition and experiment control	140
4.7	Data management	146

4.8	Integration and commissioning	155
5	Science Analyses	163
5.1	Reconstruction of the CMB lensing potential	163
5.2	Power-spectrum estimation	164
5.3	Component maps and cross-correlation with other surveys	166
5.4	Parameter estimation	166
5.5	Creating source and cluster catalogs	166
5.6	Transient and time-domain analyses	167
5.7	Sky simulations	167
5.8	Implementation	168
5.9	Validation and verification	168
6	Project Overview	169
6.1	Introduction	169
6.2	Scope, work breakdown structure, and cost	169
6.3	Cost, schedule, and risk	176
6.4	Risks and opportunities	179
6.5	Value management—alternative analysis and selection	180
6.6	Operations plan	182
6.7	R&D and pre-conceptual design	182
	Appendices	187
A	Science Forecasting	187
A.1	Ultra-deep field targeting the degree-scale signature of gravitational waves	187
A.2	Deep and wide field for measuring light relic density	207
A.3	Mapping baryons: angular resolution and allocation of detectors across frequency for the LATs	212
A.4	Gamma-ray bursts, sensitivity, angular resolution, and cadence	214
B	Options	219
B.1	Re-use of existing telescopes and site infrastructure	219

B.2	Readout	219
B.3	Kinetic-inductance detectors	221
B.4	Coupling to detectors	222
B.5	Commercial detector fabrication	222
B.6	Other developments	223
C	List of Abbreviations	225
	References	231

Science Case

CMB-S4 will probe fundamental physics and astrophysics with a millimeter-wavelength survey of unprecedented depth over a large sky area. These measurements will represent a leap forward in the study of the cosmic microwave background (CMB), which in the half century since its first discovery, has again and again transformed our understanding of the early Universe.

CMB experiments on the ground, on balloons, and in space have provided conclusive evidence that our Universe evolved from an early hot, dense state. They have determined the age and the composition of our Universe with percent-level precision. They have provided the strongest evidence that dark matter cannot consist of non-luminous baryonic matter. They have measured the polarization of the CMB, and shown us in a model-independent way that the fluctuations we see in CMB intensity were already present at the time of recombination. To be consistent with Einstein's theory of general relativity, some process must have generated these fluctuations long before the moment when our Universe became filled with a hot and dense plasma; i.e., long before the "hot Big Bang" began.

Measurements of the CMB continue to provide us with a remarkable opportunity to study the Universe over its history. Although CMB observations have already harvested nearly all the information accessible in the primary temperature anisotropies, current experiments have only begun to make precise measurements of the CMB polarization. The secondary anisotropies, such as the weak gravitational lensing of the CMB by large-scale structure and scattering by the thermal and kinematic Sunyaev-Zeldovich (tSZ and kSZ) effects, also hold tremendous promise to improve our understanding of astrophysics and cosmology.

CMB-S4 will push beyond these new frontiers in CMB science. It will exploit the enormous potential of CMB measurements to once again transform our understanding of the early Universe and of particle physics. In doing so, it will fulfill the goals set out in the 2010 Astronomy and Astrophysics Decadal Survey and the 2014 report of the Particle Physics Project Prioritization Panel. Specifically, it will search for primordial gravitational waves and for light relic particles. CMB-S4 will also constrain neutrino properties and provide critical measurements of the evolution of cosmic structure from the early Universe to the present day, advancing a key goal of the 2010 Decadal Survey.

CMB-S4 will also be built as an astronomical survey machine. To date, CMB experiments have produced the only wide-area mm-wave sky surveys, providing new insights in astrophysics. CMB-S4 will image a large fraction of the sky, including much of the Galaxy, with high-fidelity measurements of intensity and linear polarization, on scales from arcminutes to many degrees.

These wavelengths reveal a wide variety of astrophysical effects. Compton scattering of mm-wave light on electrons in the late Universe provides key information about the ionization history of the Universe and thermal evolution of diffuse gas. Individual galaxy clusters can be detected through this Compton scattering, and large samples of galaxies can be detected through either thermal emission from their dust or by synchrotron emission from active galactic nuclei (AGN). By surveying large areas of sky, these surveys are able to probe large volumes and detect large samples of objects that are too rare to be found in smaller-area surveys.

CMB-S4 will provide measurements on a wide variety of timescales. Typical CMB experiments survey the same patch of sky repeatedly to build up the signal-to-noise ratio. This leads to sampling on timescales

as fast as multiple times per second, with multiple observations extending over several years. CMB-S4's unprecedented depth means that it should observe many objects in the mm-wave transient and variable sky, including gamma-ray bursts and blazars. It should also observe the thermal emission from moving Solar System objects such as asteroids, dwarf planets, and unknown bodies in the outer Solar System.

CMB-S4 will generate measurements with broad and diverse scientific applications. The following sections will give a tour of some of these applications and touch on many ideas in particle physics, cosmology, astrophysics, and astronomy.

1.1 Summary of science goals

We have organized the rich and diverse set of CMB-S4 scientific goals into four themes:

1. *primordial gravitational waves and inflation;*
2. *the dark Universe;*
3. *mapping matter in the cosmos;*
4. *the time-variable millimeter-wave sky.*

The first two science themes relate to fundamental physics and have been primary drivers of the CMB-S4 concept from the beginning. The final two themes have emerged through engagement with the wider astronomical community and studies of broader scientific opportunities made possible by a millimeter-wave survey of unprecedented depth and breadth, and together we refer to them as our Legacy Survey themes. We summarize the four science themes in this short section and present the full science case in the remainder of the chapter.

Primordial gravitational waves and inflation. We have a historic opportunity to open up a window to the primordial Universe. If the predictions of some of the leading models for the origin of the hot big bang are borne out, CMB-S4 will detect primordial gravitational waves. This detection would provide the first evidence for the quantization of gravity, reveal new physics at the energy scale of grand unified theories, and yield insight into the symmetries of nature and possibly into the properties of quantum gravity. Conversely, a null result would rule out large classes of models and put significant strain on the leading paradigm for early-Universe cosmology, the theory of cosmic inflation.

Cosmic inflation refers to a period of accelerated expansion prior to the hot big bang. During this epoch, quantum fluctuations were imprinted on all spatial scales in the cosmos. These fluctuations seeded the density perturbations that developed into all the structure in the Universe today. While we cannot yet claim with high confidence that the Universe underwent cosmic inflation, the simplest models of inflation are exceptionally successful at matching the data. Specifically, these predictions include small mean spatial curvature and initial density perturbations drawn from a nearly Gaussian distribution with a variance that is slightly larger on large scales than on small scales. Each of these predictions has been verified to high precision.

Tantalizingly, the observed (weak) scale-dependence of the amplitude of density perturbations has quantitative implications for the detection of primordial gravitational waves. In the simplest class of models, the amplitude of primordial gravitational waves is comparable to the deviation from scale invariance,

quantified by $n_s - 1$. However, *all* inflation models that naturally explain the observed $n_s - 1$ value, and that also have a characteristic scale larger than the Planck mass, generate primordial gravitational waves above the 95% confidence upper limit that CMB-S4 can set (see Fig. 7 for example models). A well-motivated sub-class within this set of models is detectable by CMB-S4 at 5σ .




Figure 1. Forecast of CMB-S4 constraints in the n_s - r plane for a fiducial model with $r = 0.003$. Also shown are the current best constraints from a combination of the BICEP2/Keck Array experiments and Planck [5]. Models that naturally explain the observed departure from scale invariance separate into two viable classes: monomial and plateau. The monomial models ($V(\phi) = \mu^{4-p}\phi^p$) are shown for three values of p as blue lines for $47 < N_* < 57$ (with the spread in N_* reflecting uncertainties in reheating, and smaller N_* predicting lower values of n_s). This class is not completely ruled out by the data, but is disfavored. The plateau models divide into those with plateaus near the scalar field origin, for which we include the quartic hilltop (green band) as an example, and those with plateaus away from the origin, for which we include the \tanh form (gray band) as an example, as this form arises in a sub-class of α -attractor models [6]. Some particular realizations of physical models in the plateau class are also shown: the Starobinsky model [7] and Higgs inflation [8] (small and large orange filled circles, respectively) and fibre inflation [9] (purple line). The differing choices of N_* for Higgs and Starobinsky reflect differing expectations for reheating efficiency.

Because the Universe has expanded by a tremendous amount since the time when primordial perturbations were imprinted, CMB observations can probe physics at extraordinarily small length scales, *up to 10^{10} times smaller than those probed in terrestrial particle colliders*. The CMB provides a unique window to test new phenomena at these length scales. The observational requirement is also clear: we must measure the polarization to high precision on angular scales from several arcminutes to degrees.

Primordial gravitational waves source an odd-parity fluctuation pattern in the polarization across the sky, called “ B modes” by analogy with electromagnetism (while scalar density perturbations source only an even parity polarization pattern, called “ E modes”). To measure primordial B modes, we must observe at multiple frequencies to remove Galactic foreground contamination and also measure small angular scales to remove the (non-primordial) gravitational lensing-generated polarization B modes. The CMB-S4 reference design has sufficient sensitivity to detect or tightly constrain the degree-scale B modes generated by gravitational waves in many models, and to measure the amount of gravitational waves (tensor perturbations), detecting or setting an upper limit on the tensor-to-scalar ratio r . With an order of magnitude more detectors than precursor observations, and exquisite control of systematic errors, we will improve upon limits from previous

observations by a factor of 5, allowing us to either detect primordial gravitational waves or rule out a broad class of models with a super-Planckian characteristic scale.

Complementary to the search for gravitational waves, CMB-S4 will provide exquisite measurements of primordial *density* fluctuations via *E* modes. The polarization sensitivity will surpass current measurements of *E*-mode polarization, which are far from being sample-variance-limited. Because polarization has lower Galactic foregrounds than temperature, we will improve measurements across the angular scales already observed in temperature, and push to yet smaller angular scales. These polarization measurements will significantly extend and enhance searches for non-power-law features in the primordial power spectrum, small variations in the equation of state, and small departures from Gaussianity. The CMB is the most robust observable for non-Gaussianities to date and CMB-S4 will provide the tightest constraints on the most compelling signatures, improving the constraints from the *Planck* satellite. Non-Gaussianities can also arise in models with undetectably small gravitational wave production, and provide an independent handle on the early Universe. Non-Gaussianity can also be measured via cross correlation of the CMB-S4 mass map with galaxy surveys, a measurement that has the potential to rule out a large class of inflationary models.

The dark Universe. In the standard cosmological model, about ninety five percent of the mass-energy density of the Universe is in dark matter and dark energy. With CMB-S4 we can address numerous questions about these dark ingredients, such as: How is the dark (plus baryonic) matter distributed on large scales? Does the dark matter have non-gravitational interactions with baryons? And are there additional unseen components beyond dark matter and dark energy?

Light relic particles are one well-motivated possibility. These additional components to the dark sector are light particles produced in the early Universe, and are sometimes referred to as “dark radiation.” Many extensions of the Standard Model (SM) predict such light relics, including axion-like particles and sterile neutrinos [10, 11, 12]. For large regions of the unexplored parameter space, these light particles are thermalized in the early Universe. To date, CMB observations by the *Planck* satellite can probe light particles that departed from equilibrium (“froze out”) as early as the first 50 micro-seconds of the Universe. With CMB-S4 we can push back this frontier by over a factor of 10,000, to the first fractions of a nanosecond.

CMB-S4 will achieve sensitivity to relics that froze out well before the quark-hadron phase transition (the epoch when the Universe cooled sufficiently that quarks became locked into hadrons like neutrons and protons). The contribution of light relics to the energy density leads to observable consequences in the CMB temperature and polarization anisotropy. This is often parameterized with the “effective number of neutrino species,” N_{eff} . The collective influence of the three already-known light relics (the three families of neutrinos) has already been detected at high significance. Current data are only sensitive enough to detect additional relics that froze out after the quark-hadron transition, and Stage-3 CMB experiments can only push somewhat into that epoch, so CMB-S4’s ability to probe times well before that transition is a major advance.

In addition to precise constraints on N_{eff} , CMB-S4 will give an *independent* high-precision measurement of the primordial helium abundance, Y_p . This is particularly useful since Y_p is sensitive to N_{eff} a few minutes after reheating, while the CMB power spectrum is affected by N_{eff} prior to recombination, about 370,000 years later. Measuring the radiation content at these well-separated times provides a window onto any non-trivial evolution in the energy density of radiation in the early Universe. Furthermore, N_{eff} and Y_p are sensitive to neutrino physics and physics beyond the Standard Model in related, but different ways, allowing even finer probes of particle physics, especially in the neutrino and dark sectors.

CMB-S4 will also enable a broader exploration of the dark Universe in combination with other probes, often significantly enhancing them by breaking their intrinsic degeneracies. It will improve or detect various




Figure 2. *CMB-S4-enabled exploration of light relics, axions, neutrino mass, and dark matter properties. In each case, there is a window in the mass of the relevant particle where the CMB is particularly sensitive. Each such region is shown in a color (or colors) representing the observable(s) that drives the constraint. The primary CMB anisotropies at high- ℓ (blue) is particularly sensitive to light relics, as discussed in Section 1.3.1, and properties of the dark matter discussed in Section 1.3.4. Low- ℓ modes (green) most directly impact constraints on dark matter annihilation. CMB lensing reconstruction (red) is a sensitive probe of matter in the late Universe, particularly effects that suppress clustering power. Axion dark matter can create additional modulations of the CMB polarization angles (yellow) through their coupling to photons. A detection of primordial gravitational waves (pink) would severely constrain the QCD axion because of the implied high scale of inflation.*

possibilities for the dark matter properties beyond the simplest cold-dark-matter scenario, as described in Fig. 2. It will add to dark energy constraints through precision measurements of the primordial power spectrum (where dark energy physics enters through projection effects), through precision measurements of the lensing convergence power spectrum, through the CMB-lensing-derived mass calibration of galaxy clusters, and through CMB lensing tomography.

Mapping matter in the cosmos. Matter in the Universe can be sorted into two categories, “normal” or “baryonic” matter that is described by the Standard Model of particle physics and “dark” matter that has only been observed to interact gravitationally. Observations indicate there is more than five times more dark matter than baryonic matter, and most of the baryonic matter is in the form of hot ionized gas rather than cold gas or stars. CMB-S4 will be able to map out normal and dark matter separately by measuring the fluctuations in the total mass density (using gravitational lensing) and the ionized gas density (using Compton scattering).

Gravitational lensing of background sources by intervening gravitational potentials leads to detectable distortions that can be used to reconstruct fluctuations in the mass density. Virtually all of the the density fluctuations within the observable Universe leave an imprint when the CMB is the background source. The map resulting from CMB lensing reconstruction will be wide-area, highly sensitive, and extremely well-calibrated.




Figure 3. Example lensing-deflection maps (top) and thermal SZ (Compton y -maps, bottom) reconstructed with *Planck* (left) and CMB-S4 data (right). The center panels show a 25 deg^2 patch of the all-sky lensing-deflection field in the *WebSky* simulations (top) and Compton y (bottom). The left panels show Wiener-filtered maps of the signal after adding (Gaussian) noise and residual foregrounds with levels corresponding to the *Planck* 2018 lensing-deflection map (top) and the *Planck* 2015 “Needlet Internal Linear Combination” t SZ map (bottom). The right panel shows analogous Wiener-filtered maps with noise expected for CMB-S4 (top) and residual foregrounds determined by the CMB-S4 + *Planck* t SZ noise power spectrum in Fig. 20. The significantly higher fidelity of the CMB-S4 reconstruction is evident.

On its own, we can use this map to precisely measure the amplitude of large-scale structure at intermediate redshifts, with important applications to dark energy, modified gravity, and studies of neutrino masses. In concert with catalogs of objects, we can use this map to weigh samples (of e.g., galaxies and galaxy clusters) to as high a redshift as such sources can be found. The technique of CMB lensing tomography, enabled by CMB-S4 and galaxy catalogs from—for example—the Large Synoptic Survey Telescope (LSST), will allow for the creation of mass maps in broad redshift slices out to redshifts as high as 5, making possible new precision tests of cosmology. Such results explore the connection between visible baryons and the underlying dark-matter scaffolding. In conjunction with cosmic-shear surveys (e.g., LSST) that measure the low-redshift mass distribution, a map of the high-redshift mass distribution can be constructed, gaining new insight into the first galaxies. By calibrating cluster masses at high redshift, the abundance of galaxy clusters can be used as an additional probe of dark energy and neutrino masses.




Figure 4. *CMB-S4 survey area, in Galactic coordinates, with the ecliptic plane marked as the solid line.*

Low-redshift structure acts to lens both the CMB and the images of intermediate- z galaxies. Detailed comparisons provides a valuable cross-check on galaxy shear measurement calibration and enables geometric tests using the longest possible lever arm.

Most of the baryons in the late Universe are believed to be in a diffuse ionized plasma that is difficult to observe. This ionized plasma can leave imprints in the CMB through Compton scattering, the so-called Sunyaev-Zeldovich effects. The two leading variants are either a spectral distortion from hot electrons interacting with the relatively cold CMB (thermal SZ or tSZ), or a general redshift or blueshift of the scattered photons due to coherent bulk flows along the line of sight (kinematic SZ or kSZ).

The nature of the scattering makes the tSZ independent of redshift. With the deep and wide field covering a large amount of volume and the ultra-deep field imaging lower-mass clusters, CMB-S4 will be an effective probe of the crucial regime of $z \gtrsim 2$, when galaxy clusters were vigorously accreting new hot gas while at the same time forming the bulk of their stars. The CMB-S4 catalog will be more than an order of magnitude larger than current catalogs based on tSZ or X-ray measurements, and will contain an order of magnitude more clusters at $z > 2$ than will be discovered with Stage 3 CMB experiments. CMB-S4 will also measure the diffuse tSZ signal everywhere on the sky and make a temperature-weighted map of ionized gas which can be used to measure the average thermal pressure profiles around galaxies and groups of galaxies.

CMB-S4 will measure the kSZ effect, which will be combined with data from other surveys to make maps of the projected electron density around samples of objects. Applications of these maps include measuring ionized gas as a function of radius, directly constraining the impact of feedback from active galactic nuclei and supernovae on the intergalactic medium and constraining theories of modified gravity with the bulk flow amplitude as a function of separation.

Even without overlapping galaxy catalogs, the kSZ signal adds extra small-scale power that is significantly non-Gaussian. Some of this excess power and non-Gaussianity will be coming from the relatively local Universe where the galaxy catalogs overlap, but there should also be a substantial signal coming from the epoch of reionization. By directly probing the ionized gas distribution, these measurements are completely

complementary to the measurements of the neutral gas that can be obtained with redshifted Ly- α or redshifted 21-cm studies.

In the course of its survey, CMB-S4 will catalog the emission from galaxies in the mm-wave band, including AGN and dusty star-forming galaxies. The matter in our own Galaxy will also be mapped in intensity and linear polarization over a large fraction of the sky, with extracted images of synchrotron and dust emission with high fidelity on scales ranging from arcminutes to several degrees.

The time-variable millimeter-wave sky. There have been relatively few studies of the variable sky at mm-wavelengths, with only one systematic survey done to date (by a CMB experiment [13]). Known contributors to the time-varying sky are transient events, Solar System objects, and variable AGN (especially blazars).

Targeted follow-up observations of gamma-ray bursts, core-collapse supernovae, tidal disruption events, classical novae, X-ray binaries, and stellar flares have found that there are many transient events with measured fluxes that would make them detectable by CMB-S4. A systematic survey of the mm-wave sky with a cadence of a day or two over a large fraction of the sky would be an excellent complement to other transient surveys, filling a gap between radio and optical searches. Gamma-ray burst afterglows can be detected within a few hours of the burst in many cases, and there is a possibility of capturing mm-wave afterglows that have no corresponding gamma-ray trigger either from the geometry of relativistic beaming and/or from sources being at very high redshift.

Thermal emission from planets, dwarf planets, and asteroids have been measured at these wavelengths, and since such sources have high proper motions, they should be easily differentiated from the relatively stationary extrasolar sky. Using the thermal emission rather than reflected light has several complementary aspects: the fall-off with distance is less severe, providing unique information on possible large objects in the distant reaches of the Solar System; the physical information available is also very different, measuring long-wavelength emissivity rather than optical reflectivity; and with long-time baselines for observation it will be possible to build up rotation curves for a large number of objects, enabling detailed comparison with the optical and infrared versions.

CMB-S4 will play an active role in multi-messenger astronomy. Accreting black holes are known to be highly variable. A CMB survey can provide a long baseline with high time sampling in both intensity and linear polarization. This will create a mm-wave archive for multi-messenger astronomy, in particular for future blazars that are discovered to be sources of high-energy neutrinos (such as the blazar TXS 0506+056, thought to be associated with the IceCube event IC170922A). With a large catalog of time-variable blazars, it will be possible to derive detailed variability statistics over several years with nearly daily monitoring for both the detected objects and the sources that are observed to *not* be neutrino sources. Additionally, the natural wide-area nature of the survey will make it straightforward to search for gravitational wave sources that happen to be poorly localized. Although the first binary neutron star merger, GW170817, was not detected at millimeter wavelengths, this was likely due to the low density of the merger environment [14]. There is reason to expect, based on observations of short gamma-ray bursts, that at least some mergers can occur in denser environments, which will make them fainter in the optical, *but* enhance their mm emission [15].




Figure 5. Filled circles show 5σ limiting magnitudes for transient surveys (ASAS-SN, Zwicky Transient Facility, Large Synoptic Sky Survey, Australia SKA Pathfinder, all in black, CMB-S4 weekly flux limit in red) over a large fraction of the sky. Diagonal lines indicate constant νS_ν , lines separated by factors of 100: green shows a Neptune-mass planet at 500 AU; and blue lines show SEDs corresponding to a quasar or blazar, normalized to be representative in flux of the population measurable in daily CMB-S4 maps. Open circles show the coadded 5σ depth for past and near-future surveys that cover large fractions of the sky, with CMB-S4 again shown in red. The ultra-deep CMB-S4 survey (not shown) is a factor of 2-6 deeper than the deep and wide survey, depending on the wavelength.

1.2 Primordial gravitational waves and inflation

The first science theme addresses physical processes in the very early Universe responsible for the origin of all structure.

CMB-S4 has the potential to detect a pristine relic from the primordial Universe, and to unravel the mechanism responsible for the generation of primordial perturbations. Primordial gravitational waves at the time of recombination leave an imprint on the polarization pattern of the cosmic microwave background.

Cosmologists widely regard inflation, a period of nearly exponential expansion, as the most compelling paradigm for the very early Universe. Many of the predictions of the simplest models of inflation have been verified. For example, we have detected at high significance the predicted small departure from scale invariance of the density perturbations. The sign and size of the departure from scale invariance implies that measurements of or constraints on the gravitational wave amplitudes by CMB-S4 will provide key information about the early Universe, as explained below.

Because gravitational waves are fluctuations in the spacetime metric rather than the density, such a detection would open a new window into the early Cosmos and transform our understanding of fundamental physics. For the foreseeable future, precise measurements of CMB polarization are our only way to detect primordial gravitational waves. Even an upper limit from CMB-S4 would provide invaluable insights into the first moment of our Universe.

In addition, CMB-S4 will provide unprecedented constraints on primordial fluctuations in general. The simplest models of inflation predict that the initial fluctuations in the number densities of the various components (dark matter, baryons, photons, and neutrinos) are completely correlated. Such fluctuations are adiabatic. Other models predict fluctuations in some components that are anti-correlated with other components. Such fluctuations are “isocurvature” fluctuations, and they could come in a wide variety of possible forms. To date, no experiment has detected isocurvature fluctuations, and CMB-S4 will substantially improve upon current constraints.

Inflation produces nearly Gaussian fluctuations, but many models make a definite prediction for a small amount of non-Gaussianity. Discovery of that non-Gaussianity would present a true breakthrough, on par with a detection of primordial gravitational waves. *Planck* [16, 17] showed that the non-Gaussianity of the primordial fluctuations is indeed small; however, impressive as these constraints are, they do not yet reach the theoretically motivated thresholds that CMB-S4 can achieve. CMB-S4 will have multiple ways to measure non-Gaussianity, particularly in conjunction with large-scale structure measurements, opening an important range of parameter space.

1.2.1 Primordial gravitational waves

Stars and galaxies formed through gravitational collapse from primordial density perturbations that were generated in the very early Universe. The detailed mechanism that created these density perturbations remains unknown, but one may expect it to not only produce density perturbations but also gravitational waves. The amplitude of the gravitational wave signal provides crucial information about the processes responsible for the origin of structure. By measuring or constraining the amplitude of primordial gravitational waves CMB-S4 will provide invaluable information about the origin of all structure in our Universe.

Our quantitative discussion of primordial gravitational waves begins with the line element for a perturbed Friedmann-Lemaître-Robertson-Walker (FLRW) universe, which in the ADM (Arnowitt-Deser-Misner) formalism [18] is given by

$$\begin{aligned} ds^2 &= -N^2 dt^2 + h_{ij}(dx^i + N^i dt)(dx^j + N^j dt), \\ h_{ij} &= a^2(t)[e^{2\mathcal{R}}\delta_{ij} + \gamma_{ij}]. \end{aligned} \quad (1.1)$$

The Hamiltonian and momentum constraints determine the lapse N and the shift N^i in terms of the dynamical scalar and tensor degrees of freedom \mathcal{R} and γ_{ij} . In general h_{ij} may also contain vector perturbations, but these rapidly decay and can be neglected unless they are actively sourced, e.g., by cosmic strings.

Because the equations of motion for the fluctuations are invariant under translations, and the fluctuations are small enough to be treated perturbatively, it is convenient to work with the Fourier transforms

$$\mathcal{R}(t, \mathbf{x}) = \int \frac{d^3k}{(2\pi)^3} \mathcal{R}(t, \mathbf{k}) e^{i\mathbf{k}\cdot\mathbf{x}} + \text{h.c.} \quad \text{and} \quad \gamma_{ij}(t, \mathbf{x}) = \sum_{\lambda} \int \frac{d^3k}{(2\pi)^3} \gamma_{\lambda}(t, \mathbf{k}) e_{ij}(\mathbf{k}, \lambda) e^{i\mathbf{k}\cdot\mathbf{x}} + \text{h.c.}, \quad (1.2)$$

where $e_{ij}(\mathbf{k}, \lambda)$ is the transverse-traceless polarization tensor for the graviton, λ labels the polarization states of the gravitational waves, and ‘h.c.’ stands for the Hermitian conjugate.

In a universe that is dominated by matter or radiation the expansion rate $H = \dot{a}/a$ decays more rapidly than the momentum k/a redshifts. So at early times the modes are ‘outside the horizon,’ $k \ll aH$, and are time-independent. As modes ‘enter the horizon,’ $k \gg aH$, the modes oscillate.

The statistical properties of the scalar and tensor fluctuations, \mathcal{R} and γ_{λ} , at times when the modes are outside the horizon, provide the link between late-time observations and the primordial era. For a universe that is statistically homogeneous and isotropic, and in which the primordial fluctuations are adiabatic and Gaussian, the information about the statistical properties is contained entirely in the 2-point correlation functions

$$\begin{aligned} \langle \mathcal{R}(\mathbf{k}) \mathcal{R}(\mathbf{k}') \rangle &= (2\pi)^3 \delta^3(\mathbf{k} + \mathbf{k}') \frac{2\pi^2}{k^3} \Delta_{\mathcal{R}}^2(k), \\ \langle \gamma_{\lambda}(\mathbf{k}) \gamma_{\lambda'}(\mathbf{k}') \rangle &= (2\pi)^3 \delta_{\lambda\lambda'} \delta^3(\mathbf{k} + \mathbf{k}') \frac{2\pi^2}{k^3} \frac{1}{2} \Delta_{\gamma}^2(k), \end{aligned} \quad (1.3)$$

where the factor of 1/2 in the last line accounts for the fact that the measured power includes contributions from each of the two graviton polarizations.

All current observations are consistent with $\Delta_{\gamma}^2(k) = 0$ and angular power spectra given by

$$C_{XX,\ell} = \int \frac{dk}{k} \Delta_{\mathcal{R}}^2(k) \left| \int_0^{\tau_0} d\tau S_X(k, \tau) u_{X,\ell}(k(\tau_0 - \tau)) \right|^2, \quad (1.4)$$

where $S_X(k, \tau)$ with $X = T, E$ are source functions that encode the evolution of the modes, and $u_{X,\ell}$ are functions that encode the geometry of the Universe. In a spatially flat universe $u_{T,\ell}$ is, for example, a spherical Bessel function, $u_{T,\ell} = j_{\ell}$. The spectrum of primordial perturbations is nearly scale-invariant, i.e.,

$$\Delta_{\mathcal{R}}^2(k) = A_s \left(\frac{k}{k_*} \right)^{n_s - 1 + \frac{1}{2} \left. \frac{dn_s}{d \ln k} \right|_{k=k_*} \ln(k/k_*) + \dots} \quad \text{with} \quad n_s \approx 1 \quad \text{and} \quad \left. \frac{dn_s}{d \ln k} \right|_{k=k_*} \approx 0. \quad (1.5)$$

If primordial gravitational waves are present, they also contribute to the angular power spectra of temperature and polarization anisotropies. Given that the scalar perturbations are nearly scale invariant, we may expect the same for the tensor perturbations and hence parameterize their power spectrum as

$$\Delta_\gamma^2(k) = A_r \left(\frac{k}{k_*} \right)^{n_t + \frac{1}{2} \left. \frac{dn_t}{d \ln k} \right|_{k=k_*} \ln(k/k_*) + \dots} \quad \text{with} \quad n_t \approx 0 \quad \text{and} \quad \left. \frac{dn_t}{d \ln k} \right|_{k=k_*} \approx 0. \quad (1.6)$$

We can then define the tensor-to-scalar ratio as $r = A_t/A_s$. At linear order in perturbation theory only the tensor perturbations generate B -mode polarization, and a search for primordial gravitational waves can be cast as a search for B -mode polarization in the cosmic microwave background radiation. Angular power spectra for anisotropies in the CMB temperature and polarization are shown in Fig. 6. B -mode anisotropies created by gravitational waves are shown for two representative values of the tensor-to-scalar ratio and are significantly fainter than both temperature and E -mode anisotropies, highlighting the experimental precision necessary to detect them.

Primordial gravitational waves are not the only source of degree-scale B modes, since weak gravitational lensing of the cosmic microwave background by large-scale structure converts E - to B -mode polarization. This lensing contribution dominates over the primordial signal for $r \lesssim 0.01$. CMB-S4 will rely on precision measurements of E and B modes on smaller angular scales to remove 90% of this lensing contribution.

In addition, as we will discuss, emission from thermal dust and relativistic electrons in our galaxy lead to B -mode polarization in the microwave sky, and instrumental effects convert temperature or E -mode polarization to B -mode polarization. Both must be controlled to unprecedented levels.

Implications for inflation Inflation, a period of accelerated expansion of the early Universe, is the leading paradigm for explaining the origin of the primordial density perturbations that grew into the CMB anisotropies and eventually into the stars and galaxies we see around us. Accelerated expansion requires matter with an energy density that dilutes relatively slowly as the Universe expands. In inflationary models, such an energy density is usually obtained via the introduction of a new field ϕ , called the inflaton. In the simplest scenarios its Lagrangian density is given by

$$\mathcal{L} = -\frac{1}{2} g^{\mu\nu} \partial_\mu \phi \partial_\nu \phi - V(\phi), \quad (1.7)$$

where $V(\phi)$ is the potential energy density.

The overall evolution of the Universe is well-described by a FLRW line element

$$ds^2 = -dt^2 + a^2(t) \left[\frac{dr^2}{1 - kr^2} + r^2 d\Omega^2 \right], \quad (1.8)$$

where $k = 0$ for a flat spatial geometry, $k = \pm 1$ allows for spatial curvature, and the time evolution is specified by the scale factor, $a(t)$. The Hubble parameter, $H = \dot{a}/a$, gives the rate of expansion of the Universe.

A period of inflation will drive the spatial curvature close to zero, in good agreement with current observations. As a consequence, we will assume spatial flatness and set $k = 0$ for most considerations; however, we discuss constraints on curvature obtainable with CMB-S4 in Sect. 1.2.2.2.

In addition to a flat universe, the simplest models of inflation predict that the primordial density perturbations are adiabatic, very nearly Gaussian, and nearly scale-invariant, in agreement with existing observations.




Figure 6. Theoretical predictions for the temperature (black), E-mode (red), and tensor B-mode (blue) power spectra. Primordial B-mode spectra are shown for two representative values of the tensor-to-scalar ratio: $r = 0.001$ and $r = 0.05$. The contribution to tensor B modes from scattering during the recombination epoch peaks at $\ell \approx 80$ and from reionization at $\ell < 10$. Also shown are expected values for the contribution to B modes from gravitationally lensed E modes (green). Current measurements of the B-mode spectrum are shown for BICEP2/Keck Array (light orange), POLARBEAR (orange), and SPTPol (dark orange). The lensing contribution to the B-mode spectrum can be partially removed by measuring the E-mode polarization and exploiting the non-Gaussian statistics of the lensing.

In single-field slow-roll inflation, the gauge-invariant combination of metric and scalar-field fluctuations that is conserved outside the horizon has the power spectrum

$$\Delta_{\mathcal{R}}^2(k) = \frac{1}{2\epsilon M_{\text{P}}^2} \left(\frac{H}{2\pi} \right)^2 \Big|_{k=aH}, \quad (1.9)$$

where $\epsilon = -\dot{H}/H^2$ is the first slow-roll parameter, and $M_{\text{P}} = 1/\sqrt{8\pi G}$ is the reduced Planck mass.

In addition to producing primordial density perturbations, the rapid expansion of spacetime creates primordial gravitational waves that imprint a characteristic polarization pattern onto the CMB. Their power spectrum is given by

$$\Delta_{\gamma}^2(k) = \frac{8}{M_{\text{P}}^2} \left(\frac{H}{2\pi} \right)^2 \Big|_{k=aH}, \quad (1.10)$$

and is a measure of the expansion rate of the Universe during inflation. Together with the Friedmann equation a detection of primordial gravitational waves would then allow us to infer the inflationary energy scale.¹

¹In some models of inflation the relation between the tensor-to-scalar ratio and the energy scale of inflation are broken because there are additional sources of gravitational wave production [19]. However, the signal in these models is highly non-Gaussian and would not be mistaken for quantum fluctuations.




Figure 7. Forecast of CMB-S4 constraints in the n_s - r plane for a fiducial model with $r = 0.003$. For comparison, we also show the current best constraints from a combination of the BICEP2/Keck Array experiments, Planck [5], and BAO data. These are compared to several theoretical models. Chaotic inflation with $V(\phi) = \mu^{4-p}\phi^p$ for $p = 2/3, 1, 2$ are shown as blue lines for $47 < N_* < 57$ (with smaller N_* predicting lower values of n_s). The Starobinsky model and Higgs inflation are shown as small and large orange filled circles, respectively. The green band shows the predictions for a quartic hilltop model for which the potential throughout the inflationary period is described by $V(\phi) \approx V_0(1 - (\phi/M)^4)$ before developing a minimum. The gray band shows the prediction of a sub-class of α -attractor models [6], and the purple line shows fibre inflation [9].

CMB-S4 will be able to detect primordial gravitational waves for $r > 0.003$, and a detection would point to inflationary physics near the energy scale associated with grand unified theories. As a consequence, a detection would provide additional evidence in favor of the idea of the unification of forces, and would probe energy scales far beyond the reach of the Large Hadron Collider (LHC) or any conceivable collider experiment. Selected models within reach of CMB-S4 are shown in Fig. 7 together with current limits and constraints expected for CMB-S4.

The knowledge of the energy scale of inflation would have broad implications for many other aspects of fundamental physics, including ubiquitous ingredients of string theory like axions and moduli, fields that control the shapes and sizes of the compact dimensions.

Because the polarization pattern is due to quantum fluctuations in the metric of spacetime generated during inflation, a detection would also provide evidence that gravity is quantized. Moreover, because many of the models that predict a signal that is strong enough to be detected with CMB-S4 are based on symmetry principles and inflation occurs at high energy and large inflaton field range, a detection would also provide some insights into the nature of quantum gravity.

In the absence of a detection, the upper limit of $r < 0.001$ at 95% CL achievable by CMB-S4 is nearly two orders of magnitude stronger than current limits and about a factor of 5 stronger than projected limits for Stage-3 experiments. As explained in more detail in Ref. [3], such an upper limit would exclude large classes of inflationary models. In summary, there are only two classes of model that naturally explain the observed value of n_s and are consistent with current data. The first class consists of monomial models with potentials that during inflation are well approximated by $V(\phi) = \mu^{4-p}\phi^p$. Representatives in this class are shown in




Figure 8. Forecast of CMB-S4 constraints in the n_s - r plane for a fiducial model with $r = 0$. Also shown are the current best constraints from a combination of the BICEP2/Keck Array experiments and Planck [5]. The Starobinsky model and Higgs inflation are shown as small and large orange filled circles. The lines show the classes of model that naturally explain the observed value of n_s . The corresponding potentials all either polynomially or exponentially approach a plateau. The scale in field space over which the potential approaches the plateau is referred to as the “characteristic scale” (see Ref. [3] for more details). We show different values, $M = M_{\text{P}}/2$, $M = M_{\text{P}}$, $M = 2M_{\text{P}}$, and $M = 5M_{\text{P}}$. Longer dashes correspond to larger values of the scale M . The Planck scale plays an important role because the gravitational scale and the characteristic scale share a common origin. The number of e-folds N_* chosen for the figure corresponds to nearly instantaneous reheating, which leads to the smallest values for r for a given model. Other reheating scenarios predict larger values of r and are easier to detect or exclude.

Fig. 7, and the entire class of models is shown in Fig. 8. The second class consists of models in which the potential $V(\phi)$ approaches a plateau, either polynomially or exponentially. The potential for models in this class has a characteristic scale over which the potential varies [3].² The sensitivity of CMB-S4 is chosen to exclude all models in this class with a characteristic scale that exceeds the Planck scale. The Planck scale constitutes an important threshold because the scale of gravitational interactions and the characteristic scale may share a common origin and be linked to each other, such as in the Starobinsky model [7], in Higgs inflation [8], or more general models involving non-minimally coupled scalar fields. As a consequence, even in the absence of a detection CMB-S4 would significantly advance our understanding of inflation, and would dramatically affect how we think about the theory. The classes of model that naturally predict the observed value of n_s , together with current constraints and constraints expected for CMB-S4, are shown in Fig. 8.

1.2.2 Primordial density perturbations

CMB-S4 can also seek to characterize the primordial Universe by searching for well-motivated signatures in the scalar fluctuations, in the primordial power spectrum, and non-Gaussianities.

²This characteristic scale was introduced in Ref. [3] and should not be confused with the field range or the energy scale of inflation. For a discussion see Refs. [3] and [20].

We will begin with discussions of the overall shape of the primordial power spectrum, and generalize to investigate the possibility of measuring features in the primordial power spectrum.

Currently, constraints on non-Gaussianity rely on measurements of the CMB bispectrum, and CMB-S4 will improve on these measurements. Below, we will summarize some well-motivated bispectrum shapes [21, 22]. In addition, the precise measurement of B -mode polarization opens up another possible window to explore the physics of the early Universe. Non-Gaussianities that measure interactions between tensors and scalars, can, for the first time, be constrained with sensitivity that has the potential to exceed that of scalar non-Gaussianities [23]. Although the presence of these couplings is predicted to be small [24, 25] in the simplest inflation models, such a detection would present evidence of exciting new physics [26, 27]. In addition, theoretical consistency conditions between scalar and tensor non-Gaussianities provide a compelling observational target [28, 29, 30].

In addition to the bispectrum, fluctuations in the galaxy density are also a powerful probe of primordial non-Gaussianity, since the galaxy power spectrum will be biased on large scales in the presence of primordial local non-Gaussianity. Recently, two independent techniques have been proposed to use CMB data to precisely measure the scale-dependence of this galaxy bias on large scales. Precise measurements require “cosmic variance cancellation,” which uses multiple tracers of the same field, to remove cosmic variance on this bias. We will present results of forecasts on the local type of non-Gaussianity using CMB-S4 data combined with data from the LSST.

1.2.2.1 Primordial Power Spectrum

Currently, our best constraints on the shape of the primordial power spectrum on large scales come from measurements of the CMB temperature anisotropies by the *Planck* satellite. The current value of the slope of the primordial curvature power spectrum is $n_s = 0.965 \pm 0.004$ [31]. The forecasted uncertainty on the slope for a model of Λ CDM + running of the tilt using CMB-S4 noise, including the CMB temperature and polarization fields, is $\sigma_{n_s} = 0.002$, reducing the error from *Planck* by a factor of 2. Slow-roll models of inflation predict a running of the tilt of order $(n_s - 1)^2 \sim 10^{-3}$. With CMB-S4, the forecasted uncertainty on the running is $\sigma_{n_{\text{run}}} = 0.0029$ (see Eq. 1.5), which combined with large-scale structure constraints could be tightened even further. For comparison, we find $\sigma_{n_s} = 0.0027$ and $\sigma_{n_{\text{run}}} = 0.0038$ for Simons Observatory.

Features in the CMB temperature power spectrum [32, 33, 34, 31] have been claimed and interpreted as possible evidence for new physics during inflation. Evidence of a feature would provide us with yet another unique signature that can be used to understand the physics of the early Universe [35]. Specific forms of such features are typically handled on a case-by-case basis. However, for model-independent analyses, it is desirable to have a simple and accurate prescription that relates CMB observables to the shape of the inflationary potential. Various approaches exist in the literature for reconstructing this primordial power spectrum (see e.g., [36, 37, 38, 39, 40, 41, 42, 43])

In a series of papers [44, 45, 46, 47, 48], a formalism (known as “generalized slow roll”) has been developed to test the hypotheses of slow-roll and single-field inflation in a general and model-independent way. In this formalism, there is a single source function G' that is responsible for the observable features and it is simply related to the local slope and curvature of the inflationary potential in the same way the tilt is for the case of standard slow roll:

$$G' \approx 3 \left(\frac{V'}{V} \right)^2 - 2 \frac{V''}{V}, \quad (1.11)$$

where primes denote derivatives with respect to the inflation field. This framework can be used to map constraints from the CMB onto constraints on the shape of the inflationary potential, without assuming any specific model of inflation.




Figure 9. Reconstructed primordial power spectrum with CMB-S4 noise curve including atmospheric noise co-added with Planck noise, noise from point sources, and lensing reconstruction. We parameterize our source function $\delta G'$ in the range $0.5 < s/\text{Mpc} < 4870$ with a total of 40 parameters, equispaced in logarithmic scale. Also shown are the constraints coming from the Simons Observatory baseline, for the same fraction of the sky observed by CMB-S4 (as discussed in Ref. [49]).

Here we adopt the generalized slow-roll approximation to accommodate order unity variations in the power spectrum from slow-roll predictions. Under this formalism the primordial power spectrum $\Delta_{\mathcal{R}}^2(k)$ can be written in terms of a single unknown source function G' , and the rest are known functions. We perform a Fisher matrix analysis varying the six standard Λ CDM cosmological parameters (A_s , n_s , τ , θ , Ω_c , Ω_b) and 40 coefficients that parametrize the source function.

The ordinary slow-roll approximation corresponds to $G'(\ln s) = 1 - n_s$, and results in a power-law curvature power spectrum. We restrict our parameterization of fluctuations, $\delta G'$ (around the constant $G' = 1 - n_s$) to $0.5 < s/\text{Mpc} < 4870$. We sample $\delta G'$ for a total of 40 parameters, equispaced on a logarithmic scale. We then construct a smooth function using the natural spline of these points. In Figure 9 we show the forecasted reconstructed initial curvature power spectrum for CMB-S4.

1.2.2.2 Spatial Curvature

The theory of inflation predicts an almost flat universe with curvature $|\Omega_K| \sim 10^{-4}$ (from contributions from large-scale modes). If departures from this value were measured at a significant level, this would give

us information about the physics of inflation. For instance, it could indicate that the inflaton field was not slowly rolling when the largest scales exited the horizon.

Combining CMB-S4, *Planck* TT on the part of the sky that is not observed by CMB-S4, and *Planck*'s low- ℓ ($\ell < 30$) measurement of TT over the full sky (along with a *Planck* prior on τ of $\sigma(\tau) = 0.007$), the constraints on the curvature reach $\sigma(\Omega_K) = 2.4 \times 10^{-3}$, which is a factor of 2.5 better than current constraints from the CMB, namely $\sigma(\Omega_K) = 6.5 \times 10^{-3}$ using CMB temperature, polarization, and lensing data from the *Planck* satellite. This constraint could be sharpened even further by combining CMB and baryon acoustic oscillation (BAO) data.

1.2.2.3 Isocurvature

Measurements of CMB temperature and polarization power spectra indicate that the primordial initial conditions are primarily adiabatic; that is, if we define

$$S_{ij} \equiv \frac{\delta n_i}{n_i} - \frac{\delta n_j}{n_j}, \quad (1.12)$$

then the adiabatic condition is that $S_{ij} = 0$. Here the species labels i, j can denote baryons, cold dark matter (CDM), photons (γ), or neutrinos. Number densities are denoted by n_i . Adiabatic perturbations are produced if the initial perturbations in all species are seeded by the inflaton. Otherwise, the initial conditions are a mixture of adiabatic and isocurvature perturbations, for which $S_{i\gamma} \neq 0$. These initial conditions determine the acoustic peak structure and large-scale amplitude of CMB anisotropies [50, 51, 52, 53, 54, 55], which can thus probe additional fields present during inflation. Present data still allow subdominant CDM density isocurvature (CDI), baryon density isocurvature (BDI), or neutrino density isocurvature (NDI) initial conditions (described in Refs. [56, 57, 58, 59, 34, 31]) at the roughly 1% level, depending on precise assumptions made about the spectral shape and multi-mode admixtures of isocurvature. CMB-S4 could shed light on inflationary physics by detecting isocurvature fluctuations. As specific examples of the physics that can be measured with CMB-S4, we consider two well-studied scenarios: the curvaton; and compensated isocurvature perturbations (CIPs). Axion-type isocurvature is discussed in Section 1.3.4.5.

The curvaton scenario is a model in which a sub-dominant field σ acquires vacuum fluctuations during inflation, sources ζ (the total gauge-invariant curvature perturbation), and then decays [60, 61, 62, 63, 64]. Curvaton candidates include sneutrinos, among others [65, 66, 67, 68, 69, 70, 71, 72]. Depending on whether a species i (or its set of quantum numbers) is produced by, before, or after curvaton decay, perturbations in i are offset from ζ , with distinct density perturbation ζ_i in the species i on surfaces of constant curvature. Here $i \in \{\nu, b, c\}$. The resulting isocurvature perturbations [63, 64, 73] have entropy fluctuation given by

$$S_{i\gamma} = \begin{cases} -3\zeta - 3(\zeta_\gamma - \zeta), & \text{if } i \text{ is produced before } \sigma \text{ decay,} \\ 3\left(\frac{1}{r_D} - 1\right)\zeta - 3(\zeta_\gamma - \zeta), & \text{if } i \text{ is produced by } \sigma \text{ decay,} \\ -3(\zeta_\gamma - \zeta), & \text{if } i \text{ is produced after } \sigma \text{ decay.} \end{cases} \quad (1.13)$$

The parameter r_D is the fractional energy density in the curvaton when it decays.

The mixture of isocurvature modes is determined by whether or not baryon number, lepton number, and CDM are produced before, by, or after curvaton decay. Curvaton-type isocurvature is distinct from axion isocurvature, because it is correlated (or anti-correlated) with ζ . If lepton number is produced by curvaton decay, the lepton chemical potential ξ_{lep} sets the size of NDI modes [64, 74, 75]:

$$S_{\nu\gamma} = -\frac{135}{7} \left(\frac{\xi_{\text{lep}}}{\pi}\right)^2 \zeta_\gamma. \quad (1.14)$$

There are 27 distinct curvaton scenarios, since baryon number, lepton number, and CDM could each be produced before, by, or after curvaton decay. Viable models are those in which one of baryon number or CDM is produced by curvaton decay, and those in which *both* baryon number and CDM are produced after curvaton decay. For curvaton-decay scenarios, we use the notation $(b_{y_b}, c_{y_c}, L_{y_L})$, where the subscripts run over $y_i \in \{\text{before, by, after}\}$ and L denotes lepton number. For example, $(b_{\text{before}}, c_{\text{by}}, L_{\text{by}})$ is a model in which baryon number is produced before curvaton decay, CDM by curvaton decay, and lepton number by curvaton decay. Current isocurvature limits favor values of $r_D \approx 1$, except for models in which baryon number is produced by curvaton decay and CDM before (or vice versa), which favor central values of $r_D \approx 0.16$ ($r_D \approx 0.84$).

The current limits [76] on r_D are shown in Table 1-1, along with a forecast of CMB-S4's sensitivity to r_D via isocurvature modes. There is dramatic improvement in constraints on the $(b_{\text{by}}, c_{\text{before}}, L_{\text{by}})$ and $(b_{\text{before}}, c_{\text{by}}, L_{\text{by}})$ scenarios because of the accompanying NDI perturbations. One unusual case is $(b_{\text{after}}, c_{\text{after}}, L_{y_L})$, where isocurvature just constrains the combination [76] $\chi_D \equiv [1 + \xi_{\text{lep}}^2 / (\pi^2) (1/r_D - 1)]^{-1}$, while the independent constraint on ξ_{lep}^2 is driven by the CMB limit on total relativistic energy density (N_{eff}).

Isocurvature scenario	<i>Planck</i>	CMB-S4
	$\Delta r_D / r_D^{\text{adi}}$	$\Delta r_D / r_D^{\text{adi}}$
$(b_{\text{by}}, c_{\text{before}}, L_{y_L})$	0.03	0.005
$(b_{\text{before}}, c_{\text{by}}, L_{y_L})$	0.01	0.004
$(b_{\text{by}}, c_{\text{after}}, L_{y_L})$	0.04	0.01
$(b_{\text{after}}, c_{\text{by}}, L_{y_L})$	0.008	0.002
$(b_{\text{by}}, c_{\text{by}}, L_{y_L})$	0.007	0.002
	$\Delta \chi_D / \chi^{\text{adi}}$	$\Delta \chi_D / \chi^{\text{adi}}$
$(b_{\text{after}}, c_{\text{after}}, L_{y_L})$	0.003	0.0004
	$\Delta \xi_{\text{lep}}^2$	$\Delta \xi_{\text{lep}}^2$
$(b_{\text{by}}, c_{\text{before}}, L_{\text{by}})$	0.02	0.002
$(b_{\text{before}}, c_{\text{by}}, L_{\text{by}})$	0.4	0.04
$(b_{\text{by}}, c_{\text{after}}, L_{\text{by}})$	0.3	0.04
$(b_{\text{after}}, c_{\text{by}}, L_{\text{by}})$	0.3	0.04
$(b_{\text{by}}, c_{\text{by}}, L_{\text{by}})$	0.3	0.04
$(b_{\text{after}}, c_{\text{after}}, L_{\text{by}})$	0.3	0.04

Table 1-1. Isocurvature constraints on r_D and ξ_{lep}^2 , both at 95% CL using *Planck* TT+BAO+LowP data [76] in viable curvaton decay-scenarios, and Fisher forecasts for CMB-S4 sensitivity.

If baryon number is produced before curvaton decay, while CDM is produced by curvaton decay, CMB 3-pt correlations (see Sec. 1.2.2.4) are excited, with $f_{\text{NL}}^{\text{local}} \approx 5$, also within the reach of CMB-S4. As a result, CMB-S4 provides a useful set of consistency checks if certain curvaton signatures are observed [76].

If baryon number is produced by curvaton decay, but CDM is produced before (or vice versa), a relatively large compensated isocurvature perturbation S_{bc} (CIP) is produced between the baryons and CDM. Curvaton-generated CIPs are proportional to ζ , $S_{\text{bc}} = A\zeta$, where $A \approx 17$ in the $(b_{\text{by}}, c_{\text{before}}, L_{y_L})$ scenario and $A \approx -3$ for $(b_{\text{by}}, c_{\text{before}}, L_{y_L})$. CIPs would induce non-Gaussianities in the CMB [77, 78, 79, 80]. At CMB-S4 sensitivity [80], the threshold for a 95% CL detection is $A \approx 10$, and so a CIP test of the $(b_{\text{by}}, c_{\text{before}}, L_{y_L})$

scenario is within reach of CMB-S4, a substantial improvement over *Planck* sensitivity. For uncorrelated CIPs, we find that the sensitivity of CMB-S4 to a scale-invariant (SI) angular power spectrum of uncorrelated CIPs is $\Delta_{\text{cl}} = 0.003$ at the 95% CL with current parameters [80], assuming *Planck* cosmology [34]. Here Δ_{cl} is the rms CIP amplitude on cluster scales. This is a significant improvement over the upper limit of $\Delta_{\text{cl}} \leq 0.077$ from the *Wilkinson Microwave Anisotropy Probe* (WMAP) [79], or the *Planck* upper limit $\Delta_{\text{cl}} \leq 0.064$ obtained in Ref. [81].

1.2.2.4 Primordial non-Gaussianity: the bispectrum

Unlike the scalar and tensor power spectra, higher-order correlations of the scalar modes are directly sensitive to the dynamics and field content responsible for inflation (and its alternatives) [82]. While non-Gaussian correlations are small in conventional single-field slow-roll inflation, there exist many other possibilities for the nature of inflation that give strikingly different predictions when we move beyond the power spectrum. The constraints on non-Gaussianity from the WMAP [83, 84] and *Planck* [16, 17] satellites currently place the most stringent limits on a wide range of mechanisms for inflation; however, these measurements are not sufficiently sensitive to suggest that a particular mechanism is favored by the data.

The space of non-Gaussian signals from inflation can broadly be grouped into two conceptual categories that generate distinguishable features in the correlation functions. These are signals that: (1) indicate non-trivial self-interactions of the effective inflaton fluctuation; or (2) indicate interactions with degrees of freedom other than the inflaton. For the first category, self-interactions that respect the time-translation invariance during inflation lead to qualitatively different predictions from self-interactions that violate it. For the second category the signatures qualitatively depend on the mass of the additional degrees of freedom (fields). Additional light degrees of freedom can fluctuate significantly and may have large self-interactions. Heavy degrees of freedom do not fluctuate appreciably, but may come into existence by quantum fluctuations and decay into inflaton quanta, generating non-Gaussian correlations [85, 86, 29, 87]. Alternatively, they may become excited by the dynamics of the inflaton and their backreaction on the inflationary dynamics may lead to non-Gaussian correlations [88, 89, 90, 91].

Constraints on non-Gaussianity are often expressed in terms of the correlator of three scalar modes, described by the bispectrum $B_{\mathcal{R}}(\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3)$, defined through

$$\langle \mathcal{R}(\mathbf{k}_1)\mathcal{R}(\mathbf{k}_2)\mathcal{R}(\mathbf{k}_3) \rangle = (2\pi)^3 \delta(\mathbf{k}_1 + \mathbf{k}_2 + \mathbf{k}_3) B_{\mathcal{R}}(\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3). \quad (1.15)$$

The structure of particle interactions relevant for inflation provides both a general organizing principle for this functional space and several specific well-motivated forms of the bispectrum. We now briefly review the classification of scalar non-Gaussianity from inflation and comment on non-Gaussian signatures that are especially relevant for large-field inflation.

Our convention for extracting a normalization of the amplitude, f_{NL} , follows Refs. [21] and [16], i.e.,

$$B_{\mathcal{R}}(k_1, k_2, k_3) = \frac{3}{5}(4\pi^4)2A_s^2 f_{\text{NL}} F(k_1, k_2, k_3), \quad (1.16)$$

where F is referred to as the shape of the bispectrum, which in a scale invariant universe $\propto k^{-6}$. Different physical effects during inflation lead to different shapes for the bispectrum.

All bispectra that come from fluctuations of the field that drives inflation (“single-clock” scenarios) most strongly couple Fourier modes k of similar amplitude. The “squeezed limit” of these bispectra (the coupling of modes $k_1 \ll k_2 \sim k_3$) is very restricted. A large fraction of the parameter space for scenarios involving interactions during inflation that respect the underlying shift symmetry (i.e., are approximately scale-invariant) are well captured by the equilateral [21] and orthogonal shapes [83] shapes. This includes scenarios

in which inflaton fluctuations have non-trivial self-interactions [92, 93, 94, 95, 96, 83] or couplings between the inflaton and other (potentially massive) degrees of freedom [97, 98, 99, 100, 101, 89, 85]. Single-field slow-roll inflation necessarily produces $f_{\text{NL}}^{\text{equil}} < 1$ [102] and therefore any detection of $f_{\text{NL}}^{\text{equil}} \geq 1$ would rule out this large class of models. A detection would imply that inflation is a strongly coupled phenomenon and/or involved more than one field [103, 104, 105]. In single-field inflation, the amplitude of the non-Gaussianity typically suggests a new energy scale, M , such that $f_{\text{NL}}^{\text{equil}} \propto A_s^{-1/2} (H/M)^2$ [96, 106]; at this energy scale self-interactions become strongly coupled and current limits on non-Gaussianity [17] translate into $M > \mathcal{O}(10)H$. In the presence of additional hidden sectors, the amplitude of non-Gaussianity scales with the strength of the coupling between the inflaton and these additional fields, usually suppressed by an energy scale Λ . Current limits give $\Lambda > (10-10^5)H$ [107, 108], with the variation depending mostly on the dimension of the operator coupling the two sectors. For $r > 0.01$, these constraints require some of the interactions to be weaker than gravitational. CMB-S4 would further tighten existing constraints on a wide variety of interactions of the inflaton with itself and any other fields that are excited during inflation.

When light degrees of freedom other than the inflaton contribute to the observed scalar fluctuations, coupling between modes of very different wavelengths is allowed. Historically, the most well-studied signature of this type comes from the “local” shape, which couples short wavelength modes $k_2 \sim k_3$ to a long wavelength mode $k_1 (\ll k_2, k_3)$. A significant detection of this shape would rule out all models of single-clock inflation [109]. In addition, such a signal would open the door to significant cosmic variance on all scales, from coupling of fluctuations within our observed volume to any super-Hubble modes [110, 111, 112]. Indeed, there would be room for a significant shift between the observed amplitude of scalar fluctuations (and so the observed r) and the mean value of fluctuations on much larger scales [113]. Any scenario that predicts local non-Gaussianity, together with fluctuations on scales much larger than our observed volume, predicts a probability distribution for our observed $f_{\text{NL}}^{\text{local}}$, but many well-motivated scenarios also predict a small mean value; these include the simplest modulated reheating scenario [114] and ekpyrotic cosmology [115], both of which predict mean values of $f_{\text{NL}}^{\text{local}} \approx 5$.

Currently the strongest constraints on the local shape come from the *Planck* 2015 temperature and polarization analysis that finds $f_{\text{NL}}^{\text{local}} = 0.8 \pm 5.0$ [17]. CMB-S4 will improve these limits by a factor 2 for local and orthogonal shapes.³ Note that these forecasts assume that CMB-S4 can be combined with large-scale measurements from *Planck*, the impact of secondaries [116] can be reduced, and temperature and polarization can be sufficiently delensed to remove large contributions to the variance [117]. Even under these assumptions, the constraints are insufficient to reach the most compelling theoretical threshold around $|f_{\text{NL}}^{\text{local}}| \lesssim 1$ [104]. Nonetheless, the landscape of inflationary models is large and a detection would instantly present a monumental discovery similar to that of gravitational waves. CMB-S4 could, for example, provide hints for the mean level of non-Gaussianity expected from modulated reheating scenario or ekpyrotic cosmology at roughly the 2σ level. The simplest curvaton scenario, which predicts $f_{\text{NL}} = -5/4$ [63], will be out of reach using just the bispectrum measured by CMB-S4, but we will show in the next subsection that combining CMB-S4 with LSST data is likely to reach this significant threshold.

Table 1-2 shows the forecasted constraints on the local, equilateral, and orthogonal shapes from CMB-S4 using bispectrum measurements. In fact we expect that CMB-S4 would likely provide the strongest constraints on equilateral and orthogonal templates that will be available for the foreseeable future. Since limits on non-Gaussianity provide a unique and fundamental insight into the nature of inflation, this increased sensitivity would provide a non-trivial improvement in our understanding of the Universe.

³Note that constraints on orthogonal and local non-Gaussianities in the CMB scale almost as expected, i.e., as the square root of the number of observed modes. Unfortunately, some information is lost in projection, which particularly affects equilateral non-Gaussianity, scaling more like the fourth root of the number of modes.

Perhaps of special interest for CMB-S4 are non-Gaussian signatures that would be expected in models of large-field inflation, which could lead to non-trivial features in the bispectrum. For example, in models in which the inflaton is an axion, there is an approximate discrete shift symmetry [118, 119, 120] that naturally lead to periodic features in the bispectrum. Often these models predict counterparts in the power spectrum where the associated features are expected to be detected first [121], but this need not be the case [122]. An attempt has been made [17] to look for resonant and local features in the bispectrum; however, at present, these models have not yet been constrained systematically.

Correlators including at least one B mode will benefit significantly from the improved sensitivity of CMB-S4. In particular, the three point correlation $\langle \mathcal{R}(\mathbf{k}_1)\mathcal{R}(\mathbf{k}_2)\gamma_\lambda(\mathbf{k}_3) \rangle$ can be constrained using $\langle BTT \rangle$, $\langle BTE \rangle$ and $\langle BEE \rangle$.

The details of the tensor-scalar-scalar correlator are contained in the bispectrum $B_{\mathcal{R}\mathcal{R}\gamma_\lambda}(\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3)$, defined by pulling out the appropriate polarization structure associated with the tensor mode:

$$\langle \mathcal{R}(\mathbf{k}_1)\mathcal{R}(\mathbf{k}_2)\gamma_\lambda(\mathbf{k}_3) \rangle = (2\pi)^3 \delta(\mathbf{k}_1 + \mathbf{k}_2 + \mathbf{k}_3) B_{\mathcal{R}\mathcal{R}\gamma_\lambda}(k_1, k_2, k_3) e_{ij}(\mathbf{k}_3, \lambda) \hat{k}_1^i \hat{k}_2^j, \quad (1.17)$$

where $e_{ij}(\mathbf{k}, \lambda)$ is the transverse-traceless polarization tensor (see Eq. 1.2). The amplitude and momentum dependence of the bispectrum $B_{\mathcal{R}\mathcal{R}\gamma_\lambda}$ can be parametrized by [23]

$$B_{\mathcal{R}\mathcal{R}\gamma_\lambda}(k_1, k_2, k_3) = 16\pi^4 A_s^2 \sqrt{r} f_{\text{NL}}^{\mathcal{R}\mathcal{R}\gamma} F(k_1, k_2, k_3). \quad (1.18)$$

In the bottom half of Table 1-2 we show the results of forecasts for $\sqrt{r} \tilde{f}_{\text{NL}}$ using local-, equilateral-, and orthogonal-like templates. A detection of this correlation would be an immediate indication of some deviation from the simple inflationary paradigm [25, 123]. There are known possibilities that would generate a scalar-scalar-tensor bispectrum with larger amplitude and/or different shape: different symmetry patterns (e.g., solid inflation [124] or gauge-flation [125, 126]); gravitational waves not produced as vacuum fluctuations; or multiple tensors (e.g., bigravity [25]). Any non-zero tensor amplitude could also be sourced by a higher-order massive spin field that couples to two scalars and one graviton (see for example Ref. [28] for a discussion of such signatures).

1.2.2.5 Primordial non-Gaussianity: large-scale structure

CMB-S4 can also contribute to the measurement of local primordial non-Gaussianities in an indirect way, by complementing a galaxy survey to exploit so-called sample-variance cancellation. By combing CMB-S4 data with measurements of large-scale structure, for example with LSST [129], impressive constraints could be achieved on $f_{\text{NL}}^{\text{local}}$. We briefly review this method and forecast constraints on the local parameter.

On linear scales, galaxy bias relates the amplitude of matter perturbations δ_m and galaxy perturbations δ_g . Local non-Gaussianities lead to a scale-dependent correction to the galaxy bias b_g on large scales, proportional to $f_{\text{NL}}^{\text{local}}/k^2$ [130]. If one can obtain a measurement of both the biased galaxy field and the unbiased matter field, one can cancel out the stochastic nature of the modes to measure $b_g = \delta_g/\delta_m$ without sample variance [131]. Of course the measurement of both modes will in practice be noisy and limit the effect of sample-variance cancellation.

Two different methods have recently been proposed for how to use high-resolution CMB maps to measure the matter field δ_m in conjunction with a large-scale galaxy survey for sample-variance cancellation: the reconstructed CMB lensing potential [132]; and the kSZ reconstruction [133]. Interestingly these two techniques trace different modes, since lensing probes transverse modes and kSZ velocities probe radial modes of the underlying matter distribution. We now discuss each technique in more detail.

Shape: $\langle\mathcal{R}\mathcal{R}\mathcal{R}\rangle$ $\langle TTT\rangle, \langle TTE\rangle, \langle TEE\rangle, \langle EEE\rangle$	Current	CMB-S4 goal	Conservative	CV-limited
f_{sky}	75%	43%	43%	100%
$\sigma(f_{\text{NL}}^{\text{local}})$	5	1.9 (5.3)	2.1	0.8 (7.1)
$\sigma(f_{\text{NL}}^{\text{equil}})$	43	22.1 (-0.4)	23.5	11.8 (-1.9)
$\sigma(f_{\text{NL}}^{\text{ortho}})$	21	9.0 (-5.0)	10.6	4.4 (-6.3)
Shape: $\langle\mathcal{R}\mathcal{R}\gamma\rangle$ $\langle BTT\rangle, \langle BTE\rangle, \langle BEE\rangle$	Current	CMB-S4 goal	Conservative	CV-limited
f_{sky}	69%	3%	3%	100%
$\sigma(\sqrt{r}\tilde{f}_{\text{NL}}^{\text{local}})$	28	0.79	1.2	0.052
$\sigma(\sqrt{r}\tilde{f}_{\text{NL}}^{\text{equil}})$...	16	24	1.7
$\sigma(\sqrt{r}\tilde{f}_{\text{NL}}^{\text{ortho}})$...	4.4	7.4	0.41

Table 1-2. Forecast constraints on non-Gaussianity. $\langle\mathcal{R}\mathcal{R}\mathcal{R}\rangle$: Current constraints are from Planck . Forecasted constraints for CMB-S4 goal use Planck noise at low ℓ and bispectra are computed using the Planck 2018 cosmology [127]. Forecasts do not include foregrounds or lensing contributions to the covariance, but the signal bias from lensing is given in brackets ($T + E$, see Ref. [128]); some secondaries are expected (see Ref. [116]), in particular in the local and orthogonal shapes constrained using temperature data alone. We quote constraints without using $\langle TTT\rangle$ in the conservative column for reference, where $\sigma_{f_{\text{NL}}} = ((\sigma_{f_{\text{NL}}}^{\text{All}})^{-2} - (\sigma_{f_{\text{NL}}}^T)^{-2})^{-1/2}$. Some of these can be cleaned using e.g., ILC techniques, while others would require simulations. Deprojection of tSZ was performed in Ref. [49], which suggested minimal impact, so we do not expect the constraints to be impacted as much as removing $\langle TTT\rangle$ entirely, but further study is required, especially for contributions that cannot be removed using multi-frequency information such as kSZ. For polarization, apart from the reionization-lensing bispectrum, leakage of secondaries is currently less well studied. Constraints can be combined over 35% of the sky with Planck sensitivity, improving constraints by $\mathcal{O}(10\%)$. You can compare the results to the cosmic variance limited constraints with $\ell_{\text{max}} = 5000$ in the last column. $\langle\mathcal{R}\mathcal{R}\gamma\rangle$: Current constraint are derived from temperature only and shape considered is not exactly $f_{\text{NL}}^{\text{local}}$. There are currently no constraints on equilateral and orthogonal non-Gaussianities $\langle\gamma\mathcal{R}\mathcal{R}\rangle$. Forecasts add the small-aperture telescopes over 3% of the sky. We assume 13% residual lensed B-mode power. Preliminary studies show that foregrounds could leak into $\langle BTE\rangle$ and $\langle BEE\rangle$, and we show how projected constraints are affected when these combinations are not included in the conservative column. Lensing contributions are expected to be present in these observables as well, but neither they nor foregrounds are included in these forecasts.

Kinematic Sunyaev-Zeldovich effect The large-scale radial velocity field can be reconstructed by a process called kSZ tomography [134, 135]. This is possible because the kSZ temperature is proportional to the line-of-sight integral over the product of electron density ρ_e and radial velocity v_r . If one uses galaxies as tracers for the electron distribution, one can construct a quadratic estimator that combines small-scale CMB temperature fluctuations and the small-scale galaxy distribution, to estimate the large-scale velocity field. The noise in this kSZ velocity reconstruction is independent of the scale of the perturbation. In linear theory the velocity field is related to the matter density field by a factor faH/k , where f is the growth rate, resulting in a density reconstruction noise proportional to k^2 that can be much smaller than galaxy shot noise (which is independent of k) on large scales.

The kSZ-derived matter density can then be cross-correlated with a galaxy tracer of the same large-scale density field [133], as a function of scale, to determine the scale dependence of the clustering bias. Since the same modes are being measured in each survey, there is no sample variance in this comparison, greatly improving the precision on large scales.

We forecast $\sigma(f_{\text{NL}}^{\text{local}}) = 0.57$ for the combination of CMB-S4 kSZ and LSST, using the forecasting pipeline of Ref. [133], which can be compared to $\sigma(f_{\text{NL}}^{\text{local}}) = 1.45$ from LSST alone (assuming no internal sample-variance cancellation within the LSST data). For comparison, for Simons Observatory plus LSST the sensitivity is forecasted to be $\sigma(f_{\text{NL}}^{\text{local}}) = 1.0$ [136].

CMB Lensing An alternative method to measure $f_{\text{NL}}^{\text{local}}$ using sample-variance cancellation is to cross-correlate the CMB-S4 CMB lensing measurement with galaxy surveys such as LSST [137, 132]. Since the CMB lensing convergence field is determined by all the structure between the CMB last-scattering surface and the observer, we need to include galaxies at high redshift, $z > 3$, to obtain a large cross-correlation coefficient between CMB lensing and galaxies, which is required for sample-variance cancellation to work efficiently. In addition to the LSST galaxies at $z \leq 3$ specified in the last section, we therefore include additional LSST galaxies at $z = 3-7$, based on extrapolating [132] recent Hyper Suprime-Cam observations [138, 139] of dropout galaxies [140] in that redshift range (not including these high-redshift galaxies would degrade the forecasted $f_{\text{NL}}^{\text{local}}$ precision reported below by about a factor of 2). We split the LSST galaxies into six tomographic redshift bins at $z = 0-0.5, 0.5-1, 1-2, 2-3, 3-4$, and $4-7$, and estimate the ability of CMB-S4 to measure local non-Gaussianity.

The result of this forecast depends strongly on the minimum multipole ℓ_{min} of both the LSST galaxy density field and the CMB lensing convergence, since the signal scales as $1/k^2$. For $\ell_{\text{min}} = 2, 8, 15, 20$, and 40 we obtain $\sigma(f_{\text{NL}}^{\text{local}}) = 0.72, 0.99, 1.3, 1.5$, and 2.2 , respectively, assuming $f_{\text{sky}} = 0.3$ overlap with LSST. Similarly to the kSZ forecast above, this method therefore allows us to probe the $\sigma(f_{\text{NL}}^{\text{local}}) \approx 1$ regime. With the Simons Observatory, the same forecast yields $\sigma(f_{\text{NL}}^{\text{local}}) = 1.2, 1.3, 1.5, 1.7$, and 2.6 , which is 60% larger than CMB-S4 for $\ell_{\text{min}} = 2$ and about 15% larger for $\ell_{\text{min}} \geq 15$.

1.3 The dark Universe

The second science theme relates to the fundamental physics of invisible components of the Universe.

CMB-S4 will probe the fundamental physics of components that are difficult or impossible for us to observe directly. It will test for the presence of light, relativistic relic particles, beyond our Standard Model of particle physics, that were thermally produced in the early Universe. To date, CMB observations by the *Planck* satellite can probe light particles that departed from equilibrium (“froze out”) as early as the first $\simeq 50$ micro-seconds of the Universe. With CMB-S4 we can push back this frontier by over a factor of 10,000, to the first fractions of a nanosecond.

CMB-S4 will achieve sensitivity to relics that froze out well before the quark-hadron phase transition (the epoch when the Universe cooled sufficiently that quarks became locked into hadrons like neutrons and protons). The contribution of light relics to the energy density leads to observable consequences in the CMB temperature and polarization anisotropy. This is often parameterized with the “effective number of neutrino species,” N_{eff} . The collective influence of the three already-known light relics (the three families of neutrinos) has already been detected at high significance. Current data are only sensitive enough to detect additional relics that froze out after the quark-hadron transition, and Stage-3 CMB experiments can only push somewhat into that epoch, so CMB-S4’s ability to probe times well before the transition is a major advance.

CMB-S4 can measure the summed mass of the neutrino species. Together with terrestrial experiments this may shed light on the mechanism responsible for neutrino mass, one of the biggest mysteries of the Standard Model. The total mass unambiguously determines the absolute mass scale of neutrinos, and if the sum is $\lesssim 0.1\text{eV}$, CMB-S4 can furthermore disfavor the inverted mass hierarchy. Together with neutrino-less double beta-decay experiments, CMB-S4 may provide evidence to help determine if neutrinos are Dirac or Majorana particles, or point to unknown physics in the neutrino sector. The main neutrino observables are CMB lensing, lensing cross-correlation, and the abundances of galaxy clusters.

Current cosmological observations already require the existence of dark energy and dark matter, but these have not been observed in the laboratory. CMB-S4 can target the various predictions that dark-energy models make for the clustering and growth of matter fluctuations at late times, examining the secondary anisotropies from gravitational lensing and interactions with the gas in galaxies and galaxy clusters. Constraints on cosmic birefringence could probe the microphysics of the dark energy.

CMB-S4 can test dark-matter models and constrain parts of the parameter space that are inaccessible to laboratory experiments. In particular, CMB observations directly probe the physics of dark matter throughout cosmic history, and do not rely on assumptions about the local dark-matter phase-space distribution within the Milky Way. CMB-S4 can place constraints on a variety of scenarios, including dark matter that interacts with baryons or with dark radiation, or consists of ultra-light axion-like particles.

1.3.1 Light relics

A natural and important question is whether there are additional, unknown forms of radiation in the Universe (i.e., additional relativistic species). This radiation density leaves a measurable imprint on the CMB and can be determined with high precision by CMB-S4. Radiation also changes the Universe at late times, by altering the amplitude of clustering, the scale and phase of the baryon acoustic oscillations (BAOs), and abundances of light elements. Precision measurements of the radiation content of the Universe with the CMB

is both a window into the dark sector and a necessary tool for calibrating measurements of the lower-redshift Universe.

The possibility of additional radiation is also compelling, in the contexts of both particle physics and cosmology [11, 141, 12]. New light particles may arise in the form of axions and sterile neutrinos, or can appear as a byproduct of new symmetries that would explain the small mass of the Higgs boson [142, 143, 144, 145, 146, 147, 148, 10, 149, 150, 151, 152, 11, 153, 154, 141, 155, 12, 156, 157, 158, 159]. Furthermore, light particles can thermalize in the early Universe for wide ranges of unexplored parameter space, leading to an observable level of additional radiation. Light particles may form the dark matter (e.g., axions) or part of a dark sector, they can mediate forces in the dark and visible sectors, or they can result from the decay of new heavier particles. These possibilities may also play a role in explaining the discrepancies observed in the Hubble constant H_0 measurements [160, 161, 162, 163, 164], the amplitude of fluctuations σ_8 [165, 166, 167, 168], and clustering on small scales [169, 127].

CMB-S4 will provide a transformative measurement for the radiation density, particularly for particles that would have thermalized in the early Universe. Any such particle adds at least a percent-level contribution to the radiation energy density; the precise amount, per degree of freedom, is determined entirely by its decoupling temperature. CMB-S4 will be orders of magnitude more sensitive to decoupling temperatures than current experiments, and can reach targets that are not reachable by other means.

Cosmic neutrino background. The cosmic neutrino background is one of the remarkable predictions of the hot Big Bang scenario. In the very early Universe, neutrinos were kept in thermal equilibrium with the Standard Model plasma. As the Universe cooled, neutrinos decoupled from the plasma. A short time later, the relative number density and temperature in photons increased, primarily due to the transfer of entropy to photons from annihilating electron-positron pairs. The background of cosmic neutrinos persists today, with a temperature and number density similar to that of the CMB. Their energy density ρ_ν is most commonly expressed in terms of the effective number of neutrino species,

$$N_{\text{eff}} = \frac{8}{7} \left(\frac{11}{4} \right)^{4/3} \frac{\rho_\nu}{\rho_\gamma}, \quad (1.19)$$

where ρ_γ is the energy density in photons. This definition is chosen so that $N_{\text{eff}} = 3$ in the SM if neutrinos had decoupled instantaneously prior to electron-positron annihilation. The neutrino density ρ_ν receives a number of corrections from this simple picture of decoupling, and the best available calculations give $N_{\text{eff}}^{\text{SM}} = 3.045$ in the SM [170, 171, 172].

Cosmology is sensitive to the gravitational effects of neutrinos, both through their mean energy density [173, 174, 175, 176] and their fluctuations, which propagate at the speed of light in the early Universe due to the free-streaming nature of neutrinos [176, 177, 178]. A radiation fluid whose fluctuations do not exceed the sound speed of the plasma [179, 180] could arise from large neutrino self-interactions [181, 182, 183, 184, 185], neutrino-dark sector interactions [186, 187], or dark radiation self-coupling [188, 189, 190]. Such a radiation fluid can be observationally distinguished from free-streaming radiation, and can serve as both a foil for the cosmic neutrino background and a test of new physics in the neutrino and dark sectors [177, 191, 192].

Neutrinos are messengers from a few seconds after the Big Bang and provide a new window into our cosmological history. While these relics have been detected in cosmological data, higher precision measurements would advance the use of neutrinos as a cosmological probe. Furthermore, the robust measurement of the neutrino abundance from the CMB is crucial for inferring cosmic parameters, including the expansion history using BAOs [193], the neutrino masses [2], and H_0 [164].

Other light relics. *A measurement of the value of N_{eff} provides vastly more information than just the energy density in cosmic neutrinos.* The parameter N_{eff} is a probe of any particles that have the same gravitational influence as relativistic neutrinos, which is true of all kinds of (free-streaming) radiation. Furthermore, this radiation could have been created at much earlier times when the energy densities were even higher than in the cores of stars or supernovae, shedding light on the physics at new extremes of temperature and density, and on our early cosmic history.

New light particles that were thermally produced in the early Universe contribute to the neutrino density ρ_ν and increase N_{eff} above the amount from neutrinos alone. The presence of any additional species can therefore be characterized by $\Delta N_{\text{eff}} \equiv N_{\text{eff}} - N_{\text{eff}}^{\text{SM}}$. Since all such thermalized particles behave in the same way from a cosmological point of view, this parametrization captures a vast range of new physics, e.g., axions, sterile neutrinos, dark sectors, and beyond [150, 194, 155, 3].

Constraints on N_{eff} are broadly useful and, most importantly, allow the exploration of new and interesting territory in a variety of well-motivated models. This can be seen with a simple example: dark-matter-baryon scattering. For low-mass (sub-GeV) dark matter, current data allow for relatively large scattering cross-sections [195]. If they scatter through a Yukawa potential, which is a force mediated by a scalar particle, this force is consistent with fifth-force experiments and stellar cooling if the mediator has a mass around 200 keV. However, the particle that mediates the force necessarily⁴ contributes $\Delta N_{\text{eff}} \geq 0.09$ when it comes into thermal equilibrium with the Standard Model [196]. Excluding this value would require that the strength of the interactions is small enough to prevent the particle from reaching equilibrium at any point in the history of the Universe, which, consequently, limits the scattering cross-section, as shown in the left panel of Fig. 10. This measurement is sensitive to 10–15 orders of magnitude in cross-section that are not probed by direct constraints from cosmology and astrophysics, and is five orders of magnitude stronger than meson decay searches. We see that *cosmological measurements of ΔN_{eff} are an extremely sensitive probe of dark-sector physics that are complementary to more direct tests, both in the laboratory and with astrophysical observations* [196, 197].

More generally, the contribution to N_{eff} from any thermalized new particle is easy to predict because its energy density in equilibrium is fixed by the temperature and the number of internal states (e.g., spin configurations). Under mild assumptions (see e.g., Ref. [201] for a detailed discussion), the contribution to ΔN_{eff} is determined by two numbers, the last temperature at which it was in equilibrium, T_{F} , and the effective number of spin degrees of freedom, g_s , according to

$$\Delta N_{\text{eff}} = g_s \left(\frac{43/4}{g_\star(T_{\text{F}})} \right)^{4/3}. \quad (1.20)$$

The function $g_\star(T_{\text{F}})$ is the number of effective degrees of freedom (defined as the number of independent states with an additional factor of 7/8 for fermions) of the SM particle content at the temperature T_{F} . This function appears in the formula for ΔN_{eff} because it determines how much the photons are heated relative to a new light particle due to the annihilation of the heavy SM particles as the Universe cooled (see the right panel of Fig. 10). CMB-S4 is expected to reach a precision of $\sigma(N_{\text{eff}}) = 0.03$, which would extend our reach in T_{F} by several orders of magnitude for a particle with spin $s > 0$ and be the first measurement sensitive to a real scalar ($s = 0$) that decouples prior to the QCD phase transition.

To understand the impact of such a measurement, recall that equilibrium at temperature T arises when the production rate Γ is much larger than the expansion rate $H(T)$. At high temperatures, production is usually fixed by dimensional analysis, $\Gamma \propto \lambda^2 T^{2n+1}$, where λ is the coupling to the Standard Model with units of [Energy]^{- n} . The particle is therefore in equilibrium if $\lambda^2 \gg M_{\text{P}}^{-1} T^{-2n+1}$. There are two important features of this formula: (i) the appearance of the Planck scale M_{P} implies that we are sensitive to very weak

⁴The mediator with a mass of 200 keV is too heavy to contribute to N_{eff} , but it must decay to sub-eV mass particles, which will increase N_{eff} , in order to avoid more stringent constraints.




Figure 10. *Left:* Limits on the dark-matter-baryon cross-section σ_{bDM} for a Yukawa potential. Future cosmological constraints will restrict $\Delta N_{\text{eff}} < 0.09$ and, therefore, exclude cross-sections large enough to thermalize the (200 keV-mass) particle mediating the force [196]. This limit is compared to the direct bound on baryon-dark-matter scattering from the CMB [198] and to the constraints on dark forces from the Bullet Cluster [199]. The strongest current constraint is from the absence of meson decays to the mediator [200]. *Right:* Contributions of a single massless particle, which decoupled from the SM at temperature T_F , to the effective number of relativistic species, $N_{\text{eff}} = N_{\text{eff}}^{\text{SM}} + \Delta N_{\text{eff}}$, with the SM expectation $N_{\text{eff}}^{\text{SM}} = 3.045$ from neutrinos. The dashed line shows the 2σ limit from a combination of current CMB, BAO, and Big Bang nucleosynthesis (BBN) observations [127]. The purple line shows the projected sensitivity of CMB-S4 and illustrates its power to constrain light thermal relics. The displayed values on the right are the observational thresholds for particles with different spins and arbitrarily large decoupling temperatures.

couplings ($M_{\text{P}}^{-2} = 8\pi G_N$); and (ii) for $n \geq 1$ it scales like an inverse power of T . As a result, sensitivity to increasingly large T_F implies that we are probing increasingly weak couplings (lower production rates) in proportion to the improvement in T_F (not ΔN_{eff}). These two features explain why future measurements of ΔN_{eff} can be orders of magnitude more sensitive than terrestrial and astrophysical probes of the same physics [155, 3].

The impact of the current measurement of N_{eff} and the projected sensitivity of CMB-S4 is illustrated in Fig. 10. Anticipated improvement in measurements of N_{eff} translate into orders of magnitude in sensitivity to the temperature T_F . This temperature sets the reach in probing fundamental physics. Even in the absence of a detection, future cosmological probes would place constraints that can be orders of magnitude stronger than current probes of the same physics, including for axion-like particles [155] and dark sectors [202, 158, 196, 159]. It is also worth noting that these contributions to N_{eff} asymptote to specific values of $\Delta N_{\text{eff}} = 0.027$, 0.047, and 0.054 for a massless (real) spin-0 scalar, spin-1/2 (Weyl) fermion, and spin-1 vector boson, respectively (see Fig. 10).

Even without new light particles, N_{eff} is a probe of new physics that changes our thermal history, including processes that result in a stochastic background of gravitational waves [203, 204, 205]. Violent phase transitions and other nonlinear dynamics in the primordial Universe could produce such a background, peaked at frequencies much larger than those accessible to B -mode polarization measurements of the CMB or, in many cases, direct detection experiments such as LIGO and LISA [206, 207, 208, 209, 210]. For particularly violent sources, the energy density in gravitational waves can be large enough to make a measurable contribution to N_{eff} [210, 211, 212].

In addition to precise constraints on N_{eff} , cosmological probes will provide an *independent high-precision measurement of the primordial helium abundance* Y_{p} , due to the impact of helium on the free electron density prior to recombination. This is particularly useful since Y_{p} is sensitive to N_{eff} a few minutes after the Big Bang, while the CMB and matter power spectra are affected by N_{eff} prior to recombination, about 370,000 years later. Measuring the radiation content at these well-separated times provides a window onto any nontrivial evolution in the energy density of radiation in the early Universe [213, 214, 215, 216]. Furthermore, N_{eff} and Y_{p} are sensitive to neutrinos and physics beyond the Standard Model in related, but different ways, which allows for even finer probes of new physics, especially in the neutrino and dark sectors.

1.3.1.1 Observational signatures

The effect of the radiation density on the damping tail drives the constraint on N_{eff} in the CMB in $\Lambda\text{CDM} + N_{\text{eff}}$ models. The largest effect is from the impact on the mean free path of photons, which introduces an exponential suppression of short wavelength modes [217]. In a more detailed analysis, the effect on the damping tail is subdominant to the change to the scale of matter-radiation equality and the location of the first acoustic peak [175], both of which are precisely measured. As a result, the effect of neutrinos on the damping tail is more accurately represented by holding the first acoustic peak fixed. This changes the sign of the effect on the damping, but the intuition for the origin of the effect remains the same. At the noise level and resolution of CMB-S4, this effect is predominately measured through the TE spectrum at $\ell > 2500$.

In addition to the effect on the expansion rate, perturbations in neutrinos (and dark radiation) affect the photon-baryon fluid through their gravitational influence. The contributions from neutrino fluctuations are well described by a correction to the amplitude and the phase of the acoustic peaks in both temperature and polarization [176]. The phase shift is a particularly compelling signature, since it is not degenerate with other cosmological parameters [176, 177]. This effect is the result of the free-streaming nature of neutrinos, which allows propagation speeds of effectively the speed of light (while the neutrinos are relativistic). Any free-streaming light relics will also contribute to the amplitude and phase shift of the acoustic peaks. Furthermore, light relics that do not freely stream can in principle be distinguished from those that do by measuring their differing impacts on the damping tail and phase shift [177, 192].

Planck has provided a strong constraint on $N_{\text{eff}} = 2.92^{+0.18}_{-0.19}$ using temperature and polarization data [127]. Recently, the phase shift from neutrinos has also been established directly in the *Planck* temperature data [218]. This provides the most direct evidence for the presence of free-streaming radiation in the early Universe, consistent with the cosmic neutrino background.

1.3.1.2 Target

One of the features that makes N_{eff} a compelling theoretical target is the degree to which broad classes of models fall into two basic levels of ΔN_{eff} . As illustrated in the left panel of Fig. 11, any species that was in thermal equilibrium with the Standard Model degrees of freedom produces a characteristic correction to N_{eff} that depends only on its spin and its freeze-out temperature. For freeze-out after the QCD phase transition, one finds $\Delta N_{\text{eff}} \gtrsim 0.3$. Freeze-out before the QCD phase transition instead produces $0.07 > \Delta N_{\text{eff}} > 0.027$ per effective degree of freedom. The first case has been tested by the data from the *Planck* satellite. The second category, which is sensitive to freeze-out temperatures ranging from hundreds of MeV to the reheating temperature, falls into the level of sensitivity attainable by CMB-S4.




Figure 11. *Left: Constraints in the Y_p – N_{eff} plane from current and future cosmological surveys, compared to the predictions of standard BBN and current astrophysical measurements of Y_p [219]. Right: Inferred reach in T_F for a given sensitivity to N_{eff} . The vertical limits on T_F assume we are using a vector ($g = 2$) to relate 2σ N_{eff} limits to T_F . We see that at CMB-S4 sensitivity, constraints have a reach two orders of magnitude higher in T_F than either Planck or the Simons Observatory. Bands show error estimates based on Ref. [220].*

The right panel of Fig. 11 shows how a measurement at CMB-S4 sensitivity will dramatically improve our knowledge of physics at higher temperatures (i.e., at earlier times, weaker coupling). CMB-S4 at design sensitivity would place an upper limit of $\Delta N_{\text{eff}} < 0.060$ at 2σ , which translates into $T_F = 70$ (30) GeV for a light vector (Weyl fermion). This dramatic improvement in sensitivity to T_F over existing limits from *Planck* $T_F = 0.1$ (0.1) GeV, and projected constraints from Simons Observatory, $T_F = 0.3$ (0.3) GeV, is a direct result of the particle content of the Standard Model, which changes dramatically from quarks and gluons to baryons and mesons around the QCD phase transition (0.1–1 GeV) and thus significantly dilutes the energy densities of particles that decoupled at higher temperatures. Reaching $\Delta N_{\text{eff}} < 0.06$ –0.07 leads to dramatic improvements in the sensitivity to particles with spin. At this level of sensitivity, the zoo of particles present around the QCD phase transition is insufficient to dilute the energy below experimental sensitivity, and we see a very sharp corresponding rise in the T_F curves. Instead we are limited only by the dilution due to a few heavy particles and thus it is easier to make significant gains.

1.3.2 Neutrino mass

The 2015 Nobel Prize in Physics recognized the discovery of neutrino oscillations, which demonstrated that neutrinos have mass. However, the overall scale of the neutrino masses and some mixing parameters of the full set that determine neutrino oscillations are still not measured. Cosmology offers a unique view of neutrinos; they were produced in large numbers in the high temperatures of the early Universe and left a distinctive imprint in the cosmic microwave background and on the large-scale structure of the Universe. Therefore, CMB-S4 and the DESI and LSST surveys together will have the power to detect properties of neutrinos that complement those probed by large terrestrial experiments such as the Deep Underground Neutrino Experiment (DUNE), and those searching for beta-decay and neutrinoless double-beta decay.

Neutrino flavor oscillations—oscillations between electron, muon, and tau neutrinos—are described by a model where the three neutrino flavor states are a linear combination of three mass states. The matrix that

relates the two is the Pontecorvo-Maki-Nakagawa-Sakata (PMNS) matrix. The PMNS matrix depends upon six real parameters: three mixing angles, θ_{12} , θ_{23} , θ_{13} ; and three CP -violating phases, δ , α_1 , α_2 .

Experiments have measured the three mixing angles of U_{PMNS} and the two mass-squared splittings, $\Delta m_{21}^2 = 7.54 \times 10^{-5} \text{ eV}^2$ and $|\Delta m_{13}^2| \approx 2.4 \times 10^{-3} \text{ eV}^2$ [221]. Despite this, fundamental questions about neutrino mass and mixing remain: (i) the absolute mass scale; (ii) the relative ordering, or *hierarchy* of the masses; (iii) whether neutrinos are Dirac or Majorana particles; and (iv) and whether neutrinos violate charge and parity (measuring δ_{CP}). CMB-S4 will measure the sum of the neutrino masses,

$$M_\nu \equiv \sum m_\nu \quad (1.21)$$

with sufficient sensitivity to be relevant to these open issues. The value of M_ν unambiguously determines the absolute mass scale of neutrinos, and if M_ν is determined to be $\lesssim 0.1 \text{ eV}$, CMB-S4 can furthermore rule out the inverted mass hierarchy at increasing statistical significance depending on $\sigma(M_\nu)$. There are also circumstances in which the combination of CMB-S4 constraints and neutrinoless double-beta decay measurements can point to neutrinos being Dirac particles. Finally, the observables available to CMB-S4 to probe neutrinos can also provide information about new physics in the neutrino sector: the existence of new sterile neutrino states, neutrino interactions and decays, or a nonstandard thermal history for the Universe.

In the standard cosmology, neutrinos were in thermal equilibrium with photons, electrons, and positrons until the temperature of the Universe dropped below $T \sim 1 \text{ MeV}$. At this time the weak interaction rate became smaller than the expansion rate of the Universe and neutrinos decoupled. Assuming the standard thermal history, the neutrino temperature can be related to the CMB temperature, and we can predict a relic abundance of $n_\nu \approx 56 \text{ cm}^{-3}$ for each neutrino and antineutrino state. The CMB-inferred total radiation density in the early Universe is in excellent agreement with a contribution from neutrinos with sub-eV mass and this expected relic abundance. As discussed in 1.3.1, CMB-S4 constraints on the radiation density will put stringent tests on the thermal history of neutrinos.

The neutrino mass-square splittings give lower limits on the masses of two neutrino states of $\geq \sqrt{\Delta m_{21}^2} \approx 0.01 \text{ eV}$ and $\geq \sqrt{|\Delta m_{13}^2|} \approx 0.05 \text{ eV}$. These lower limits, in combination with the temperature T_ν , imply that the energy of at least two species of relic neutrinos is today dominated by their rest mass, rather than their momentum. In combination with the inferred relic abundance, this predicts a neutrino contribution to the energy budget of $\Omega_\nu h^2 \gtrsim 0.0006$. This Ω_ν contributes to the matter budget of the Universe at late times.

Since neutrinos were relativistic for much of the history of the Universe and still have relatively large kinetic energy today, their gravitational clustering is qualitatively different from that of cold dark-matter particles or baryons. On large scales, neutrino clustering is identical to that of CDM and baryons. On smaller scales, where the neutrino velocity is important, neutrinos free stream out of gravitational potentials, leaving the CDM and baryons behind; since the free-streaming neutrinos' energy still contributes to the expansion, this causes a suppression to the growth of structure on small scales. The scale separating these two regimes is the neutrino free-streaming scale, defined by [222, 223] $k_{\text{fs}}(a) \equiv \sqrt{\frac{3}{2}} a H(a) / v_\nu(a) \approx 0.04 a^2 \sqrt{\Omega_{\text{m}} a^{-3} + \Omega_\Lambda} (m_\nu / 0.05 \text{ eV}) h / \text{Mpc}$ in comoving coordinates. For a fixed CDM and baryon density, massive neutrinos thus induce a suppression in the matter power spectrum given by [222, 224, 225, 226]

$$\frac{P_{\text{mm}}(k \gg k_{\text{fs}} | \Omega_\nu)}{P_{\text{mm}}(k \gg k_{\text{fs}} | \Omega_\nu = 0)} \approx 1 - 6\Omega_\nu / \Omega_{\text{m}}. \quad (1.22)$$

Cosmological data sets are primarily sensitive to the net suppression, and therefore the neutrino energy density Ω_ν . Since the neutrino number density is separately constrained and is the same for each mass eigenstate, these measurements of Ω_ν determine M_ν .

A central goal of the CMB-S4 experiment is to achieve a robust detection of the neutrino mass for the minimum mass sum implied by oscillation data, $M_\nu = 58 \text{ meV}$ (or $\Omega_\nu h^2 = 0.0006$). In the next section, we will discuss methods through which CMB-S4 can directly constrain the neutrino mass.

1.3.2.1 CMB-S4 observables for M_ν

The main effect of massive neutrinos on large-scale structure is to suppress the matter power spectrum on scales smaller than the neutrino free streaming scale. By measuring the magnitude of this suppression, we can therefore determine the neutrino mass from cosmology.

Typically, this suppression is measured by comparing the amplitude of structure below the free-streaming scale to the nearly unsuppressed amplitude of structure seen in the CMB; however, a constraint on the optical-depth parameter τ is required to determine the primordial amplitude from the CMB, and the precision of this optical-depth constraint can therefore be a limiting factor. Beyond the typically limiting optical-depth degeneracy, another degeneracy arises with the matter density (on which structure growth also depends), though this is more easily overcome by combining CMB data with baryon acoustic oscillation measurements.

CMB Lensing: One of the cleanest methods to probe the matter power spectrum, and hence the neutrino mass, is CMB lensing. Along their path to our telescopes, the photons of the CMB are gravitationally deflected by the entire mass distribution through which they pass, remapping the unlensed temperature and polarization as $X(\hat{\mathbf{n}}) = X^{\text{unl}}(\hat{\mathbf{n}} + \nabla\phi(\hat{\mathbf{n}}))$ where the CMB $X = T, E, B$ and the lensing potential ϕ is a direct gravitational probe of the projected mass distribution. The lensing potential power spectrum is related to the projected matter power spectrum by $C_\ell^{\phi\phi} = \int_0^{\eta^*} d\eta W^2(\eta) P(k = \ell/\eta)$, where W is a geometric projection kernel that depends mainly on distance ratios and η^* is the distance to the CMB last scattering surface. CMB lensing is a relatively clean probe because it arises from linear or mildly non-linear scales that are minimally affected by baryonic feedback and because the source redshift is perfectly known. While non-Gaussianity introduced by extragalactic and Galactic foregrounds must be taken into consideration, such systematics are typically small to begin with and can be further mitigated by improved lensing estimators and multifrequency cleaning.

Lensing cross-correlations: A related probe that can also give insight into the neutrino mass is the measurement of cross-correlations of CMB-S4 lensing with upcoming galaxy surveys such as LSST or Euclid. Since a lensing-galaxy correlation $C_\ell^{\phi g}$ is proportional to bP_{mm} , where b is galaxy bias, an unknown astrophysical parameter. The galaxy power spectrum scales as $C_\ell^{\text{gg}} \propto b^2 P_{\text{mm}}$, the combination of these two probes provides (to some extent) an independent measure of the amplitude of structure, and hence of the neutrino mass. While the optical-depth degeneracy remains, it is somewhat reduced by the galaxy shape constraint in C_ℓ^{gg} , as well as the sensitivity to low-redshift growth. Furthermore, since the cross-correlations can be broken up into tomographic redshift bins, cross-correlation measurements can help break degeneracies between dark-energy parameters and neutrino masses. While some systematic errors are nulled in cross-correlation, extragalactic foregrounds remain in temperature; in addition, the measurement of galaxy clustering can come with considerable challenges such as photometric redshift uncertainty.

Galaxy Clusters: Via the thermal Sunyaev-Zeldovich effect, CMB-S4 will be able to find tens of thousands of galaxy clusters out to high redshifts. The masses of these galaxy clusters can be obtained via an SZ-mass scaling relation calibrated by both weak lensing (e.g., with LSST) and CMB lensing (internally). This will allow a determination of the mass function, which is a sensitive probe of the amplitude of structure on small scales, and hence will provide a precise probe of the neutrino mass. The redshift-dependence of the galaxy cluster abundance is expected to reduce degeneracy between neutrino and dark-energy parameters. Although clusters are complex objects with rich astrophysics complicating precise cosmological measurements, great progress has been made in mass calibration and in the mitigation of other uncertainties. Clusters represent a

physically very different probe to CMB gravitational lensing; the robustness of neutrino mass measurements will be substantially improved by having independent measurement techniques with entirely different sources of systematic errors.

1.3.2.2 CMB-S4 forecasts for M_ν

Forecasts within $\nu\Lambda$ CDM: In this section we will present the results of forecasts for the constraints on neutrino mass arising from multiple probes within a minimal model of Λ CDM + neutrino mass. The forecast errors on the sum of the neutrino masses from different CMB-S4 derived probes are shown in Fig. 12, plotted as a function of a prior assumed on the CMB optical depth. It can be seen that all probes perform quite similarly in the absence of an estimate of τ that is more precise than the current state of the art with *Planck* ($\sigma(\tau) = 0.007$). All probes are limited by the incomplete knowledge of the CMB optical depth; this well-known result strongly motivates improvements in external determinations of the CMB optical depth. An alternative possibility, albeit one that has not yet been demonstrated in data, is to make use of measurements of the kSZ trispectrum to constrain reionization and obtain an estimate of the optical depth. Preliminary estimates indicate that tight constraints of order $\sigma(\tau) = 0.002$ (and powerful neutrino mass measurements) might be possible with CMB-S4 kSZ, though such constraints would be less model-independent than those arising from the large-scale CMB. Small improvements in neutrino mass constraints can also be seen by adding probes with redshift resolution or the ability to probe the galaxy-spectrum shape, because these provide information that is not degenerate with τ .

We note that using only polarization lensing-reconstruction preserves nearly all the signal-to-noise ratio in a neutrino mass measurement. While we adopt the CMB lensing constraints as our conservative baseline, we find that for low $\sigma(\tau)$, CMB-S4 cluster-cosmology forecasts (with an LSST weak-lensing mass calibration) appear even more powerful than CMB lensing, motivating more research in this area.

In all cases, even for present-day optical-depth constraints, an approximately 3σ detection of the minimal mass sum is possible for CMB-S4. The fact that multiple different CMB-S4 probes should be able to achieve similar constraints will allow for cross-checks and thus increase the reliability of any claimed detection. Furthermore, unlike with prior experiments, such as Simons Observatory, we will make a near-equally significant detection using only lensing of CMB polarization; we do not need to rely on temperature-based reconstructions where extragalactic foreground systematics are a possible concern. CMB-S4 will thus provide definitive measurements of the neutrino mass sum that are more reliable than those from prior experiments. If the optimistic constraints on the optical depth of order $\sigma(\tau) \approx 0.002$ can be achieved, with either internal kSZ analyses or a large-scale cosmic variance limited, external CMB experiment, at least 5σ detections of the minimal mass sum will become possible with CMB-S4.

Forecasts with more general dark-energy assumptions: To examine the strength of the dependence of our constraints on the assumption of a particular cosmological model, we relax the assumption that the dark energy must be a cosmological constant and perform our forecasts again, this time freeing the dark-energy equation of state parameter w . The resulting constraint contours are shown in the right panel of Fig. 12. While all probes show some degeneracy between the neutrino mass and the dark-energy equation of state, this degeneracy is, to some extent, reduced in the three probes that resolve growth of structure as a function of redshift, resulting in improved neutrino mass constraints. That is because the different redshift dependences of the energy densities of dark energy and neutrinos allow the effects to be distinguished. We conclude that, especially with tomographic probes, powerful CMB-S4 measurements of neutrino mass are feasible even when allowing for somewhat more general dark-energy physics.

Before proceeding, it is useful to put a CMB-S4 detection of neutrino mass in context with anticipated constraints on M_ν from other cosmological data sets. At present, the tightest bounds on M_ν are achieved by




Figure 12. Left panel: Forecasted constraints on the neutrino mass sum for several different CMB-S4-derived cosmological probes, written here in terms of the significance of a detection of the minimum value consistent with oscillation data $M_\nu = 58\text{meV}$. Neutrino mass constraints are degenerate with the optical depth to the CMB, τ . Constraints here are shown as a function of the assumed $1 - \sigma$ errors on τ . All probes give similar results for the current values $\sigma(\tau) \approx 0.007$; constraints improve rapidly as the error on τ decreases. We have marked on the plot both the current τ constraint and an optimistic value derived from either kSZ analysis or an external large-scale CMB measurement. Constraints are derived from: the CMB-S4 lensing power spectrum (blue solid line); the lensing power spectrum using only polarization reconstruction (blue dashed); the CMB-S4 lensing power spectrum with the addition of LSST gold sample galaxy cross-correlations and clustering (red); CMB-S4 Sunyaev-Zel’dovich selected clusters with LSST weak lensing mass calibration (green). In all cases, DESI BAO and CMB-S4 primary CMB constraints are also assumed. Right panel: Forecast errors for joint constraints on neutrino mass and the dark energy equation of state. Probes with redshift resolution generally show somewhat less degeneracy of neutrino constraints with equation of state values, as well as smaller errors in general.

combining information in the primary CMB with a measure of the amplitude of structure at late times (in effect determining M_ν by comparing σ_8 and A_s). The amplitude of structure at late times can be determined from a variety of probes, including the CMB-S4 observables of CMB lensing and SZ cluster abundance discussed above, but also using data from galaxy surveys such as the abundance of galaxy clusters, the Lyman- α forest power spectrum, redshift-space distortions (RSDs), or weak-gravitational-lensing statistics. At present the tightest constraints on M_ν are from *Planck* CMB lensing, yielding $M_\nu < 120\text{meV}$ at 95% confidence [127]), with the latest data from the Lyman- α forest giving competitive constraints of $M_\nu < 140\text{meV}$ at 95% confidence. To date, other probes of the amplitude of structure give interesting but less stringent constraints, e.g., RSDs from the Baryon Oscillation Spectroscopic Survey (BOSS), $M_\nu < 230\text{meV}$ [227], and $< 260\text{meV}$ from DES lensing and galaxy clustering [228]. However, in the next decade a number of experiments are anticipated to constrain or detect M_ν at a level comparable to expectations for CMB-S4—for example, LSST, the Dark Energy Spectroscopic Instrument (DESI), *Euclid*, *WFIRST*, the Square Kilometer Array (SKA), and Simons Observatory are forecasted to reach 1σ limits on M_ν of order 10–40 meV when combined with current CMB data [136, 229, 230, 231]). In the event of no improvement in our knowledge of τ , each probe may only reach a $2\text{--}3\sigma$ detection of $M_\nu = 58\text{meV}$. Multiple independent approaches will therefore be critical for establishing the value of the neutrino mass from cosmology. The role of CMB-S4 is therefore as a particularly robust cross-check of any data on M_ν from other experiments and, in the event of improvements to our knowledge of τ , as an experiment that provides a high-significance measure of M_ν .

Scenario	$m_{\beta\beta}$	m_{β}	M_{ν}	ΔN_{eff}	Conclusion
Normal hierarchy	$< 2\sigma$	$< 2\sigma$	60 meV	0	Normal neutrino physics; no evidence for BSM
Dirac neutrinos	$< 2\sigma$...	350 meV	0	Neutrino is a Dirac particle
Sterile neutrino	$< 2\sigma$	$< 2\sigma$	350 meV	> 0	Detection of sterile neutrino consistent with short-baseline data
Diluted neutrinos	0.25 eV	0.25 eV	< 150 meV	< 0	Modified thermal history (e.g., late decay)
Exotic Neutrinos	0.25 eV	0.25 eV	< 150 meV	0	Modified thermal history; (e.g., neutrino decay to new particle)
Excluded	0.25 eV	0.25 eV	500 meV	0	Already excluded by cosmology
Dark radiation	$< 2\sigma$	$< 2\sigma$	60 meV	> 0	Evidence for new light particles; normal hierarchy for neutrinos
Late decay	$< 2\sigma$	$< 2\sigma$	60 meV	< 0	Energy injection into photons at temperature $T \lesssim 1$ MeV

Table 1-3. Relation between neutrino experiments and cosmology. We include the measurement of the Majorana mass via neutrinoless double-beta decay ($m_{\beta\beta}$) or a kinematic endpoint (m_{β}) compared to the cosmological measurement of the sum of the masses M_{ν} and the CMB measurement of N_{eff} . Here “ $< 2\sigma$ ” indicates an upper limit from future observations. For observations on the timescale of CMB-S4, we use $\sigma(m_{\beta\beta}) \approx 0.075$ eV and $\sigma(m_{\beta}) \approx 0.1$ eV. For ΔN_{eff} the use of ≥ 0 indicates a significant deviation from the Standard Model value.

1.3.2.3 CMB-S4 measurements of M_{ν} : detection scenarios and implications

CMB-S4 will make a cosmological measurement of M_{ν} that is complementary to terrestrial measurements of neutrino mass and other properties. In this section, we put the CMB-S4 measurement in context with laboratory neutrino experiments and present several detection scenarios and their implications.

Determining the neutrino mass scale: While CMB-S4 will determine M_{ν} through gravitational effects, terrestrial measurements of the neutrino mass using radioactive decay are kinematic and determine an effective electron (anti-)neutrino mass, which is an incoherent sum of mass eigenstates. Current measurements from Mainz [232] and Troitsk [233] limit this mass to < 2.0 eV. The Karlsruhe TRItium Neutrino (KATRIN) experiment [234] is expected to improve this limit by a factor of 10 and Project 8 may reach a limit of 0.04 eV [235]. Within the standard neutrino mass and cosmological paradigm, the kinematic and cosmological measurements of neutrino mass are connected through the PMNS neutrino mixing matrix. The combination of cosmological and terrestrial neutrino mass measurements can therefore test our cosmological neutrino model. A discrepancy could point to new physics, for example a modified thermal history through neutrino decay

Lepton number violation – Majorana or Dirac neutrinos: One of the most exciting connections between cosmological measurements of neutrino mass and terrestrial experiments is the complementarity between cosmological neutrino mass measurements and the search for neutrinoless double-beta decay (NLDBD). NLDBD is a hypothetical decay mode of certain nuclei, where two neutrons convert to two protons and two electrons with no emission of neutrinos. The observation of NLDBD would be transformational—




Figure 13. *Left:* Majorana effective neutrino mass $m_{\beta\beta}$ versus M_ν in the scenario where NLDBD is mediated by light neutrino exchange. The area enclosed by the blue and red solid lines indicate the allowed 95% ranges from neutrino oscillation experiments [236] for normal ordering (NO) and inverted ordering (IO) assuming complete ignorance of the Majorana phases. The vertical blue and red bands show the forecasted 1σ constraints on M_ν from CMB-S4 for minimal mass NO and IO. The horizontal band shows the sensitivity of future NLDBD experiments. A CMB-S4 detection of M_ν , in combination with a detection of $m_{\beta\beta}$ can constrain the Majorana phases. *Right:* Sum of individual the neutrino masses as a function of the electron-neutrino effective mass m_β for the NO (blue) and the IO (red). Again, the area enclosed by the blue and red solid lines indicate the allowed 95% ranges from neutrino oscillation experiments [236] for NO and IO. The horizontal bands show the future cosmological constraints around each ordering (assuming the mass of the lightest neutrino state $m_{\text{lightest}} = 0$ eV). Also shown are the anticipated limits on m_β from the KATRIN experiment in the case of no detection.

it would demonstrate that neutrinos are Majorana particles and reveal a new lepton-number-violating mechanism for mass generation. This new physics could potentially explain both the smallness of neutrino masses and matter-antimatter asymmetry in the Universe.

To illustrate the connection between NLDBD searches and M_ν , consider the simplest case where NLDBD is mediated by exchange of light Majorana neutrinos. Within the context of this mechanism, these experiments determine an “effective neutrino mass,” $m_{\beta\beta}$, given in terms of the PMNS mixing matrix, including two unknown Majorana phases, and the individual neutrino masses. In this scenario, the absence of a signal in next generation NLDBD searches, combined with a cosmological measurement constraining $M_\nu > 100$ meV (corresponding to either the inverted hierarchy or a minimum neutrino mass of 50 meV), would strongly point to neutrinos being Dirac particles (see Fig. 13). On the other hand, if NLDBD is observed, cosmological measurements of M_ν are then sensitive to the Majorana phases. Perhaps the most interesting scenario would be if cosmological and NLDBD measurements cannot be jointly explained by the exchange of light Majorana neutrinos and new physics is required.

Neutrino mass ordering and CP violation: In the case of normal ordering with non-degenerate neutrino mass, the CMB-S4 measurement of M_ν will provide a 2–4 σ determination of the neutrino mass ordering. Fully characterizing neutrino mass ordering and CP violation are major goals of the terrestrial neutrino physics program [237]. In the scenario where the neutrino mass spectrum is normally ordered and non-degenerate, CMB-S4 would be a strong complement to terrestrial experiments by providing a measurement of neutrino ordering that is independent of oscillation parameters and δ_{CP} . Under all circumstances, the combination of CMB-S4 with terrestrial determinations of neutrino ordering will provide a definitive measurement of the neutrino mass spectrum.

Sterile neutrinos: Mechanisms of introducing neutrino mass often include sterile neutrinos, with both Majorana and Dirac terms potentially contributing (e.g., Ref. [238]). A number of recent neutrino oscillation experiments have reported anomalies that are possible indications of four or more neutrino mass eigenstates with mass splittings $\mathcal{O}(1\text{ eV})$. To explain these anomalies, these neutrinos typically have relatively large mixing angles and would therefore thermalize in the early Universe, affecting primordial nucleosynthesis [239], changing N_{eff} by $\mathcal{O}(1)$, and M_ν by $\mathcal{O}(1\text{ eV})$. Introducing new neutrino interactions or a modified thermal history for neutrinos can accommodate additional light sterile neutrinos without violating current constraints on N_{eff} and M_ν . As discussed in the light relics section, CMB-S4 will provide additional constraints on neutrino-self interactions, potentially confirming or ruling out these scenarios.

Interestingly, there exist tensions in some combinations of CMB and LSS data sets that could be alleviated with the presence of massive neutrinos, extra neutrinos, or both. CMB-S4 could shed light on the sterile neutrino mass and vacuum flavor-mixing parameters invoked to explain the experimental neutrino anomalies. Telltale signatures in N_{eff} , M_ν , and Y_p can allow CMB-S4 to probe this larger parameter space.

1.3.3 Dark energy

The discovery 20 years ago that the expansion of the Universe is accelerating presented a profound puzzle to physics, one that we have not yet solved. Our current framework can explain these observations only by invoking a new substance with unique properties (dark energy) or by changing the century-old, well-tested theory of general relativity developed by Einstein.

As the current epoch of acceleration is much later than the epoch from which the photons in the CMB originate, the behavior of dark energy or modifications of gravity do not significantly influence the properties of the primordial CMB. However, the CMB can still inform us about the properties of dark energy through two fundamental pathways. First, the projection from physical scales to the angular scales observed in a CMB map involve distance, which is affected by the expansion history along the line of sight to the last scattering surface. It is therefore sensitive to dark energy. The other pathway, involving secondary anisotropies, is even more important. These include weak gravitational lensing, as well as interactions with free electrons leading to the thermal and kinematic Sunyaev-Zeldovich effects. Therefore, as discussed in the dark-energy submission [240] to the 2020 Decadal Survey of Astronomy and Astrophysics (and endorsed by all major dark energy experimental collaborations), the CMB is a recognized probe of dark energy.

In the simplest model acceleration is driven by a cosmological constant. Although theoretically unnatural, this model does satisfy current constraints, so a clear target for CMB-S4 is to test the various predictions this model makes at late times. Using gravitational lensing of the CMB, the abundance of galaxy clusters, and cosmic velocities, CMB-S4 will measure both the expansion rate H and the amount of clustering, quantified by the parameter σ_8 , as a function of time. The constraints from CMB-S4 alone will be at the 0.1% level on each and, when combined with other experiments, will reach well below this level, particularly when the power of CMB-S4 is also utilized to calibrate other probes. CMB-S4 constraints will be among the most powerful tests of the cosmological constant—more crucially, the simultaneous sensitivity to the expansion and the growth will allow us to distinguish the dark-energy paradigm from modifications to general relativity. There are many models for acceleration in the latter class, and CMB-S4 will be generically useful in constraining them.

1.3.3.1 Canonical probes: lensing-convergence power spectrum and galaxy clusters

In the basic Λ CDM cosmological model, the dark energy is assumed to be the energy of the vacuum, or equivalently an inert component with equation of state $w = p/\rho = -1$. The most common phenomenological extensions of this model are models with an expansion history deviating from the Λ CDM model in a way that is parameterized by the dark energy having a time-varying equation of state. A common parameterization is

$$w(z) = w_0 + w_a \frac{z}{1+z}. \quad (1.23)$$

Changing the energy content of the Universe has two main consequences. The first is that the expansion history changes and therefore distance measures as a function of redshift change. This is how most of the low-redshift probes, such as baryon acoustic oscillations (BAOs) and type Ia supernovae probe dark energy. For CMB-S4 this effect affects the projection of the correlations at the surface of last scattering into the angular power spectrum. For example, changing w while keeping other parameters fixed will shift the positions of acoustic peaks in the CMB power spectrum. However, the precise measurement of the basic cosmological parameters in the presence of free w and w_a interacts in a non-trivial way with the low-redshift constraints of dark energy, which leads to improved constraints on the w - w_a plane through changing both the shape and amplitude of the observable CMB temperature and polarization power spectra, as well as breaking degeneracies with other parameters. For example, precise determination of parameters relevant for the pre-recombination era affects the precision with which the BAO ruler is known.

In Fig. 14 we show the improvement of constraints once CMB-S4 spectra (both CMB and lensing reconstruction) are added to more direct probes of dark energy. This specific forecast was made using the `GoFish` package⁵ [241] and we show three representative cases. The first is the determination of expansion history from the DESI experiment based on measurements of BAOs [242], arguably the most robust probe of dark energy. In blue we show the same contour based on “3×2-point” analysis of the LSST data [243, 244], i.e., the combination of the photometric galaxy clustering, galaxy-shear cross-correlations, and shear-shear auto-correlations. This probe has a very different set of systematic errors, but a comparable total constraining power. We note that both surveys will produce further constraints. For example DESI will perform analysis of the full shape of the redshift-space power spectrum while LSST will also use other probes, such as supernovae, strong lensing and galaxy clusters. Nevertheless, this choice presents a nominal set of constraints. Upon addition of the CMB-S4 temperature, polarization, and lensing reconstruction power spectra, these constraints tighten significantly, which we show as the red ellipse in Fig. 14. We quantify this in terms of a figure of merit (FoM) that is defined to be inverse area under the ellipse ($1/\sqrt{|C|}$, where C is the 2×2 C matrix of w , w_a), which is very similar in spirit to the Dark Energy Task Force figure of merit [245]. We additionally let the neutrino mass $\sum m_\nu$ be a free parameter (see also Sect. 1.3.2), since it is known to be degenerate with dark-energy parameters [246]. We find that *Planck*+DESI BAO and LSST 3×2 analyses have comparable FoMs of approximately 70. The combined FoM of these experiments is around 200, consistent with relatively modest degeneracy breaking. However, it increases to 400 with the addition of CMB-S4 data. CMB-S4 hence contributes to a doubling of the figure of merit. We want to emphasize that part of this improvement is due to ability of CMB-S4 to break the degeneracy with the neutrino mass.

Besides the effect on the expansion history, dark energy also affects the rate at which structures form in the Universe. Most importantly, in standard general relativity, these are intrinsically linked—the given expansion history determines the growth rate. By measuring the growth over cosmic history we can in principle falsify this picture, which would be a strong indication of new physics in the gravitational sector and a truly revolutionary discovery.

CMB-S4 will measure growth using the following two main methods.

⁵<https://github.com/damonge/GoFish>

- **CMB lensing reconstruction:** Change in the growth of structures will affect the amplitude (and to a lesser extent also the shape through geometric effects) of the CMB lensing convergence power spectrum. This method is particularly powerful in combination with auto-spectra of low redshift tracers as we show in Fig. 15;
- **Galaxy clusters:** The abundance and clustering of galaxy clusters are independent probes of growth. CMB-S4 will find clusters through the tSZ effect using proven matched-filtering techniques. The tSZ is a powerful tool to find and count clusters for CMB-S4 because the detection efficiency is nearly independent of redshift for an instrument with arcminute-scale beams, and the selection function is well behaved and simple to model. The utility of cluster abundances as cosmological probes is limited by systematic uncertainties in the the observable-to-mass scaling relations. CMB-S4 will use CMB halo lensing (or optical weak lensing) to calibrate the observable-to-mass relation, aided by the well-understood selection function. The tSZ selected clusters will require optical surveys to confirm and provide redshifts for the low redshift clusters ($z \lesssim 1.5$) and near-IR follow-up observations for the remaining clusters.

We show how CMB lensing reconstruction can be converted into redshift-resolved measurements of the clustering amplitude σ_8 in the left panel of Fig. 15.




Figure 14. Improvement on the standard w_0 - w_a parameters for the combination of: (i) Planck prior and expansion history measurements (DESI BAO) in black; (ii) LSST 3×2 -point function measurements (including auto- and cross-correlation of galaxy number density and shear field fluctuations) in blue; and (iii) the combination of the first two items, with CMB-S4 power-spectrum measurements (temperature, polarization, and weak lensing reconstruction, but importantly no galaxy cluster information) in red.

1.3.3.2 Kinematic Sunyaev-Zeldovich (kSZ) effect

The kSZ effect is the Doppler shift of CMB photons caused by scattering off the plasma in late-time galaxies and clusters, which are moving with respect to the CMB. This Doppler shift causes a slight change in the measured CMB temperature, while preserving the blackbody spectrum, and therefore cannot be isolated by a multi-frequency analysis, unlike the tSZ effect. However, the broad frequency coverage of CMB-S4 is ideal for reducing other foregrounds and thus the noise on the kSZ signal.

The kSZ-induced temperature fluctuation is proportional to the galaxy *radial* peculiar velocity v_r , and its optical depth to Thomson scattering τ_g (proportional to the number of free electrons in the galaxy itself): $(\Delta T/T)_{\text{kSZ}} = -\tau_g v_r$.

Thanks to its high resolution, CMB-S4 will allow unprecedented measurements of velocity fields through the kSZ effect. A measurement of the peculiar velocity can in turn constrain the amplitude of matter density fluctuations $\sigma_8(z)$ at the effective redshift of the sample, as well as the linear growth factor f , which is sensitive to dark-energy or modified-gravity models [247, 248, 249, 250]. These measurements will be able to distinguish dark energy from interesting models of modified gravity and will provide complementary constraints to redshift-space distortions and weak lensing measurements, probing larger physical scales. Figure 15 shows constraints on $f\sigma_8$ expected from CMB-S4, together with a galaxy sample from the upcoming DESI survey. We assume an overlap f_{sky} of 0.2, resulting in a total survey volume of 116 Gpc^3 , containing 19.6 million galaxies from all DESI galaxy samples.

One caveat is that the amplitude of the signal is proportional to the optical depth of the galaxy sample used in the analysis, τ_g . This is not known a priori and is in principle degenerate with $f\sigma_8$, as can be seen in the “clustering+kSZ” curve in Fig. 15. There are several proposed ideas for breaking the optical-depth degeneracy, which range from tSZ or X-ray measurements of the galaxies/clusters themselves [251, 252] to measuring RSDs in conjunction with the kSZ effect [253]. The effect of degeneracy breaking of the optical depth can be seen in the blue solid curve in Fig. 15. Scale-dependence in growth, either caused by dark-energy perturbations (for example from a non-standard speed of sound), or by screening mechanisms in modified gravity, is not subject to the optical-depth degeneracy and can be constrained directly.

In addition, kSZ measurements can be important for dark-energy studies in a more indirect way. The high resolution of CMB-S4 allows for a direct measurement of the optical-depth profile $\tau_g(\theta)$, since the large-scale velocity field is constant over the size of a galaxy. In turn, $\tau_g(\theta)$ is proportional to the gas profile of the galaxy sample, allowing for calibration of baryon effects on the power spectrum, a leading systematic for weak lensing surveys aimed at measuring dark energy [254, 255].

1.3.3.3 Cosmic birefringence

In addition to probing the global dynamics of dark energy, as discussed in previous sections, CMB-S4 will also be sensitive to the imprint of new parity-violating physics within the dark-energy sector: TB and EB correlations that arise from the effect of cosmic birefringence. Detection of these parity-violating correlations would have paradigm-changing implications for cosmological scale physics, and may present a unique handle to probing microphysics of dark energy.

The simplest dynamical way to model the accelerated expansion of the Universe is to invoke a new slowly evolving scalar field that dominates its energy budget (the quintessence models for dark energy). Such a field could couple to photons through the Chern-Simons term in the electromagnetic Lagrangian, causing rotation of the linear polarization of photons propagating cosmological distances—an effect known as “cosmic birefringence” [258]. In the case of the CMB, such rotation converts some of the primordial E mode into B




Figure 15. Constraints on the growth parameter from CMB-S4 from two independent sets of measurements. Left panel: Constraints on the matter amplitude σ_8 in tomographic redshift bins (indicated by the positions of points) from the combination of LSST galaxies and CMB-S4 lensing, assuming a fixed Λ CDM cosmology. Relaxing this assumption does not diminish our ability to measure departures from the fiducial model. Right Panel: Constraints on $f\sigma_8$ from the kSZ effect as a function of the size of prior on τ_g assuming a single redshift bin centered at $z = 0.75$. The straight dashed line shows results from RSDs coming from the DESI experiment. The signal from kSZ in galaxy clustering alone is plotted in orange and the combination of everything in blue. We see that even in the absence of an informative prior, the constraining power is twice that of DESI alone. For a sufficiently tight prior on τ_g , the CMB-S4 data alone can surpass DESI RSD measurements.

modes, producing characteristic TB and EB cross-correlations in the CMB maps [259, 260]. Even though there is no firm theoretical prediction for the size of this effect,⁶ if observed, it would be a clear “smoking-gun” evidence for physics beyond the Standard Model, in the form of a new scalar field.

Previous studies have constrained both a uniform rotation angle α , as well as anisotropic rotation described by a power spectrum. The current tightest bound on the uniform rotation angle is $\alpha < 0.5^\circ$ at 68% CL derived from *Planck* data [261]. The effect of a uniform cosmological polarization rotation is degenerate with systematic uncertainty on the overall instrument polarization-angle calibration. Assuming this overall angle can be calibrated to arbitrary precision, CMB-S4 would improve the constraint on uniform cosmological rotation to $\alpha < 0.2'$.

A promising way to pursue the search for cosmic birefringence is via the quadratic estimator formalism [262], which explores the off-diagonal (mode-coupling) EB correlations on small angular scales. The existing constraints were derived under the assumption of a scale-invariant rotation spectrum, for which $A \equiv L(L+1)C_L^{\alpha\alpha}/2\pi$ is independent of L . Such spectra could originate from fluctuations in a spectator scalar field present during inflation [263]. The best current bound, obtained from sub-degree scale polarization measurements with BICEP2/Keck [257], is $A < 0.11 \text{ deg}^2$ (and essentially the same limit from *Planck* data [264]). More accurate measurements of polarization anisotropy at higher resolution and over a wider range of scales will significantly improve on this.

Figure 16 shows the forecasted noise (assuming no rotation) in the rotation-angle power spectrum $C_L^{\alpha\alpha}$ for CMB-S4 (with effective noise of $1.81 \mu\text{K-arcmin}$, and a resolution of $1.4'$), along with the predicted 95% CL

⁶One exception, for example, are the “Axiverse” models that arise in the context of string theory and feature scalar fields that produce sky-averaged rotation angles on the order of $1/137$ radians (see Ref. [146]).




Figure 16. Black lines represent model-independent projected noise for the cosmic-birefringence rotation-angle spectrum for CMB-S4. The noise assumes no rotation signal and is calculated in three different ways: (a) assuming no delensing and using the forecasted noise in the ILC (dot-short dash line); (b) assuming 80% delensing and with forecasted noise remaining after the foreground subtraction (solid line); and (c) assuming 100% delensing and perfect foreground subtraction (long dash line). The colored lines show an example of a specific model for birefringent rotation—the scale-invariant power spectrum. The amplitude of this rotation power spectrum is set to the current 95% CL bound from POLARBEAR [256] (green dot) and BICEP2/Keck [257] (blue short dash), and to the projected bound for CMB-S4 with the noise calculated under assumption (b) (red dot-long dash line).

bound on the amplitude A of the scale-invariant spectrum. The improvement on the current constraints is several orders of magnitude.

1.3.4 Dark matter

CMB measurements have great power in testing dark matter (DM) models and constraining parts of the parameter space that are inaccessible to laboratory experiments. In particular, the CMB directly probes the physics of cosmological DM throughout cosmic history and does not rely on assumptions about the local DM phase-space distribution within the Milky Way. Given current null results of targeted searches for well-motivated candidate models, broad scans of all possibilities are warranted; CMB-S4 will enable such an approach to the DM problem. We highlight three broad classes of DM models that are of particular interest to searches with CMB-S4.⁷ In Sect. 1.3.4.1, we discuss light DM that interacts with baryons, in Sect. 1.3.4.2, we discuss models that involve DM interactions with other new light degrees of freedom, and in Sect. 1.3.4.3, we focus on ultra-light axion-like DM particles.

⁷In this document, we omit discussion of DM annihilation, since the detection sensitivity is mostly saturated by current and upcoming measurements [49].

1.3.4.1 Scattering with baryons

Traditional nuclear-recoil-based direct-detection experiments are only sensitive to WIMPs (weakly-interacting massive particles) with masses above about a GeV, and with such low interaction cross-sections⁸ that they can penetrate the heavy shielding of detector targets [267, 265, 266]. Reducing the amount of shielding can lift the “ceiling” on the sensitivity of direct searches toward higher cross-sections [268], and new strategies are being explored to expand their sensitivity to low DM particle masses, below a GeV [195]. An entirely complementary way to probe sub-GeV DM is to search for evidence of its interactions in cosmological data. Since lower DM particle masses translate to a higher number density of scattering centers, the CMB is particularly sensitive to light particles. In addition, the CMB is sensitive to all cross-section magnitudes near and above the nuclear scale.

In the scenario in which DM scatters with protons in the early Universe, a drag force between the two cosmological fluids damps the acoustic oscillations and suppresses power in density perturbations on small scales. As a result, the CMB temperature, polarization, and lensing power spectra are suppressed at high multipoles, with respect to those in a Λ CDM universe (see the left panel of Fig. 17 for illustration). This effect has been used to search for evidence of DM-proton scattering for the case of heavy DM, using CMB and Lyman- α forest measurements [269, 270, 271].

Recently, Ref. [272] presented the first cosmological search for DM particles with any mass down to a keV (orders of magnitude below the mass thresholds of direct-detection experiments), and was followed by a number of related studies [273, 274]. In particular, Ref. [275] used *Planck* data to derive the first CMB limits on the non-relativistic effective theory of DM-proton scattering—a framework developed by the direct-detection community to characterize all available phenomenologies for scattering through a heavy mediator [276, 277, 278, 279]. Similarly, Refs. [280, 281, 274] have searched for interactions that can be parameterized by a power-law dependence of the interaction cross-section on the relative particle velocity, and reported improved limits on a wide range of models. Furthermore, Ref. [280] developed an improved treatment of non-linear effects that arise in the calculation of post-recombination scattering signals. This and related studies have enabled a robust investigation of a scenario where only a fraction of DM interacts—and tightly couples—with baryons (while the rest behaves just like the standard cold DM fluid), leading to the first robust cosmological limits on an interacting DM sub-component [280, 282]. In addition, analyses of CMB data provided essential consistency-tests of recent claims that the anomalous 21-cm signal reported by the EDGES collaboration [283] could be explained with late-time DM-baryon scattering [284]; see, for example, Refs. [285, 281].

These state-of-the-art CMB tests of DM-baryon interaction physics, which resulted in a broad exploration of new parameter space, relied primarily on moderate-resolution temperature measurements from *Planck*. Since the DM signal is increasingly prominent at higher multipoles (see Fig. 17), high-resolution measurements with CMB-S4 will thus substantially improve sensitivity to DM-baryon interactions.

In Fig. 17, we present current and projected upper limits on the cross-section for scattering on protons, as a function of DM mass, for a spin-independent velocity-independent interaction (chosen as our fiducial model). Areas above the curves are excluded at the 95% confidence level. We compare current limits obtained from *Planck* (from Ref. [272]) with forecasts for Stage-3 experiments (with a white noise level of $10\mu\text{K-arcmin}$ and resolution of $1.4'$), CMB-S4 (with specifications listed in Sect. 4), and a cosmic-variance-limited experiment (with vanishing noise levels). We note the large improvement factor that CMB-S4 delivers over the current limits, for the entire DM mass range considered.

⁸For reference, one of the current largest experiments of this class, Xenon1T, is sensitive to cross-sections roughly in the range 10^{-47} – 10^{-31} cm^2 [265, 266].

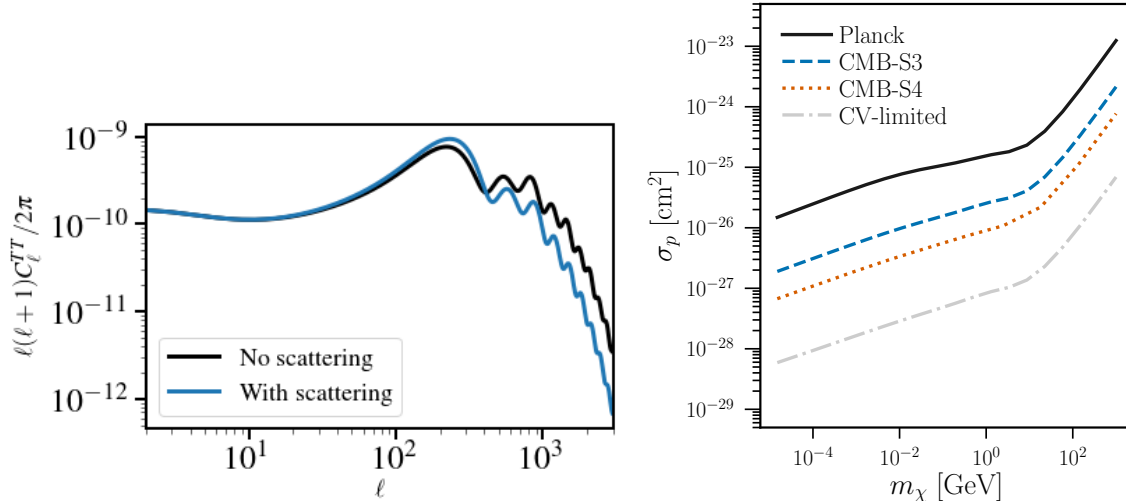


Figure 17. *Left:* Illustration of the effect of a velocity-independent spin-independent contact interaction between dark matter and baryons (with a cross-section 100 times higher than the current upper limit from Planck) on the CMB temperature power spectrum (blue), compared to the CDM case (black). *Right:* Upper limits on the DM-proton interaction cross-section as a function of DM mass, for spin-independent velocity-independent scattering. In case of a null detection, areas above the curves are excluded at the 95% confidence level. Shown are the current limits from Planck [272] and forecasts for Stage-3 experiments (CMB-S3, such as AdvACT [287]), CMB-S4, and a cosmic-variance-limited experiment (CV-limited); see also [286].

Most of the constraining power in the case of CMB-S4 comes from the lensing anisotropy, while the temperature and polarization anisotropies contribute roughly equally to the projected constraint [286]. Furthermore, Ref. [136] has shown that an increase in sky coverage beyond about 10% (keeping all other parameters fixed) only marginally improves the projected sensitivity. Finally, Ref. [286] has also shown that DM-baryon scattering is easily distinguishable from most other new-physics effects sought by CMB experiments (the neutrino mass, new light degrees of freedom, and DM annihilation) once the lensing anisotropy is measured at the level of CMB-S4. Therefore, beyond its ability to cover vast open portions of DM parameter space, CMB-S4 also holds promise as a DM discovery tool.

Small-scale CMB anisotropy measurements enabled by CMB-S4 will probe DM-baryon interactions at the time when the Universe was much less than a thousand years old. These measurements will be sensitive to particle masses outside the detection limits of most existing direct-detection experiments; furthermore, unlike all Earth-based experiments, CMB analyses are independent on the assumptions about the local phase-space distribution of DM. Compared to the Simons Observatory forecasts [136], CMB-S4 will further improve the sensitivity to DM scattering by a factor of 4–5 (for a DM mass of 1 GeV and a velocity-independent interaction)—sufficient to enable signal confirmation and detailed subsequent studies, in the case of a marginal detection with future Simons-Observatory data. On the other hand, CMB-S4 will provide important consistency checks for small-scale-structure probes of DM microphysics, enabled by upcoming galaxy surveys such as LSST. In particular, DM interactions can leave imprints on satellite galaxy populations [288] and other collapsed structures in the local Universe [289], many of which will be measured in detail in the coming decade of observations. However, tests of DM microphysics with small-scale-structure tracers face modeling and simulation challenges; CMB measurements that can capture the same physical effects on large scales and in the early Universe will be pivotal for robust inference of DM particle properties from observational data sets in the future.

1.3.4.2 Interactions with dark radiation

The exquisite sensitivity of the CMB to the depth and size of the DM gravitational potentials near the surface of last scattering makes it a particularly good probe of any new physics affecting the clustering of DM on large scales at early times. This sensitivity to DM density fluctuations is extended to lower redshifts via the weak gravitational lensing that CMB photons experience as they propagate to us. Similar to how tight coupling with photons inhibits the growth of baryon fluctuations until the epoch of hydrogen recombination, DM interacting with light (or massless) dark radiation (DR) at early times experiences a suppressed growth of structure due to the dark-radiation pressure opposing gravitational infall. Models where such interactions arise are diverse in their particle content (see e.g., Refs. [290, 291, 188]) and have been invoked to explain the apparent low amplitude of matter fluctuations measured by certain weak-lensing surveys [166, 189, 190]. They can also naturally arise in the context of self-interacting DM, which has been proposed to address possible anomalies on subgalactic scales [292].

The impact of DM-DR interaction on the CMB has been studied in detail in Refs. [291, 293] (see also Ref. [294] for a detailed derivation of the relevant Boltzmann equations). In short, the presence of extra DR mimics the presence of extra neutrino species (see Sect. 1.3.1) and affects the expansion history of the Universe, possibly modifying the epoch of matter-radiation equality, the CMB Silk damping tail, and the early integrated Sachs-Wolfe effect. However, unlike standard free-streaming neutrinos, the DR forms a tightly coupled fluid at early times, leading to distinct signatures on CMB fluctuations which include a phase and amplitude shift of the acoustic peaks (see Refs. [176, 218, 177]). In addition, the DR pressure prohibits the growth of interacting DM fluctuations on length scales entering the causal horizon before the epoch of DM kinematic decoupling. This weakens the depth of gravitational-potential fluctuations on these scales, affecting the source term of CMB temperature fluctuations. Finally, the modified matter clustering in the Universe due to the interaction with DR affects CMB lensing. For interacting DM models that are still allowed by the current *Planck* data, this latter effect is where CMB-S4 can significantly improve the constraints on these non-minimal DM theories.

Given the large array of possible DM theories to constrain, it is useful to pick a simple benchmark DM-DR model to assess the sensitivity of CMB-S4 to the effects described in the previous paragraph. A particularly useful model is one in which the interaction between DM and DR is mediated by a particle with mass of order a few MeV. Such theories have been put forward in the context of self-interacting DM [295, 296, 297] in order to obtain a velocity-dependent self-interaction cross-section that can fit both the observed density profiles of dwarf galaxies and of large galaxy clusters. For these models, the momentum-transfer rate between the DR and DM in the early Universe can be written as $\dot{\kappa}_{\text{DR-DM}} \propto \Omega_{\text{DM}} h^2 a_4 (1+z)^4$, where a_4 is a parameter that controls the strength of the DM-DR interaction that depends on the exact Lagrangian used (see Ref. [294]). Within this class of models, the most relevant parameters that can be constrained with CMB-S4 are the interaction strength a_4 , the fraction of DM interacting with DR f_{int} , and the amount of DR present in the Universe $\Omega_{\text{DR}} h^2$.

In Fig. 18, we show the projected sensitivity of a CMB-S4-like experiment to the presence of DM-DR interactions in the $\Omega_{\text{DR}} h^2 - a_4$ plane, for two values of f_{int} , namely 1% and 10%. For low values of the interaction strength a_4 , the DR decouples very early from DM and essentially behaves like an extra free-streaming relativistic species, and the constraint on $\Omega_{\text{DR}} h^2$ thus becomes equivalent to that on ΔN_{eff} discussed in Sect. 1.3.1. As the DM-DR interaction strength is increased, $\Omega_{\text{DR}} h^2$ becomes more severely constrained, to ensure that DM kinematic decoupling occurs before the modes probed by the CMB enter the causal horizon. In addition, the constraints become stronger as the fraction of interacting DM increases.

Looking ahead, the main difficulty in constraining theories in which DM interacts with DR using CMB-S4 will be modeling of non-linearities in the CMB lensing power spectrum, which can be important for multipoles




Figure 18. Dark matter-dark radiation interaction sensitivity of Planck (adapted from Ref. [291]) and an experiment like CMB-S4; for this plot, we assumed a configuration of CMB-S4 similar to that outlined in the CDT. The x-axis corresponds to the interaction strength, while the y-axis corresponds to the abundance of dark radiation. We show two different fractions of DM interacting with the DR, as indicated on each panel. The shaded regions are allowed at the 95% confidence-level. Note the different scales for the y-axes on each panel.

above 500 (in the reconstructed lensing-potential power spectrum). Recent progress in this respect is already promising [298]. The CMB constraints on DM-DR interaction could be further improved by combining CMB-S4 with probes of small-scale structure, such as the Lyman- α forest or the luminosity function of Milky Way satellite galaxies (see e.g., Ref. [299]).

1.3.4.3 Axion-like particles

The QCD axion and other axion-like particles (ALPs) are well-motivated DM candidates and can contribute to the DM density (see Ref. [141] for a recent review). One example is ultralight axions (ULAs), which are non-thermally created via vacuum realignment and have a distinctive phenomenology. ULAs are predominantly non-thermal, nonrelativistic, and do not contribute to N_{eff} .⁹

The ULAs we consider here are motivated by string theory and are associated with the geometry of the compact spatial dimensions. These axions can contribute either to the dark-matter or dark-energy budget of the Universe, depending on their particular mass, which sets the time at which the axions begin to coherently oscillate and redshift as matter. We consider axions within a range of masses $10^{-33} \text{ eV} \leq m_a \leq 10^{-20} \text{ eV}$, with negligible couplings to the Standard Model particles. We can compare these assumptions to those used in Sect. 1.3.1; the contribution of thermal axions to N_{eff} applies to any mass $m_a \lesssim 1 \text{ eV}$, including the well known QCD axion, but depends in detail on the couplings to the Standard Model particles and on the reheating temperature. Cosmological constraints on axion dark matter are thus complementary in the space of masses and couplings to the constraints on a thermal population of axions.

At the moment, the CMB (through measurements of CMB lensing, temperature, and polarization [300], using tools developed in Refs. [301, 302]), offers the best gravitational probe of the ULA density in the regime $10^{-33} \text{ eV} \leq m_a \leq 10^{-25} \text{ eV}$. We choose a fiducial value of the axion energy density consistent with these

⁹Note, however, that if these axions have couplings to ordinary matter, as described in Sect. 1.3.1, then a second, relativistic population of axions is created.




Figure 19. Constraints on ultra-light axions (ULAs). Left: Fisher-forecasted 2σ exclusion regions for the ULA mass fraction Ω_a/Ω_d for the Planck alone, Planck + Simons Observatory (SO) (as discussed in Ref. [49]), and CMB-S4, where $\Omega_d = \Omega_a + \Omega_c$. Right: Residual in lensing-convergence power spectrum C_L^{NS}/C_L^{RK} for different models including ULA DM, compared with Fisher-forecasted errors for CMB-S4.

constraints. In this window, ULAs cannot be all of the dark matter, but can be a significant component, with density comparable to baryons and neutrinos. Given the rich spectrum of particles in the Standard Model, there is no reason to expect a trivial dark sector, and it is important to explore the power of the CMB to constrain any particle species with a potentially detectable density.

CMB-S4 will push the upper edge of the CMB ULA window to $m_a \sim 10^{-23}$ eV. At higher masses still, ULAs will alter pulsar timing signatures [303, 304], suppress the clustering of neutral hydrogen at high- z (an observable probed by measurements of the Lyman- α forest flux power spectrum [305, 306, 307, 308, 309]), alter the mass spectrum of black holes through Penrose processes [146, 310, 141, 311, 312] with implications for Laser Interferometer Gravitational-Wave Observatory (LIGO) event rates, and lead to observable gravitational-wave signatures in the *Laser Interferometer Space Antenna (LISA)* band [313, 310, 311, 312] (see Ref. [314] for a recent review of ULA gravitational signatures).

1.3.4.4 Constraints on cold axion energy density

The degeneracies of ULAs with other cosmological parameters, such as N_{eff} or m_ν , vary depending on the axion mass [302]. Dark-energy like axions with masses around 10^{-33} eV change the late-time expansion rate and therefore the sound horizon, changing the location of the acoustic peaks. There are thus degeneracies of the ULA density Ω_a with the matter and curvature content. Heavier axions ($m_a \gtrsim 10^{-26}$ eV) affect the expansion rate in the radiation era and reduce the angular scale of the diffusion distance, leading to a boost in the higher acoustic peaks, which has a degeneracy with N_{eff} .

In the matter power spectrum, and thus CMB lensing power, light axions suppress clustering, suggesting a degeneracy with effects of massive neutrinos that must be broken to make an unambiguous measurement of neutrino mass using the CMB. The above-mentioned effects in the expansion rate break this degeneracy for some axion masses. There remains a significant degeneracy between Ω_a and the sum of the neutrino masses ($\sum_i m_{\nu_i}$) and massive neutrinos if $m_a \lesssim 3 \times 10^{-29}$ eV. Effort should be made to break these (and other) degeneracies using other data sets to establish if CMB-S4 data will show the unambiguous signature of neutrino mass, or hint at some other new physics.

We show the forecasted sensitivity to the axion energy density from CMB-S4, including lensing in the left panel of Fig. 19 (for fixed neutrino mass of $\Sigma m_\nu = 0.06$ eV). Adding information from the lensing reconstruction using CMB-S4 (or SO) will improve sensitivity to axion DM significantly. A percent-level measurement of the lensing deflection power at multipoles $\ell > 1000$ leads to an improvement in the error on the axion energy density of a factor of around 8 relative to the current *Planck* constraints, for an axion mass of $m_a = 10^{-26}$ eV. This will allow us to test if DM has different particle components at the percent level. Furthermore, since $\Omega_a \propto f_a^2$, this improves the expected constraint on the axion decay constant from 10^{17} GeV with *Planck* to 10^{16} GeV with CMB-S4, testing the “string axiverse” scenario [146].

Using only *Planck* data, ULAs are degenerate with CDM at $m_a = 10^{-24}$ eV, and constraints are still weak at $m_a = 10^{-25}$ eV. CMB-S4 could make a $> 5\sigma$ detection of departures from CDM for masses as large as $m_a = 10^{-25}$ eV, and could improve the lower bound on DM particle mass to $m_a = 10^{-23}$ eV and fractions $\mathcal{O}(10\%)$. Realizing this sensitivity will require better modeling of the non-linear clustering of axions [315], building upon and applying the results of previous work [316, 317, 318, 319, 320, 321, 322, 323, 324, 325, 326, 327, 328, 329, 330, 331, 332].

1.3.4.5 Axion isocurvature

The axion decay-constant, f_a , specifies the scale at which the underlying $U(1)$ symmetry is broken. If $H_I/2\pi < f_a$, then this symmetry is broken during inflation, and the axion acquires *uncorrelated* isocurvature perturbations [333, 334, 335].¹⁰ The uncorrelated CDM isocurvature amplitude is bounded by *Planck* to be $A_a/A_s < 0.038$ at 95% CL [34]. While axion isocurvature probes the same inflationary energy scale probed by primordial B modes, searches for isocurvature are independent of other constraints on the tensor-to-scalar ratio, offering a distinct test of inflationary physics, if a significant fraction of the DM is composed of axions.

The axion isocurvature amplitude is:

$$A_I = \left(\frac{\Omega_a}{\Omega_d} \right)^2 \frac{(H_I/M_{\text{Pl}})^2}{\pi^2(\phi_i/M_{\text{Pl}})^2}, \quad (1.24)$$

where $\Omega_d = \Omega_a + \Omega_c$ is the combined total primordial relic density parameter in ULAs and standard CDM. The initial axion displacement, ϕ_i , fixes the axion relic abundance such that $\Omega_a = \Omega_a(\phi_i, m_a)$ [337, 338, 339, 340, 341, 342]. Thus, if the relic density and mass can be measured by independent means, such as direct detection, a measurement of the axion isocurvature amplitude can be used to measure the energy scale of inflation, H_I .

We forecast the errors on axion isocurvature for the baseline CMB-S4 experiment with a $1 \mu\text{K}$ -arcmin noise level and a $1'$ beam; the isocurvature limit will be improved by a factor of approximately 5 compared to *Planck*, allowing for detection of axion-type isocurvature at 2σ significance in the region $0.008 < A_I/A_s < 0.038$. If the QCD axion is all of the DM, axion direct-detection experiments (e.g., ADMX [343] already in operation, or CASPER [344], in progress) and CMB-S4 could probe H_I in the range

$$2.5 \times 10^6 \lesssim H_I/\text{GeV} \lesssim 10^7 \quad (\text{S4+ADMX}), \quad (1.25)$$

$$10^8 \lesssim H_I/\text{GeV} \lesssim 4 \times 10^9 \quad (\text{S4+CASPER}), \quad (1.26)$$

where we have used the standard formulae relating the QCD axion mass and relic abundance to the decay constant [334]. Combining axion DM direct detection with CMB-S4 isocurvature measurements allows a unique probe of low-scale inflation, inaccessible to searches for tensor modes.

¹⁰We ignore the case where $H_I/2\pi > f_a$, since then no isocurvature fluctuations are excited. The limit $r_{0.05} \lesssim 0.1$ [127] implies that isocurvature is produced if $f_a > 1.6 \times 10^{13}$ GeV. This accounts for the QCD axion in the “anthropic” window (roughly half of the allowed range of f_a on a logarithmic scale), axions with GUT-scale decay constants (such as string axions [336, 146]) and axions with lower f_a in models of low-scale inflation.

We now consider isocurvature fluctuations in ULAs, which have a number of distinctive features the CMB [345, 346, 302, 300, 301, 316]. We fix the fiducial ULA fraction to 1%, such that Ω_a and m_a can be separately measured using the CMB-S4 lensing power, and thus using Eq. (1.24), a measurement of A_I is a measurement of H_I . In contrast to the QCD axion, there are masses ($m_a \lesssim 10^{-26}$ eV) for which tensor modes impose a stronger constraint on H_I than isocurvature. There are also regions of overlap between possible tensor and isocurvature measurements, where CMB-S4 could be used to determine isocurvature and axion parameters, yielding an independent measurement of H_I :

$$2.5 \times 10^{13} \lesssim H_I/\text{GeV} \lesssim 10^{14} \text{ (ULAs, CMB-S4 alone) .} \quad (1.27)$$

A broader range of scenarios is considered in Ref. [3]. Already with *Planck* data, the combination of lensing reconstruction and isocurvature limits gives the posterior range for the tensor-to-scalar ratio of $r < 0.01$ for $m_a = 10^{-24}$ eV [300].

1.4 Mapping matter in the Cosmos

The third science theme relates to surveying the contents of the Universe revealed by the millimeter-wave sky.

The deep and wide CMB-S4 field will provide a unique census of a large fraction of the sky at centimeter to millimeter wavelengths. Matter in the Universe shows up in this census in multiple ways: the gravitational potential coherently shears the image of the CMB fluctuations, some CMB photons are scattered by electrons along the way, and many objects emit their own radiation. All of these help us to map out the matter in the cosmos.

The main value of CMB-S4 for non-CMB scientists will come from the deep and wide survey (although the ultra-deep survey will also be useful for these communities). In the reference design, CMB-S4 will map 70% of the sky in total intensity and linear polarization, with a cadence of around 2 days, reaching a 5σ point-source noise well below 1 mJy in the final coadded maps, along with corresponding maps of gravitational lensing and Compton- y .

The CMB-S4 survey will complement and enhance the LSST optical survey of the same region [347], the DESI spectroscopic survey [348], the *eROSITA* all-sky X-ray survey [349], and other planned and yet-to-be-imagined surveys from both ground- and space-based facilities. By going dramatically deeper than previous CMB surveys and covering a large fraction of the sky, the potential for new discoveries is high. This section provides an overview of some of the products to be derived from CMB-S4 that will be of significant utility to the broader astronomical community.

With these data, CMB-S4 will be able to probe physical processes that govern and regulate galaxy formation in various classes of objects, and probe the circumgalactic and intracluster mediums. Secondary CMB anisotropies are an excellent probe of reionization because scattering of photons by free electrons affects the observed temperature and polarization in a predictable fashion that is sensitive to its duration and morphology (or “patchiness”). CMB-S4 will leave a legacy that includes catalogs of extragalactic sources, maps of extragalactic integrated density and pressure, and multifrequency images of Galactic polarization.

These data sets will be exploited by astronomers and astrophysicists over the decades to come.

1.4.1 Extragalactic component maps

With the broad frequency coverage of CMB-S4, we can use internal-linear-combination techniques to produce maps of isolated sky signals. These products will include maps of the CMB temperature and polarization anisotropies, the reconstructed CMB gravitational-lensing potential, Compton- y maps, the cosmic infrared background, and Galactic synchrotron and dust (in both intensity and polarization). Of particular utility to the broader astronomical community will be the lensing map and the Compton- y (or tSZ) map; we detail their properties below.

Gravitational Lensing Map: One of the main goals of the CMB-S4 project is to extract cosmological information from a reconstructed map of the large-scale structures responsible for gravitationally lensing the CMB. This will constitute a map of all matter between us and the CMB, including dark matter, with a broad redshift weighting that peaks near redshift $z = 2$. The map will cover the entire footprint of the CMB-S4 survey and will represent the true mapping of matter, with map-level signal-to-noise ratio exceeding unity,

from the largest scales in the survey down to scales of approximately 12 arcminutes (representing multipole $L \simeq 1000$). Even on smaller scales, there will still be significant statistical information.

Since the contribution to CMB lensing is very broad as a function of redshift, including appreciable weight at $z < 1$ where many optical and other surveys are sensitive, this map will have significant potential for cross-correlation with data at a range of wavelengths. Thus far, CMB lensing maps have been studied in conjunction with a very diverse group of data sets, including (in order of decreasing wavelength): extragalactic radio catalogs [350, 351, 352]; maps of the Sunyaev-Zeldovich effect [353]; the cosmic infrared background [354, 355, 356, 357]; far-infrared galaxies [358]; mid-infrared galaxies [359, 360]; mid-infrared quasars [361]; optical redshift surveys [362], including filaments [363] and voids [364]; optically-selected galaxy clusters [360]; optically-selected quasars [365, 360]; optical weak lensing [366, 367]; optical photometric catalogs [368, 369]; the Lyman- α forest [370]; X-rays [371]; and γ -rays [372]. Additionally, CMB lensing has been used to measure the masses of SZ-selected [373, 374] and optically-selected [375] galaxy clusters and groups.

The lensing map from CMB-S4 will have much lower noise levels than those used in each of these earlier studies, and will additionally have very wide sky coverage. Because it is primarily based on CMB polarization data rather than temperature data, it should also be more robust against possible foreground contamination. Future analyses can thus be performed at much higher signal-to-noise ratio, with fewer concerns about contamination. Furthermore, many new data sets in many wavelength ranges will be available on the timescale of CMB-S4 that will both be deeper and have wider sky coverage than many of the studies carried out so far, including: LSST [347] for both photometric galaxy catalogs and weak lensing; DESI [376] for spectroscopic samples; and sensitive maps of ionized gas with the SZ effect from CMB-S4 itself. In addition to being used in delensing the CMB-S4 map of the primary CMB, the CMB-S4 lensing map will also provide the ability to delens the maps from future satellite CMB missions, which are unlikely to have the same internal sensitivity to lensing. The map of CMB lensing from CMB-S4 will thus represent a lasting, legacy value to the subsets of the extragalactic astronomical community working in nearly every wavelength range.

Compton- y map: Multi-frequency CMB temperature data enable the reconstruction of tSZ (Compton- y) maps by taking advantage of the known, unique tSZ spectral function. The first such maps have been produced recently using data from *Planck* [377, 353, 378, 379]. Here, we forecast the reconstruction of a Compton- y map using data from the CMB-S4 large-aperture telescopes (LATs) in combination with *Planck* (due to the large-scale atmospheric noise from the ground, *Planck* will remain useful at low ℓ). We use a harmonic-space internal-linear-combination (ILC) method to obtain post-component-separation noise curves for the tSZ signal. The sky model includes essentially all important contributions to the microwave sky (Galactic and extragalactic), and is identical to that used in Ref. [49]. The CMB-S4 LAT noise properties are described elsewhere in this document. We include *Planck* data from 30–353 GHz, assuming white noise and Gaussian beams. The ILC approach includes the option of “deprojecting” particular contaminants with a specified SED, at the cost of increased statistical noise. We specifically consider deprojecting the CMB, a fiducial cosmic infrared background (CIB) spectrum (modified blackbody with temperature of 19.6 K and spectral index of 1.2), or both.

The results of this analysis are shown in Fig. 20. The figure shows the post-component-separation noise (per mode) of the Compton- y map produced from CMB-S4 LAT and *Planck* data for various ILC-based foreground-cleaning options, in comparison to the tSZ power-spectrum signal. The CMB-S4 noise is low enough that the Compton- y field will be mapped on a mode-by-mode basis for multipoles $3000 \lesssim \ell \lesssim 5000$, even when explicitly deprojecting the CIB foreground. In fact, even with the expected level of atmospheric noise, CMB-S4 still improves enough over *Planck* to allow mapping at $\ell \approx$ few hundred. The figure also shows a similar noise curve for the baseline Simons Observatory (SO) configuration [49], demonstrating the significant gains that CMB-S4 will achieve. The total signal-to-noise ratio of the CMB-S4 tSZ power-spectrum measurement is 1570 (standard ILC), 570 (CMB deprojected), 1500 (CIB deprojected), or 130




Figure 20. Post-component-separation noise for the Compton- y map reconstructed from the CMB-S4 LATs and *Planck*. The solid black curve shows the tSZ power-spectrum signal. The solid blue curve shows the ILC reconstruction noise, while the other curves show the noise levels for various foreground-deprojection options. The thin, solid magenta curve shows the ILC reconstruction noise for the Simons Observatory baseline configuration.

(CMB and CIB deprojected); for this estimate, we neglect the trispectrum contribution to the covariance matrix, considering Gaussian errors only. We also conservatively assume $f_{\text{sky}} = 0.45$. For the standard ILC, this is an improvement over the current *Planck* measurement by nearly two orders of magnitude. Moreover, this estimate only captures the Gaussian contributions to the tSZ signal, which is a lower bound on the total information content in the y -map, given the significant non-Gaussianity of the tSZ field. The legacy value of this map will be immense, including cross-correlations with optical, infrared, X-ray, and intensity-mapping data sets.

1.4.2 Galaxy formation and evolution

The small fraction of CMB photons scattered during their cosmic journey create secondary anisotropies that encode the evolving spatial distribution and thermal energy of diffuse ionized gas throughout the Universe. Because the CMB is a nearly uniform back light on small scales, angular features in the scattering distortions and anisotropies generated can be used to probe *all* the ionized gas within a given region, not only that gas dense enough to emit in the X-rays or which happens to contain the particular ionic species that generate observable absorption features along the particular lines of sight to background sources in instrument-specific redshift windows. With its unprecedented sensitivity and angular resolution, CMB-S4

will deliver transformative high signal-to-noise ratio maps of ionized gas over a wide range of scales and redshifts, from ionized bubbles during reionization to the intracluster medium of the most massive present-day halos. These maps, alone and in combination with large-scale structure probes at other wavelengths, will allow us to finally address a crucial longstanding question: *how does the energy injected by galaxies and supermassive black holes affect the surrounding baryonic reservoir of diffuse gas driving their evolution?*

In this section, we describe in detail the new physical insights and model constraints that CMB-S4 will provide by addressing this question, and present new quantitative forecasts for constraints on state-of-the-art numerical models of galaxy formation and reionization. We emphasize that due to the rich nature of this science, the forecasts presented here represent only a lower bound on the impact of CMB-S4 in this area; the range of applications is vast, and will certainly expand further by the start of the CMB-S4 survey.

1.4.2.1 Intracluster medium and circumgalactic medium

Understanding the physical processes that govern and regulate galaxy formation over cosmic time is a central goal in astrophysics. It has long been known that star formation is an inefficient process: less than 10% of the cosmic abundance of baryons has been converted into stars over the past 13.8 Gyr (e.g., [380, 381]). Yet early hydrodynamic simulations of galaxy formation predicted much larger star-formation efficiencies; over the past two decades it has become clear that this discrepancy is likely resolved by *feedback* processes that prevent the over-cooling of gas in dark matter halos (e.g., [382]). However, the exact mechanism(s) and even the basic energetics of these feedback processes remain poorly constrained. On the scale of galaxy groups and clusters, the most important source of feedback—i.e., energy and/or momentum injected into the intragroup/intracluster medium (ICM)—is expected to be that due to active galactic nuclei (AGN), supermassive black holes whose accretion disks drive powerful winds and jets of radiation into their surroundings.¹¹ At lower mass scales (e.g., Milky-Way-sized halos), feedback due to supernovae and stellar winds may play the dominant role in injecting feedback energy into the circumgalactic medium (CGM). This energetic feedback profoundly alters the thermodynamic structure of the ICM and CGM, heating the gas (thereby preventing star formation) and driving it to large halo-centric radii, sometimes entirely ejecting it from the halo. Such processes are likely responsible for the “missing baryons” problem, i.e., the empirical fact that low-redshift galaxies (such as the Milky Way) have a baryon deficit relative to the cosmological abundance determined from the CMB and BBN [383]. Moreover, this re-distribution of the gas leads to non-negligible changes in the dark matter distribution on scales below ≈ 10 Mpc, thus altering the matter power spectrum. The amplitude of this effect is the largest source of theoretical systematic uncertainty in cosmological constraints derived from ongoing and upcoming weak-lensing surveys (e.g., [384]). Thus, constraining the properties of AGN and supernova feedback is essential not just for understanding galaxy formation, but also for enabling next-generation constraints on dark energy and neutrino physics. Beyond feedback, the role of non-thermal pressure support sourced by gas bulk motions and turbulence in stabilizing the ICM and intragroup medium must be precisely determined in order to calibrate biases in hydrostatic-equilibrium-based estimates of cluster masses (e.g., [385, 386, 387]); however, the amplitude of this effect is currently better understood in hydrodynamical simulations than that of AGN feedback (e.g., [388, 389, 390]), and thus we focus on the latter here.

The CMB offers multiple powerful tools with which to constrain feedback processes and non-thermal pressure support in galaxies, groups, and clusters. The tSZ and kSZ effects, which are sourced by the Thomson-scattering of CMB photons off free electrons, directly probe the electron thermal pressure and momentum, respectively. Given external knowledge of the peculiar velocity field (e.g., from galaxy redshift surveys), the kSZ signal is a measure of the electron density. Both the tSZ and kSZ signals are redshift-independent, a

¹¹At high redshift ($z \gtrsim 2$), other sources of feedback (e.g., from star formation) may be of similar importance in groups and proto-clusters.

unique property that is shared by few probes in observational astrophysics. The tSZ signal allows the thermal pressure of ionized gas to be measured over a wider range of halo masses, redshifts, and halo-centric radii than any other probe. The kSZ signal is essentially the only observational tool that can directly measure the ionized gas distribution with few assumptions over a wide range of halo masses, redshifts, and halo-centric radii. Moreover, the small-scale properties of both signals are strong probes of feedback. Qualitatively, AGN and supernova feedback tends to flatten the gas pressure and density profiles of halos, while decreasing the integrated pressure and density content within the virial radius of lower-mass objects (due to the ejection of gas from these shallower potential wells, in comparison to the deeper wells of larger halos). However, the details of these predictions, including their mass, redshift, and radial dependence, depend strongly on the exact feedback model and implementation. In most simulations, the relevant physics is “sub-grid”; thus, it is essential to directly measure these quantities in order to correctly calibrate the energetics of these processes.

Here, we demonstrate explicitly how CMB-S4 will achieve this goal. We focus on extracting the tSZ and kSZ signals of halos of various masses and redshifts via cross-correlations of the deep and wide CMB-S4 LAT survey with galaxy catalogs from the DESI, Baryon Oscillation Spectroscopic Survey (BOSS), and Sloan Digital Sky Survey (SDSS) spectroscopic surveys (LSST and other photometric surveys will also provide useful catalogs for such measurements) and cluster samples from the CMB-S4 tSZ-selected catalog (the *eROSITA* X-ray-selected catalog will also be useful in this context). We closely follow the methodology of Ref. [391] to forecast constraints on the stacked electron-density and pressure profiles of these halo samples. We consider four example sets of halos, with numbers drawn from DESI [348], BOSS/SDSS [392], or the CMB-S4 tSZ cluster catalog:

- low-redshift BOSS/SDSS luminous red galaxies (LRGs), $z = 0.2$, $M_{200c} = 10^{13} M_{\odot}$, $N = 2.5 \times 10^5$;
- high-redshift DESI LRGs, $z = 1$, $M_{200c} = 10^{13} M_{\odot}$, $N = 2.5 \times 10^5$;
- low-redshift CMB-S4 clusters, $z = 0.2$, $M_{200c} = 10^{14} M_{\odot}$, $N = 1.5 \times 10^3$;
- high-redshift CMB-S4 clusters, $z = 1$, $M_{200c} = 10^{14} M_{\odot}$, $N = 1.1 \times 10^3$.

In all cases, we choose narrow mass and redshift bins: $\Delta \log_{10}(M_{200c}) = 0.1$ and $\Delta z = 0.1$. The LRG samples from these optical surveys cover the entire window $0.2 < z < 1$, so stacked measurements similar to those shown here can be performed over this entire redshift range. Samples of other galaxy types (e.g., emission-line galaxies) will extend to even higher redshift, enabling similar analyses. The CMB-S4 cluster counts assumed here are drawn from Fig. 25, and will enable galaxy cluster studies from $0.1 < z < 2$. Halo masses inferred from CMB-S4 CMB lensing data (and LSST weak lensing at low- z) will be essential in this work, relating the observed baryonic profiles to the underlying matter-density field.

We apply the component-separated CMB temperature (for kSZ) and Compton- y (for tSZ) noise curves described in the forecasting methodology appendix to a stacked aperture photometry estimator to obtain error bars on the projected gas density (kSZ) and projected thermal gas pressure (tSZ) profiles of the halos. Note that we use the “velocity-reconstruction” kSZ estimator here, which uses spectroscopic redshift information to weight the CMB map in a manner that is highly robust to foreground contamination [393, 394]. Our kSZ forecasts assume that the noise in the electron-density profiles is dominated by the noise (instrumental and foreground) from the CMB data rather than noise in the velocity reconstruction; an approximate estimate indicates $\approx 10\%$ residual noise in the velocity reconstruction using DESI. Future simulation work will be necessary to precisely calibrate this effect.

The electron-pressure and density-profile error bars calculated using the methodology described above are contrasted with theoretical predictions extracted from six cosmological hydrodynamics simulations (for more details on the simulations, see Appendix A): BAHAMAS (fiducial, “high-AGN,” and “low-AGN” models)

[395], simulations of Refs. [396, 397], EAGLE [398], and IllustrisTNG-300 [399, 400, 401]. We rescale all electron-density and pressure profiles to the same background cosmological parameter values ($\Omega_b = 0.04898$, $\Omega_m = 0.3111$, $h = 0.6766$ [127]), so that only the astrophysical model differences are reflected in the forecasts. These simulations all differ in significant ways (even at the level of using grid- or particle-based codes), particularly in their implementations of sub-grid models for AGN and stellar feedback. As a fiducial model, we simply take the mean of the predictions from all six simulations for each observable. We also perform a smooth extrapolation of the simulation predictions to large radii (beyond $2r_{200c}$, within which the one-halo term dominates), which is necessary for the line-of-sight projection calculations. However, the extrapolation is not significant, since we assume that the virial shock leads to a steep decline in the profiles at $2.5r_{200c}$.

The results of this analysis are shown in Figs. 21 and 22. Each figure corresponds to a different choice of halo mass and redshift, with the left panel showing the cumulative electron-density profile and the right panel showing the cumulative electron thermal-pressure profile (i.e., the cumulative thermal energy in electrons). Here, “cumulative” means “integrated within an aperture,” where the angular size of the aperture in arcminutes is given on the top axis of each figure. The profiles asymptote to constant values as the full electron content or thermal energy content of the halo is enclosed. The electron-density profiles are shown in terms of τ , the Thomson scattering optical depth, while the electron pressure profiles are shown in terms of Compton- y , the line-of-sight integral of the electron pressure. The figures also show the forecast signal-to-noise ratio in each case, which range from ≈ 10 to greater than 100. All of the profile forecasts and error bars are convolved to an effective beam of FWHM = 1.4 arcmin (the CMB-S4 LAT resolution at 145 GHz). The figures also show error bars for the baseline configuration of the Simons Observatory (SO) [49], including the effect of the smaller number of clusters for the $M_{200c} = 10^{14} M_\odot$ halos (6–7 times fewer clusters for these cases).

It is evident from the figures that the galaxy-formation models considered here can be distinguished at high significance with CMB-S4. Particularly important in this regard is the high angular resolution of the CMB-S4 LAT, which permits access to the inner regions of the halos at low- z and is roughly matched to the virial scale at high- z . No other astrophysical observables are sensitive to these thermodynamic properties of the gas over such large ranges in mass, redshift, and halo-centric radius. X-ray observations are limited at large cluster-centric radii, at low halo masses, and at high redshifts. However, measurements with *eROSITA* will be strongly complementary to those from CMB-S4; the combination of the two facilities will enable gas pressure and density profiles to be measured from the core to the outskirts of galaxy groups and clusters. Absorption line studies (e.g., with the Cosmic Origins Spectrograph on the *Hubble Space Telescope*) are limited to comparatively small samples of halos at low mass and low redshifts; however, they are potentially complementary to CMB-S4 measurements on stacked samples of galaxies at lower masses than considered in Figs. 21 and 22, where galaxy-formation models yield even larger differences in their predictions than seen in the figures.

The profile-based constraints shown here are only a small fraction of the total information content in the CMB-S4 dataset relevant to galaxy formation. CMB-S4 will yield tight constraints on the integrated Compton- y and τ signals as a function of halo mass and redshift (“scaling relations”). The slopes of these relations are highly sensitive to feedback models [402, 403, 404]. An optimal analysis will simultaneously incorporate constraints from the spatial distribution (profiles) and mass- and redshift-dependences of these signals. As an illustration of the constraining power of CMB-S4 in this regard, Fig. 23 shows forecasted constraints on the integrated Compton- y signal as a function of halo mass, derived from the combination of CMB-S4 data with the DESI bright galaxy sample. The plot shows the fractional constraint on the deviation of the Y - M relation from a fiducial $M^{1.79}$ power-law model. The constraints span the range of halo masses from Milky-Way-sized galaxies to massive clusters, with sub-percent constraints near a pivot halo mass of

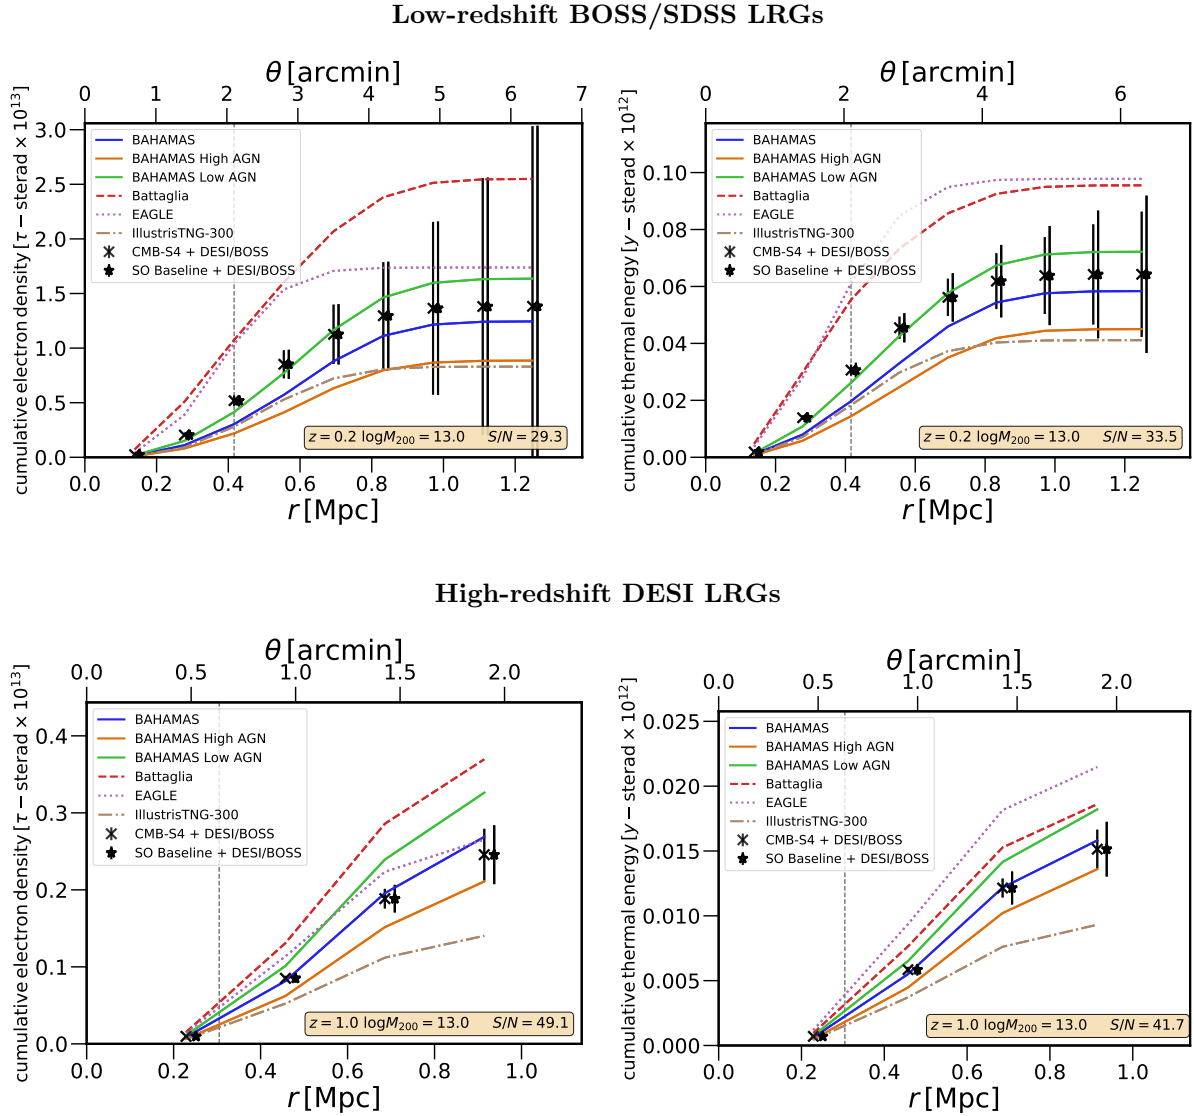


Figure 21. CMB-S4 constraints on the cumulative electron-density (left) and thermal-energy (right) profiles will distinguish between feedback models. Top row: Stacking $N = 2.5 \times 10^5$ BOSS and SDSS LRG halos of average mass $M_{200c} = 10^{13} M_{\odot}$ at $z = 0.2$. The left panel is extracted from the kSZ signal and the right panel from the tSZ signal. The lines come from density and pressure profiles around such halos measured in six cosmological hydrodynamics simulations: BAHAMAS [395] (fiducial blue, “high-AGN” orange, “low-AGN” green); Battaglia et al. [396, 397] (red); EAGLE [398] (magenta); and IllustrisTNG-300 [399, 400, 401] (brown). The data points average the predictions, and show error bars determined via stacked aperture photometry applied to component-separated maps from CMB-S4 LAT and Planck data (or SO and Planck data). The dashed vertical lines denotes r_{200c} . The insets give the CMB-S4 forecast signal-to-noise ratio. The error bars are highly correlated due to the photometry method, but the models can nonetheless be distinguished at high significance. Bottom row: The same, inferred by stacking $N = 2.5 \times 10^5$ DESI LRG halos of average mass $M_{200c} = 10^{13} M_{\odot}$ at $z = 1$. We emphasize that ionized gas properties in the low-mass, high-redshift regime shown in the bottom row cannot be easily measured with any other astrophysical probe.




Figure 22. Top: CMB-S4 constraints on the cumulative electron-density (left) and cumulative thermal-energy (right) profiles inferred by stacking $N = 1.5 \times 10^3$ CMB-S4 clusters of average mass $M_{200c} = 10^{14} M_{\odot}$ at $z = 0.2$. Bottom: The same inferred by stacking $N = 1.1 \times 10^3$ CMB-S4 clusters of average mass $M_{200c} = 10^{14} M_{\odot}$ at $z = 1.0$. The panels, curves, and data points with error bars are analogous to those shown in Fig. 21. As in the previous figure, the signals are detected at high significance, and the galaxy-formation models can be distinguished. X-ray constraints at small radii will be complementary to those shown here.

$10^{14} M_{\odot}$. These improve over current constraints by more than an order of magnitude; at Milky Way mass scales, no constraints currently exist (e.g., [404]).

Further information can be extracted by cross-correlating the CMB-S4 kSZ and tSZ data with quasars, different galaxy types (e.g., red or blue), weak-lensing maps, or filament catalogs. In the kSZ case, some of these cross-correlations will rely on a different estimator than that used here, in which the multifrequency CMB-S4 data are employed to extract the kSZ effect on small scales via component separation, filtering, and squaring of the map; forecasts indicate $S/N \approx$ few hundred for CMB-S4 in combination with LSST or *SPHEREx* using this method [405, 406]. In addition, the large-scale (two-halo) regime of the tSZ signal contains additional unique information about the evolution of the bias-weighted average thermal electron pressure of the Universe [407]. Finally, CMB-S4 will make the first detection of the polarized SZ effect [408, 409], a significant milestone in the eventual use of this signal as an astrophysical and cosmological probe. Overall, the forecasts presented here are only a first step toward exploring the rich astrophysics of galaxy formation accessible with CMB-S4.




Figure 23. Fractional constraints on the integrated pressure versus halo mass relation from a joint analysis of CMB-S4 and DESI data.

1.4.2.2 Patchy reionization

The patchiness of reionization leaves its imprint on the CMB through the kSZ effect, which refers to blackbody temperature fluctuations induced by a combination of coherent bulk flows on large scales and variations in the electron density on small scales. These “patchy kSZ” anisotropies have only recently been used to place constraints on the duration of reionization [410, 411, 412, 413]. While incremental gains in precision are expected from experiments that are underway or begin built (e.g., [414, 49]), CMB-S4 will provide definitive reionization constraints from the high- ℓ CMB, due to its much higher sensitivity. The patchy nature of reionization also generates polarization fluctuations (e.g., [415, 416]). Recent work shows this signal may be




Figure 24. *CMB-S4 constraints on the optical depth and duration of reionization in a joint analysis using the kSZ power spectrum and four-point function.*

detectable by CMB-S4 through the B -mode power at $\ell \approx 50$ –500 or by explicit reconstruction of the optical depth at the map level [417]. Additionally, it is possible to correlate these patchy polarization anisotropies induced by patchy reionization with other tracers of large scale structure such as the cosmic infrared background and CMB lensing [418, 419].

Here we highlight the power of CMB-S4 through the patchy kSZ effect with forecasts based on simulations for which the predicted shape and amplitude of the kSZ power spectrum respond naturally to the variation of unknown physical parameters over their possible values. The simulation techniques are described in Refs. [420] and [421]. The spectra are simulated at discrete “step sizes” as two parameters of the underlying model are varied away from their fiducial values: the ionization efficiency (or number of atoms ionized per atom in halos above the minimum mass), ζ ; and the mean free path of ionizing photons, λ_{mfp} . The reionization history for each of the simulations is used to determine the Thomson scattering optical depth, τ_{es} , and the duration of reionization, $\Delta z \equiv z_{75} - z_{25}$, the redshift interval over which the volume filling factor of ionized regions evolved from 25 to 75 per cent. We then calculated derivatives of the power spectrum with respect to these parameters, τ_{es} and Δz , by finite differencing. We impose a Gaussian prior on the late-time “homogeneous” kSZ contribution. This term is already known at the roughly 10% level, accounting for astrophysical and cosmological uncertainties [422, 423], and will be known better than this in the CMB-S4 era (e.g., by using the measurements described in the previous subsection). Additional uncertainty due to marginalization over non-kSZ foregrounds is incorporated with the “Deproj-0” (standard ILC) CMB-S4 LAT + *Planck* noise curves for the reference CMB-S4 configuration.

The results are shown in the right-hand panel of Fig. 24, where we have also included the kSZ four-point function, which provides additional information to separate early- and late-time kSZ contributions [424], by using the same excursion set simulated kSZ maps described above. The power spectrum is sensitive to the

duration of reionization, but this is highly degenerate with the optical depth. However, the prior knowledge of the optical depth from *Planck* large-scale EE power spectrum measurements, which is included in our analysis via a prior $\sigma(\tau) = 7 \times 10^{-3}$, breaks this degeneracy. Thus we find $\sigma(\Delta z) \approx 0.14$ from the kSZ power spectrum. The four-point function, when combined with the *Planck* EE constraint on the optical depth, performs worse on the duration than the power spectrum, as seen in the figure. However, when combining the power spectrum and four-point constraints, we find that the τ measurement from CMB-S4 alone *improves on the current Planck EE -based τ constraint by a factor of 2–3*, with $\sigma(\tau) \approx 2, 5 \times 10^{-3}$. Note that this estimate marginalizes over an unknown contribution from the low-redshift kSZ signal (albeit with a roughly 10% prior) and also assumes we have no knowledge of other foregrounds. It is thus possible that tighter constraints will result from improved measurements of low-redshift kSZ signals, as well as improved measurements of high- ℓ foregrounds.

Finally, we note that alternative scenarios, such as those in which rare quasars or very high-redshift sources play a significant role, are described by models with different parameters than we have assumed here. Thus, our forecasts by definition are uncertain and depend on the assumed model and even the fiducial parameters within that model. The high precision of the model parameter constraints we have forecasted in the standard UV-dominated scenario implies that we will be able to rule out or confirm, with high significance, other more exotic reionization scenarios such as those that include early X-ray binaries, population III sources, or rare quasars. This will especially be the case when CMB-S4 is combined with external data sets such as 21-cm and Lyman- α emitter surveys, together with independent CMB-based constraints on the optical depth from the large-scale EE power spectrum. Such discoveries are an inevitable byproduct of probing the high-precision, high-redshift frontier with CMB-S4.

1.4.3 Extragalactic Legacy Catalogs

1.4.3.1 Galaxy clusters

Clusters of galaxies, the largest gravitationally bound systems in the Universe, are powerful probes of both cosmology and astrophysics. The deep and wide millimeter-wave CMB-S4 sky survey will enable the identification of over 70,000 such systems at high significance ($> 5\sigma$) via measurements of the thermal Sunyaev-Zeldovich (tSZ) effect [425]. Cluster detection and characterization via the tSZ effect is highly complementary to techniques at other wavelengths (e.g., [426, 427]) because tSZ observables provide both low-scatter mass proxies and a detection method that is independent of cluster distance (since it is a spectral distortion of the CMB). As shown in Fig. 25, CMB-S4 will discover *an order of magnitude more high redshift systems* ($z > 1.5$) than ongoing and upcoming CMB experiments [428, 429, 49]. CMB-S4’s deep multi-band observations will enable separation of the tSZ signal from contaminating radio and infrared emission from cluster members (Sect. 1.4.1), which will be particularly important for low-mass, high-redshift clusters.

This large and well-characterized cluster sample, in conjunction with data from LSST, will significantly improve constraints on cosmological models (Sect. 1.3.3, [430]) and, when further combined with data from the next generation ESA X-ray mission *Athena* (as well as other proposed missions, including NASA’s *Lynx* and *Origins Space Telescope*), will also offer an unprecedented opportunity to understand the assembly and evolution of the massive galaxies that reside in these systems, as well as the effects of astrophysical feedback on the intracluster medium (Sect. 1.4.2).




Figure 25. Cumulative number of clusters above a fixed redshift for published (solid [431, 432, 433]), and upcoming (dashed [429, 49]) CMB cluster samples. CMB-S4 will discover an order of magnitude more of the highest-redshift ($z > 1.5$) clusters than previous surveys.

1.4.3.2 Dusty star-forming galaxies

The unique combination of resolution, depth, and area covered make CMB-S4 ideal for constructing catalogs of extragalactic mm-wave sources. A remarkable and largely unanticipated result (but see, e.g., [434]) from the previous generation of CMB surveys was the discovery of large numbers of strongly lensed dusty star-forming galaxies (DSFGs, [435]), the massive galaxies that make up the high-redshift ($z > 1$) component of the cosmic infrared background. Because these objects are selected by their dust emission, which is re-radiated UV emission from intense star formation, this population of objects is complementary to samples of galaxies found from their stellar light at optical and near-IR wavelengths (rest-frame UV at high redshift) with instruments such as *Hubble* and *JWST*. The DSFG sample from the first generation South Pole Telescope (SPT) survey includes the most distant and massive halo known within the epoch of reionization ($z=6.9$; [436, 437]), and a galaxy cluster forming just 1.5 Gyr after the big bang [438]. Most of these sources are gravitationally lensed by intervening galaxies; this magnification enables them to be studied in detail with ALMA [439, 440], providing fast redshift determinations and increased effective spatial resolution. These first massive galaxies mark the sites of the largest overdensities in the cosmic web [437] and trace the formation of the dusty, molecular interstellar medium from the metal-poor and chemically simple raw materials present in the first galaxies. With ALMA, the lensed systems are also useful as probes of the dark-matter distribution in the foreground lensing halos [e.g., 441]. In addition to ALMA, these sources will be excellent targets to followup with *JWST*, and the catalog will also have strong synergies with surveys such as those carried out with LSST, *WFIRST*, SKA, and *eROSITA*.

The sample of DSFGs from the CMB-S4 survey is expected to contain over 100,000 sources, of which around 10,000 should be strongly lensed (see Fig. 26). Based on source counts and redshift distributions measured with SPT-SZ data [442, 440], we expect thousands of these sources to lie at $z > 7$, more than 100 to be at $z > 8$, and around 1,000 to be associated with high-redshift protoclusters (Fig. 26). The beamsizes of

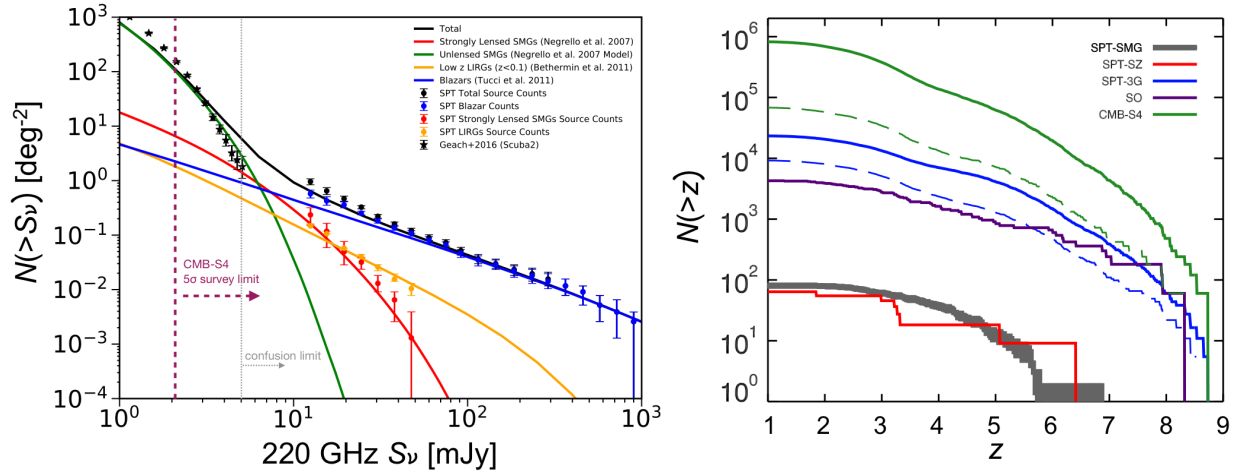


Figure 26. *Left:* Source density in the mm sky from SPT-SZ, SCUBA-2, and various galaxy-evolution models. With the expected 220-GHz detection threshold shown by the short-dashed line, CMB-S4 will detect over 100,000 extragalactic sources, including more than 10,000 strongly-lensed galaxies and thousands of galaxies at $z > 7$. *Right:* Expected number of DSFGs in the CMB-S4 survey. The gray line shows the measured distribution of DSFGs from the SPT-SZ survey [440], while the red line shown the predicted distribution from the phenomenological model of Ref. [442]. Also shown are the predicted numbers DSFGs that will be detected by the SO (purple), SPT-3G (blue), and CMB-S4 (green). The dashed lines show the number of sources detected above the confusion limit for a 6-m (green dashed line) and 10-m (blue dashed line) telescope.

CMB-S4 is well-matched to the expected scales of protocluster regions, building a catalogue far superior to the preliminary study of peaks in the CIB made possible using *Planck*'s all-sky survey [443]. The CMB-S4 high-redshift sources will open up a new window onto massive galaxy formation in the epoch of reionization, provide an important complement to the LSST and *WFIRST* surveys, and be prime targets for followup with ALMA and *JWST*. With the unique resource of the CMB-S4 dusty source catalog, we can directly observe the onset of dust production in the Universe, identify the first massive protoclusters of galaxies that represent the predecessors of the SZ cluster sample, and provide a well-understood target list for the high-redshift ALMA user community to explore galaxy formation.

An intriguing possibility is that CMB-S4 might be the first wide survey that can select both proto-clusters and genuine clusters of galaxies, using their different spectral signatures. Statistical samples of both kinds of source, and those intermediate between them, have never been available before.

Finally, the deep, high-resolution, multi-band data from CMB-S4 will be a rich source of “serendipitous” science, examples of which have already been demonstrated with the Atacama Cosmology Telescope (ACT), *Planck*, and SPT (e.g., strongly lensed DSFGs and mm transients).

1.4.3.3 Radio sources

The broad spectral coverage of CMB-S4 will also provide high signal-to-noise ratio intensity and polarization measurements of radio sources. Applying the model of Tucci et al. [444] with a confusion limit of 1 mJy at 90 GHz, over 100,000 AGN are expected to be detected at $> 5\sigma$ in the deep and wide survey field. An additional 5,000 sources will also be detected above 5σ in polarization [445]. This source sample, selected

from the highly-uniform CMB-S4 data set covering 70% of the sky, will be an order of magnitude larger than that from upcoming CMB surveys (e.g., Ref. [49]), and will be a valuable tool for statistical studies of AGN populations ([446, 447, 448, 449, 450]). The broad spectral coverage will enable measurements of source SEDs and—for the brightest polarized sources—Faraday rotation measures. These data will also be highly complementary to higher-resolution observations at other wavelengths (e.g., VLA/VLASS, ASKAP/EMU, SKA, and *eROSITA*). In the latter case, CMB-S4 data will provide an excellent opportunity both via individual and stacked source measurements, to characterize the radio/sub-millimeter properties of over 1 million AGN to be detected in the *eROSITA* X-ray survey [451].

While the majority of radio sources detected by CMB-S4 will be flat or falling-spectrum blazars, new and rare categories of sources may emerge from the large source sample. As an example of such discovery space, Ref. [452] recently reported a small intriguing sample of low- z AGN with rising high-frequency spectra that may contain cold intrinsic dust; such AGN—which will be more easily identified with the significantly higher resolution and deeper CMB-S4 data—could provide a view of the transition stage between the star-formation and AGN phases of such galaxies. CMB-S4 will similarly provide crucial higher-frequency data for source populations detected at other wavelengths. One interesting case will be the characterization of narrow-line Seyfert 1 (NLS1) galaxies for which there is limited high-frequency data (see e.g., Ref. [453]). These galaxies are assumed to be young AGN; however, based on 37-GHz observations, their detection rate is much higher than assumed by many models [453, 454]. Observations also indicate that—contrary to expectations—relativistic jets can be generated in such sources, regardless of their radio-loudness classification based on low-frequency data. A new catalog of southern NLS1 sources has recently been released [455] for which higher-frequency data are needed. CMB-S4 will be able to provide SED determinations, as well as information on the time variability of the sources (and many other well-detected galaxies, since CMB-S4 will be a powerful time-domain survey of the mm-wave sky). By combining observations of the NLS1s with data on other types of young AGN (such as compact steep-spectrum sources, GHz-peaked spectrum sources, high-frequency peakers, and compact symmetric objects), CMB-S4 will enable studies of the launching of relativistic jets, as well as the evolutionary paths that young AGN take on their way to becoming fully-evolved, powerful radio sources.

1.4.4 The Milky Way Galaxy

Galactic polarization: One area in Galaxy-scale science for which the deep polarization data from CMB-S4 will be particularly impactful is in the characterization of the interstellar medium (ISM), the polarization of which is a particularly pernicious foreground in the search for primordial B modes (Sect. 1.2.1). The ISM is known to be highly turbulent, compressible, and magnetized and while phenomenological models provide some insight into how foreground turbulence affects the dust polarization signal, they can not yet convincingly explain a number of current observations, such as the ratio of E modes to B modes. Excitingly, synthetic polarization maps from MHD turbulence simulations coupled with CMB observations are now beginning to shed light on the properties of the turbulent ISM in our galaxy.

The E -mode and B -mode power depend on the detailed properties of the turbulent density fluctuations in the ISM and it is not possible to characterize these fluctuations without knowledge of the sonic/Alfvénic Mach numbers in the ISM. To extract this information, CMB observations are coupled with other diffuse ISM tracers (e.g., 21 cm or $H\alpha$) to obtain velocity statistics. Current results are limited by *Planck*'s sensitivity and resolution, but with future CMB-S4 polarization data and velocity information we will finally be able to place tight constraints on the parameters surrounding cosmic-ray acceleration, star formation, and diffuse ISM structure formation.

Another outstanding mystery, in the area of the star formation, is whether the magnetic fields that thread molecular clouds are strong enough to inhibit movement of gas across magnetic field lines, thereby reducing the efficiency with which stars form. By studying the linearly-polarized radiation from dust grains, which tend to align with respect to their local magnetic field [456], CMB-S4 will create highly-detailed (< 1 -pc resolution) maps of the magnetic field morphology for nearly a thousand molecular clouds. Using these polarization maps we will measure the disorder in the magnetic field direction [457, 458, 459, 460], the degree of polarization [461], and determine whether there is alignment between the orientation of the magnetic field and cloud structure [462]. Applying the same analysis techniques to both the clouds observed with CMB-S4 and “synthetic observations” of numerical simulations, CMB-S4 can constrain the magnetization levels within molecular clouds. We can then compare the magnetization to the star-formation efficiency measured with *Herschel*, *Spitzer*, and the *Wide-field Infrared Survey Explorer (WISE)*. The large number of molecular clouds that can be mapped by CMB-S4 is crucial, since dust polarization is only sensitive to the magnetic field component parallel to the sky and so large numbers of molecular clouds are needed to correct for this degeneracy [463]. The results from CMB-S4 will place strong constraints on the dynamical importance of magnetic fields in different stages of the star-formation process.

Galactic intensity: Deep, multi-frequency maps over about half of the sky will be a useful resource for tracing the census of star-formation, including seldom-studied regions away from the Galactic plane. There is also the potential for adding information on particular classes of star (e.g., those with disks). Additionally, an exciting new possibility enabled by these CMB-S4 data is the search for variable Galactic sources, which we discuss in the next section.

1.5 The time-variable millimeter-wave sky

The fourth and final science theme relates to time-variable sources.

CMB-S4 will provide a unique platform to conduct a wide-field time-domain survey in the millimeter band, covering over half of the sky to few-mJy depths every two days. In this waveband, the time-variable sky is largely unexplored, with the exception of a shallow survey by *Planck*, surveys of the Galactic Plane (such as with the JCMT), targeted measurements of a few individual sources, and a single survey by SPTpol [13]; this is largely the result of limited observing time and fields of view for mm-band instruments (e.g., ALMA), which tend to focus on high-resolution observations of known objects. Despite this, a wide variety of sources are either known or believed to have particularly interesting time-variability in bands observed by CMB-S4. Expected sources include tidal disruption events, nearby supernovae, X-ray binaries, and classical novae. Particularly good candidates are γ -ray bursts and active galaxies, such as the time-variable blazar that was identified as a possible source of high energy neutrinos. The combination of high sensitivity and wide area for CMB-S4 will open a new window for time domain astronomy and multi-messenger astrophysics.

1.5.1 Gamma-ray bursts

Gamma-ray bursts (GRBs) are one of the primary time-domain science targets for CMB-S4. The spectrum of GRB afterglows is well-described by a self-absorbed synchrotron process, with a broad emission peak from approximately 100 GHz to 1 THz [464], with emission lasting on the order of one week. The existence of so-called orphan afterglows from bursts without detected prompt γ -ray emission—either because of the γ -ray instrument field of view, misalignment of the jet with Earth, or absorption of the primary γ -ray emission—is a generic prediction of GRB models, but none have ever been detected, despite a number of possible candidates (e.g., Refs. [465, 466]). The main obstacle to their detection has been that, at both short and long wavelengths, the sources are very dim and, at the frequencies where they are bright enough to be detectable (for example, in the millimeter band), either few or no blind surveys have been conducted.

CMB-S4’s observing strategy and sensitivity are expected to change this picture dramatically, delivering a factor of 2000 improvement on the only previous time-domain millimeter survey [13], which had a candidate detection, and gives an expected 1700 afterglow detections from a population model of on- and off-axis bursts (PSYCHE, [467]) over a 7-year CMB-S4 survey. Other theoretical predictions find that at all times during the survey there should be an ongoing detectable GRB afterglow [468]. Detection of these objects would serve a number of scientific goals:

- confirm measurements of the beaming angle of GRBs from jet breaks and thus the total energy budget for GRBs in the Universe;
- improve modeling of off-axis emission, and connect with gravitational-wave sources;
- constrain the existence of a large population of γ -dark GRB-like objects, which are potential sources of the TeV–PeV diffuse neutrino background observed by IceCube;
- potentially detect afterglows from GRBs made by population-III stars at high- z , during and prior to reionization.

Note that uncertainties in this last item are large, but nevertheless, CMB-S4 will enter the range of required sensitivities [469]. Detection of even one of these would provide a wealth of information about the early Universe, while a non-detection would constrain models of the first generation of star formation.

1.5.2 Fast transients

Fast radio bursts (FRBs) are a striking astrophysical phenomenon of unknown origin. They have been seen serendipitously in radio-frequency observations with Arecibo, the Canadian Hydrogen Intensity Mapping Experiment (CHIME), Green Bank, and Parkes radio telescopes (e.g., Refs. [470, 471, 472]), with around 20 detected to date. Given this event rate and estimates of sky coverage, the full-sky rate of FRBs is roughly 5000 per sky per day, if they are isotropically distributed. They are remarkably bright: the known examples all have a minimum flux density ranging from 0.3 to 2 Jy at frequencies of a few GHz, plus the original “Lorimer burst” with flux of 30 Jy [473]. They also have a very short duration; many last for less than 1 ms, with durations up to 5 ms. Most have no useful polarization information, although a few (such as FRB150807 [474]) have been observed to have linear polarization at various levels.

FRBs are consistent with having random sky locations (with the exception of two repeating sources [475, 476]). They also have dispersion measures ranging from 375 to 1600 pc cm⁻³. This leads to an arrival-time dependence on frequency proportional to ν^{-2} , which is very well measured. Electrons in the Milky Way give a total dispersion measure of around 100 pc cm⁻³, supporting the idea that FRBs are at cosmological distances of order 1 Gpc. Their frequency spectrum is unknown but limited evidence in the 1–2 GHz range suggests consistency with a flat spectrum in S_ν . If these intriguing sources do have a flat spectrum, some will be potentially detectable in microwave-background experiments, with flux densities greater than a few mJy at microwave frequencies. Current TES-based detectors have a sampling time of order 2 ms, meaning that FRBs would generally be visible only in a single time sample if their microwave and radio durations are comparable. A point-source FRB signal would appear to be a single bright microwave burst at a random position in the focal plane with the spatial profile of the experiment’s beam shape. If its spectrum is fairly flat, it would also appear in all channels of a multichroic detector simultaneously, and in both linear polarization channels. In the time stream, the burst itself will appear as a spike, with a subsequent decay according to the detector time-constants; this decay behavior distinguishes a sky signal from electronic glitches.

The detection of FRBs at microwave wavelengths would establish their frequency spectrum, while upper limits on their rate would significantly constrain the flux distribution. Either outcome would contribute substantially to understanding these mysterious extragalactic events.

1.5.3 Protostellar variability

Variability studies offer unique insights into the mass-accretion history of protostars. There is a fundamental disconnect between steady-state star-formation models and observed protostars—we know that at least some protostars go through explosive burst events during their evolution and that this may be important for the overall mass assembly (e.g., Ref. [477]). Circumstantial evidence suggests that most, if not all, protostars go through these stages; however, given the rarity of these large bursts, to catch one requires monitoring a huge number of protostars (thousands) over many years. CMB-S4 provides that opportunity.

A survey with the James Clerk Maxwell Telescope [478] has found that about 10% of protostars show low-level accretion variability, over timescales of a year or so [479]. This strongly suggests that we are probing

physical conditions in the inner, planet-forming, disk. In at least one case quasi-periodicity is seen over many years, suggesting a link to long-lived structure (i.e., companions or planets) within the disk. Additionally protostars are expected to be very active and a handful of extremely powerful flares have been observed in the radio [480]. Magnetic fields are believed to play an important role in funneling material from the disk onto the star (the last step for accretion). Detailed measurements of the properties of magnetic reconnection events are required to make significant progress in our understanding of this complex phenomenon. Follow-up spectroscopic studies of flares discovered with CMB-S4 could unveil the astrochemistry of the time-varying signatures of volatile species and UV and X-ray dissociation, as well as changes in the protoplanetary disk snow-line through evaporation of ices.

Studying protostellar variability in this waveband is a new research area. The examples we give illustrate that there are important physical processes that can be discerned through measurements of variability and transient phenomena within the Milky Way. At millimeter wavelengths, so far only one modest survey (the JCMT Transient Survey) has monitored star-forming regions, and has just begun to produce results. To draw statistically significant conclusions will require a much larger survey, just like CMB-S4 will provide.

1.5.4 Accreting binary systems

Within our Galaxy, accretion in binary systems can lead to strongly time-variable emission. Observations have shown that classical novae could be easily detectable in CMB-S4, with fluxes observed to be hundreds of mJy at mm-wavelengths [481, 482], while flaring events have been observed in X-ray binaries to be tens of mJy [483].

1.5.5 Solar System

CMB-S4 will survey about 70% of the sky approximately every two days. Moving objects, such as asteroids, dwarf planets, or as yet undiscovered bodies such as Planet Nine, can be detected in difference images using their long-wavelength thermal emission. The expected mm-wave Solar System can be seen in Fig. 27. For mm-wave imaging of particular objects, an instrument like ALMA would be far more sensitive; the particular benefit of CMB-S4 will be nearly daily measurements of large numbers of bright objects over several years, along with wide sky coverage for discovery of new or unexpected objects.

Using thermal emission rather than reflected light, mm-wave measurements are less sensitive to albedo effects, and the wavelengths are long enough that molecular absorption and emission is negligible. Furthermore, the fall-off in flux as a function of distance is less severe, with the main effect being the usual $1/d^2$ dependence; optical reflected light gets an additional factor of $1/d^2$ from the distance to the Sun. The temperature of Solar-System objects in equilibrium with solar irradiance scales only as $d^{-1/2}$, while objects with substantial internal heat or radioactive heating would decrease even more slowly.

Asteroids show substantial time variability due to rotation and viewing-angle effects. Using CMB-S4, more than 1000 known asteroids will be detected, with many of these having high enough signal-to-noise ratio to chart the flux as a function of rotational phase. Comparisons with measurements in the optical band, and also the infrared (since many of these objects have been measured as part of the *NEOWISE* project [484]), will yield new insights into the surface properties and geometry of these objects (e.g., Ref. [485]).

Dwarf planets, such as Pluto, will be easily detected, with measurements possible out to roughly 100 AU. Extending further, CMB-S4 will be able to detect any possible planets of Earth radius or larger at about




Figure 27. Millimeter-wave Solar System. Red circles show predicted asteroid flux densities at 150 GHz, crosses show estimated flux densities of known dwarf planets at 90, 150, and 220 GHz, and diagonal lines show the flux density of any possible Earth-sized planet in the outer Solar System in the CMB-S4 bands. Planet 9 is estimated to be several Earth masses at a distance of several hundred AU.

1000 AU. The recent excitement about a possible Planet Nine [486] (nominally $5\text{--}15 M_{\oplus}$ at several hundred AU) has made it clear that there is substantial discovery space for new planets in the outer Solar System. Assuming an Earth-like composition, $1 M_{\oplus}$ of rocky material should have an equilibrium temperature of $> 30\text{ K}$ just from radiogenic heat. An atmosphere would further enhance the mm-wave flux. The mm-wave-sky monitoring performed by CMB-S4 offers the potential to search for a wide variety of moving objects.

1.5.6 Multi-messenger astrophysics

A wide-area sensitive mm-wave survey can play an important role in multi-messenger astrophysics, in concert with both high-energy neutrino searches and gravitational wave astronomy.

The IceCube neutrino source TXS 0506+056 appeared to be associated with a blazar [487] that is mm-bright. With CMB-S4, this source would have had nearly daily flux measurements over many years, as well as many other similar sources that could be used to characterize the statistics of variability.

The first binary neutron star merger, GW170817, was not visible at mm-wavelengths [14], most likely due to the low density of the environment for that particular event. It is expected that at least some future events should be in denser environments that will enhance the mm-wave flux [15]. If there is no other

detectable emission, CMB-S4 would provide arcminute localization as part of regular survey operations. In addition, the mm-wave light curve can be compared with emission at other wavelengths to better understand gravitational-wave events.

We do not know what fraction of future GW sources will turn out to be optically-obscured, but visible in the millimeter—so it seems wise to have an instrument that is regularly scanning the sky in this waveband. Provided that the GW source is in the southern hemisphere, then it *will* be in the CMB-S4 deep and wide survey field, and observed every few days as a matter of normal survey operations. *This same leveraging of the CMB-S4 wide survey applies to all other future examples of transient source for which astronomers would like rapid follow-up observations.*

Science and Measurement Requirements

Introduction

The science portfolio for CMB-S4 is rich and broad, spanning areas of fundamental physics through astrophysics. Here we identify four of the most compelling and unique areas of expected scientific output from CMB-S4 to define the Level 1 (L1) Science Goals that will drive the measurement requirements and instrument design. We then present a high-level overview of the process by which we have determined these measurement requirements, under cost constraints, to achieve our L1 goals. The quantitative flowdown work is described in Appendix A. Here we summarize key results from that appendix and present our measurement requirements. The measurement requirements driven by our four L1 goals enable the great variety of additional science goals of the CMB-S4 science program that we presented in Chapter 1.

2.1 Primary science goals

There are two primary quantitative science targets, related to primordial gravitational waves and light relics.

Primordial Gravitational-Wave (PGW) Science Goal:

If $r = 0$, achieve a 95% confidence upper limit of $r \leq 0.001$. If $r \geq 0.003$, achieve a 5σ detection.

Motivation: All inflation models that naturally explain the observed deviation from scale invariance and that also have a characteristic scale equal to or larger than the the Planck mass predict $r \geq 0.001$. A well-motivated sub-class within this set of models predicts $r = 0.003$ to 0.004 . A characteristic scale near the Planck mass arises in many models whether they emerge from string theoretic considerations, effective field theory, or a minimal-new-physics approach (Higgs inflation), precisely because of the role gravity plays in the origin of the scale. An upper limit at $r = 0.001$ would point us toward more complicated solutions that introduce a non-Planckian scale. The observed departure from scale invariance is a potentially important clue that strongly motivates exploring gravitational wave amplitudes down to $r = 10^{-3}$.

Light-Relics (LR) Science Goal:

Achieve $\Delta N_{\text{eff}} < 0.06$ at 95% confidence.

Motivation: We have the opportunity with CMB-S4 to detect new light particles thermally produced in the early Universe. The contribution to N_{eff} depends on both the nature of the particle and the energy at which it was in equilibrium with Standard Model particles. A natural target is to search for new particles back to before the QCD phase transition.

With CMB-S4, any particle that was in thermal equilibrium at the beginning of the QCD phase transition can be ruled out at 95% confidence. While that sensitivity is not sufficient to detect a real scalar at an epoch earlier than the QCD phase transition, the sensitivity of CMB-S4 allows a further two order of magnitude improvement in energy sensitivity to either a Weyl fermion or vector particle.

We also have two Legacy Survey science goals that we have used to define measurement requirements.

Galaxy-Clusters (GC) Science Goal:

For galaxy cluster searches, achieve a lower mass limit that is below $10^{14} M_{\odot}$ at $z \geq 2$.

Motivation: Galaxy clusters in the local Universe appear to have formed the bulk of their stars at $z \approx 2$ -3. A catalog at these redshifts will provide new views on the astrophysics of galaxy clusters. This sensitivity will allow views of clusters similar to massive clusters that we see at $z \approx 0.5$, but at an earlier stage in their development when they were forming their stars.

Gamma-Ray Burst (GRB) Science Goal:

Measure many gamma-ray burst afterglow light curves.

Motivation: Gamma-ray burst afterglows contain a wealth of information about the central engine and the surrounding medium. The peak wavelength of the emission evolves with time, passing through mm-wavelengths a few days after the burst. Measurements made at this time will provide key information in particular about the density of the surrounding medium, uniformly for all bursts in the survey area.

There are numerous additional science goals, presented in the previous chapter. Those additional goals are all enabled by the CMB-S4 survey, but none are design drivers.

2.2 Flowdown to measurement requirements

Figure 28 presents an overview of how key measurement properties are derived via a flowdown process from our four key science drivers. The work developing the quantitative relationships summarized here is presented in greater detail in Appendix A.

2.2.1 Overview

The Primordial Gravitational-Waves and Light-Relics science drivers lead us to a two-tiered survey, with a small ultra-deep field contained within a much larger deep field. For the Light-Relics target, reduction of sample variance drives us toward observing as much sky as possible. Information about N_{eff} primarily comes, at CMB-S4 noise levels, from the TE spectrum. Sample variance in this power spectrum is determined by sky coverage and the amplitudes of the TT and EE spectra—amplitudes that are sufficiently large as to make sample variance the dominant statistical source of uncertainty on the relevant angular scales. For the Primordial Gravitational Wave science, which depends on measurement of the BB spectrum, the relevant sample variance is much smaller and can be reduced further by delensing. Galactic foregrounds are also




Figure 28. Each science theme has one science topic contributing to the definition of measurement requirements and instrument requirements. The choice of a combination of low-res maps and high-res maps is driven entirely by PGW. The choice of a two-tiered survey is driven by the combined requirements of the PGW and Light Relics drivers. The connections of all the drivers to the measurement properties are indicated with color coding. Some measurement properties are written out in more than one color to indicate that they are influenced by more than one science driver.

much brighter relative to the signal of interest, which motivates finding the cleanest regions of sky. While the optimal strategy depends in detail on how well we can delens, and on how well we can clean the foregrounds, current forecasting points to the need to go very deep on a small ($f_{\text{sky}} \approx 0.03$) patch of sky.

Therefore the full CMB-S4 data set consists of an Ultra-Deep Field with many mm-wave bands over a fairly small patch of sky, and a Deep and Wide Field (shallower, though still unprecedentedly deep), covering a large fraction of the entire sky ($f_{\text{sky}} \approx 0.7$).

The Ultra-Deep field must be covered by both low-resolution and high-resolution instruments, driven by the unique demands of isolating degree-scale CMB B modes, and of estimating the contribution to these B modes from the gravitational lensing of E modes. The rationale for this split is given in more detail in subsequent chapters, as are instrumental considerations that drive the bifurcation into small- and large-aperture telescopes (SATs and LATs).

With the basic scheme in place of a survey with deep and ultra-deep fields, with the ultra-deep field having both high- and low-resolution components, we can now list the important measurement properties for each of these three components of the CMB-S4 survey: map area; total statistical weight; distribution of weight across frequency; angular resolution; and cadence. In the subsections that follow, we describe our process of flowing down science requirements to requirements for these measurement properties, for each of the four science drivers in turn.

2.2.2 Measurement properties required to achieve PGW Science Goal

Meeting our Primordial Gravitational-Wave goals will require significant advances, not only in raw sensitivity but also in control of foregrounds and instrumental systematics. As discussed in detail in the CMB-S4 Science Book and in previous publications [1], attaining the desired levels of sensitivity to the signatures of gravitational waves in the CMB necessitates first and foremost at least an order-of-magnitude increase in the raw number of detectors on the sky compared to Stage-3 experiments. These works also noted that foreground mitigation will be crucial for CMB-S4, particularly in the pursuit of the gravitational-wave signal. To minimize contamination from Galactic foregrounds, it is clear that multiple frequency channels are required. It is known from analysis of BICEP/Keck and *Planck* data [488] that, if unsubtracted and unmodeled, Galactic dust imparts a bias to the measurement of the tensor-to-scalar ratio r at 150 GHz at a level two orders of magnitude above the target $\sigma(r)$ for CMB-S4. Synchrotron emission is expected to impart a similar bias at lower frequencies. At least one frequency channel is required to remove each of these contaminants, and multiple channels will be needed for each component if the behavior of these foregrounds is not perfectly uniform across the survey. Galactic foregrounds are also reduced by observing a patch or patches of sky with as low as possible column density of Galactic material. This consideration—and optimizations of raw sensitivity—drive a survey for the degree-scale gravitational-wave signal to small patches of sky ($f_{\text{sky}} \approx 0.03$).

The other major foreground for gravitational-wave searches, one which cannot be mitigated with frequency coverage, is the signal from gravitational lensing. Scalar density perturbations in the early Universe produce only even-parity (“ E -mode”) patterns in the CMB polarization to first order, while gravitational-wave perturbations also produce odd-parity (“ B -mode”) patterns [489, 490, 491]. This fact is essential to our ability to detect a gravitational-wave signal in the presence of the much larger signal from scalar density perturbations. Gravitational lensing of the E modes, however, produces a secondary source of B modes, at an amplitude significantly larger than the target CMB-S4 sensitivity to the gravitational-wave signal at degree scales. The lensing process can effectively be inverted with sufficiently high-fidelity measurements of the E -mode signal, to estimate the gravitational lensing potential, and a map-level prediction of the lensing contamination can then be accounted for in gravitational-wave searches [492, 493]. The level of this so-called “delensing” required for a small-area survey and the CMB-S4 target $\sigma(r)$ requires deep, high-resolution data over the same patch of sky observed for the degree-scale B modes.

The above measurement challenges were taken into account in the optimization exercises described in the CMB-S4 Science Book and the CDT report. We have extended that work for this report, and it is presented in a self-contained manner (rather than as an update) in this chapter with details in Appendix A.

If the true value of r is $\gtrsim 0.01$ then it may be the case that we can improve our characterization of the primordial B -mode signal by covering more than 3% of the sky. Thus we leave open the possibility of broadening our Ultra-Deep Field in order to re-optimize if Stage III data, or early data from CMB-S4, indicate a large value of r .

Frequency	20	30	40	85	95	145	155	220	270
Angular resolution (arcmin)	11.0	72.8	72.8	25.5	25.5	22.7	22.7	13.0	13.0
Total survey weight / 10^6 (μK^{-2})	0.12	0.69	0.43	11.0	14.1	5.7	4.8	0.71	0.24
Q/U rms (μK -arcmin)	8.4	3.5	4.5	0.88	0.78	1.2	1.3	3.5	6.0

Table 2-1. *Ultra-Deep degree-scale map noise and angular resolution measurement requirements. Survey weight is $2A/N^2$ where A is the effective sky area (in particular, $f_{\text{sky}}^{\text{noise}}$ as defined in Appendix A) and N is the Q/U map noise level. This quantity is linear in detector count and run time. Note that, as we will see in the next chapter, the 20-GHz channel is on an LAT rather than an SAT.*

As indicated in Fig. 28 the PGW science goals are the primary drivers for many of the CMB-S4 measurement properties. We now review the measurement properties of the Ultra-Deep surveys and how the requirements on them are determined by the PGW science requirements.

Map area: Assuming $r = 0$, our semi-analytic framework formally leads to an optimal map area that is less than 1% of the total sky. There are several effects, not included in the analytic optimization procedure that argue against targeting less than 1% of the sky, as discussed in Appendix A. These include the cut-sky impact on E/B decomposition and risks in becoming too heavily dependent on delensing. Folding these considerations into our semi-analytic results, we have chosen 3% as a nominal target field size for many of our optimization studies.

Broader sky coverage has advantages if r turns out to be sufficiently large. Thus we considered both a ‘‘Pole Deep’’ survey of about 3% of the sky, and a ‘‘Pole Wide’’ survey of about 7% of the sky. To cover more sky may require observations from Chile. Based on the studies in Appendix A, we may deploy some SATs to Chile if $r \gtrsim 0.01$. In this happy case, we will have the opportunity to make high-significance detections from two different sites using different observing strategies. Our reference design takes advantage of our access to these two different sites by building in deployment flexibility. The baseline plan is for all SATs to go to Pole, but if we find indications that r is sufficiently large, then we can also send some SATs to Chile.

Degree-scale total map weight and frequency distribution: Our semi-analytic framework can be used to calculate $\sigma(r)$ from observations of a given amount of sky, as a function of total detector effort (or, equivalently, total map weight), optimally distributing the detectors across the low-resolution frequency bands and a high-resolution band for delensing. The result of this calculation for $f_{\text{sky}} = 0.03$ is that achievement of $\sigma(r) = 5 \times 10^{-4}$ requires 1.8×10^6 150-GHz-equivalent detector years. Roughly seventy percent of this (1.2 million) is allocated to the degree-scale observations (with the remainder allocated for delensing) with a particular distribution across frequency as described in Appendix A.

In reality one is not able to choose the number of detectors in each frequency band in the continuously variable manner assumed in the semi-analytic calculations. For the reference design a realistic mapping of detectors into dichroic optics tubes has been carried out while seeking to maintain the band distribution as determined in the optimization calculations. This departure from the optimal distribution results in an increased requirement for the total number of 150-GHz-equivalent detector years for the degree-scale survey. The result is the reference design described in Sect. 3.3, with expected measurement performance as given in Table 2-1, with a total effort of 1.5×10^6 150-GHz-equivalent detector years delivering a total map weight, from all frequency channels combined, of $38 \times 10^6 \mu\text{K}^{-2}$.

Degree-scale angular resolution: The telescope aperture determines the angular resolution as a function of frequency, with lower frequencies having a lower angular resolution. The aperture size is set to allow 30 GHz and higher frequencies sensitivity to the primordial B -mode recombination bump at $\ell \approx 100$. Lower resolution than achieved by the reference design aperture at 30 and 40 GHz would result in reduced

Frequency (GHz)	20	30	40	95	145	225	270
Angular resolution (arcmin)	11.0	7.3	5.5	2.3	1.5	1.0	0.8
White noise level ($\mu\text{K-arcmin}$)	8.4	5.0	4.5	0.68	0.96	5.7	9.8

Table 2-2. *Measurement properties of a dedicated delensing survey that meets our measurement requirements.*

performance. As detailed in Appendix A, a 20-GHz channel on the SAT had insufficient sensitivity at $\ell \gtrsim 100$ due to insufficient resolution. Therefore the 20-GHz channel was moved to the delensing LAT.

We find the map noise levels at the given frequencies and angular resolutions in Table 2-1 are sufficient for achieving our science requirements.

Hi-res ultra-deep total map weight and angular resolution: The delensing total map weight is informed by our semi-analytic optimization. For observing 3% of the sky, one can see from Fig. 69 that achieving our target $\sigma(r)$ requires delensing at a level such that the residual lensing B modes have an rms amplitude about 30% of the uncleaned amplitude, corresponding to a residual lensing B -mode power that is about 10% of the uncleaned level. From Fig. 68 one can see that this level of delensing requires a map noise of about $0.7 \mu\text{K-arcmin}$ in the delensing survey.

The above is all for an idealized case of a delensing observation with a $1'$ beam and no foreground contamination. In practice we need greater total weight to clean or constrain foregrounds, as we discuss below.

In Fig. 68 we show forecasts for the amount of residual lensing B -mode power as a function of delensing map noise and angular resolution, again for coverage of 3% of the sky. We see that a requirement to clean out approximately 90% of the lensing-induced B -mode power is not a strong driver for angular resolution, and it is a strong driver for low map noise. Angular resolution for the delensing LAT reference design is driven primarily by our galaxy cluster science goal.

Hi-res ultra-deep frequency distribution: We have not performed a rigorous optimization of the allocation of detectors across frequency for the delensing LAT focal plane, as we have done for the SATs. For estimation of the lensing potential, and then of the lensing B modes, we expect foreground contamination to be a much lower fraction of the signal of interest than is the case for the degree-scale observations.

Our reference design has 95- and 145-GHz channels as the main CMB channels, and then a pair of lower frequencies to guard against synchrotron and free-free, and a pair of higher frequencies to guard against dust. Forecasts for lensing performance in the presence of foregrounds are made with the ILC method described in Appendix A.

We see the reference design allocation of detectors across frequency as a highly conservative one regarding control of foreground contamination. Allocation of detectors across frequency for the delensing LAT is primarily driven by our galaxy cluster science goal.

2.2.3 Measurement properties required to achieve LR Science Goal

To reach the light relics goal, a highly precise measurement of the CMB power spectrum is required, in particular for the cross-power between temperature and polarization (TE). The precision of this

Frequency (GHz)	30	40	95	145	220	270
Angular resolution (arcmin)	7.4	5.1	2.2	1.4	1.0	0.9
Total survey weight $(TT)/10^6$ [μK^2]	0.22	0.68	26.3	26.3	2.2	0.38
White noise level for TT (μK -arcmin)	21.8	12.4	2.0	2.0	6.9	16.7
White noise level E/B (μK -arcmin)	30.8	17.6	2.9	2.8	9.8	23.6

Table 2-3. Deep and Wide Field map noise and angular resolution requirements.

measurement is primarily set by the uncertainty in the maps and the number of independent samples (i.e., the area of sky for the survey).

The uncertainty in the polarization maps is expected to be dominated by instrument noise rather than by foreground uncertainties. This is because the dominant small-scale foregrounds for CMB-S4 (thermal SZ and dusty star-forming galaxies) are largely unpolarized; radio sources can be polarized, but at the sensitivity of CMB-S4 the bright sources can be individually identified and removed (or masked), with the remaining number of undetected sources expected to be too small to contribute any contamination.

Map area: The measurement property to which the standard deviation on N_{eff} is most sensitive is map area. Semi-analytic estimates of $\sigma(N_{\text{eff}})$ as a function of sky coverage, at fixed total map statistical weight, show rapid improvement as f_{sky} increases, as shown in Appendix A in Fig. 76.

Thus motivated, we explored observing strategies with successively smaller minimum observing elevations, which allow progressively greater sky coverage. There is a frequency-dependent noise penalty for observing at lower elevations, through higher air mass, due to the greater background loading of the detectors. Our observing strategy simulations took these noise penalties into account, allowing for the construction of multi-frequency maps of noise variance. These in turn were reduced to estimates of errors on CMB power spectrum measurements, and from there to estimates of $\sigma(N_{\text{eff}})$.

Exploration of observing strategies led to the discovery that avoiding the galaxy was more harmful than helpful, due to the impact of the time it takes to halt a scan and reverse its direction. The result for these galaxy-avoiding maps was they were effectively anti-apodized, with coverage building up near the edges. We have thus opted to scan right through the galaxy. In addition to opening up opportunity for a wide range of Galactic science, this scan strategy also allows us to make a choice of Galactic cut, for the N_{eff} science, after we have gathered the data.

The result of our analysis is Fig. 77 showing $\sigma(N_{\text{eff}})$ for three different choices of minimum elevation angle and as a function of the Galactic cut. The corresponding amount of sky used for light relic science is shown in Table A-19. For a 30° minimum elevation angle we can achieve our goal with a 13 to 14% galaxy cut. This means we cut out from our map the regions with Galactic contamination that are at the same level as the worst 13 to 14% of the whole sky. We believe that such a cut is sufficient for protecting us from significant foreground-modeling-induced biases in estimates of N_{eff} . We plan to investigate further with map-based simulations. We have not considered even lower elevation angles due to the rapidly increasing air mass penalties and concerns about sidelobe pickup from the ground.

For the above reasons, our reference design survey assumes the sky coverage that comes from a 30° minimum elevation angle. The coverage is inhomogeneous and thus leads to different measures of the amount of sky covered, as described in Appendix A. For what we call $f_{\text{sky}}^{\text{noise}}$ we have a total coverage of 71%, or 29,000 deg^2 . After the 13 to 14% Galactic cut this leaves about 62% for determination of N_{eff} . Other science goals, with greater sensitivity to Galactic foreground contamination, are likely to require larger Galactic masks.

Frequency distribution, total map weight, and angular resolution: The light relics science goal does not require a heavy amount of foreground cleaning. The major driver of frequency distribution is galaxy clusters. Given the distribution of detectors across frequency, and the exercise described above to determine the map area, we conclude we need a total map weight of $56 \times 10^6 \mu\text{K}^2$. Some of this statistical weight is spent on regions of the sky where Galactic emission is sufficiently bright that it is useless for light relics science, as discussed above.

Figure 75 studies the impact on changes to noise level and beam size on $\sigma(N_{\text{eff}})$. Of course, lower noise is better, as is decreased beam size. Given that the reference design hits our $\sigma(N_{\text{eff}})$ target, rather than surpasses it, we can not tolerate single changes such as an increase in noise or an increase in beam size. However, a decrease in beam size would allow for an increase in noise, or vice versa. We have not yet attempted a joint optimization of cost across all of these degrees of freedom, constrained to fixed $\sigma(N_{\text{eff}})$.

2.2.4 Measurement properties required to achieve GC Science Goal

Galaxy clusters in the local Universe appear to have formed the bulk of their stars at $z \approx 2-3$. Building a nearly mass-selected galaxy cluster catalog that extends to this period is a high priority for CMB-S4.

Measuring high-redshift galaxy clusters requires high sensitivity at several frequencies over a wide area. The clusters are detected using the distinctive frequency dependence of the thermal SZ effect. Removal of the cosmic microwave background and cosmic infrared background fluctuations is done using a combination of spatial and spectral filtering, requiring at least three frequencies at roughly comparable depth for measuring the separate components.

The CMB-S4 galaxy cluster program builds on Stage-3 experiments: SO will survey a comparable sky area with lower sensitivity, while SPT-3G is surveying about 20 times less area but with projected sensitivity that is similar to CMB-S4.

Assembling a galaxy cluster catalog at $z \approx 2-3$ requires enough sky area to have a sufficiently large cosmological volume, and enough sensitivity to detect clusters with low enough masses that the total cluster number density is sufficiently high. Galaxy clusters are rare, with the number density of the most massive clusters at any redshift exponentially suppressed, so the sensitivity is particularly important.

The sensitivity to galaxy clusters is a combination of raw detector noise, beam size, and frequency coverage. Everything else being held equal, lower noise straightforwardly improves sensitivity. At fixed detector noise, sensitivity will increase as the beam size is reduced as long as the cluster is not well-resolved; clusters at $z \approx 2-3$ will have their virialized regions subtending roughly $2'$. Multiple frequencies improve sensitivity by allowing the removal of the cosmic microwave background primary fluctuations and the cosmic infrared background fluctuations, both of which are sufficiently bright that in any single frequency map they will swamp the signal from $z \approx 2-3$ clusters.

2.2.5 Measurement properties required to achieve GRB Science Goal

CMB-S4 will be a sensitive mm-wave survey of the time-variable sky. While there are many expected sources, a particularly interesting target will be gamma-ray burst (GRB) afterglows. Many GRBs have been observed to be mm-luminous, and there is the possibility of detecting off-axis GRBs where the prompt gamma-ray emission is either too faint or relativistically beamed away from our line of sight.

The measurement requirements are that CMB-S4 have good point-source sensitivity with a cadence such that locations are re-visited several times per week.

Lower noise flux levels are helpful, allowing a larger volume to be probed. The minimum flux sensitivity is set by the range of observed fluxes for pointed follow-up of GRBs, which typically range from 1–10 mJy at frequencies of 90–230 GHz [494].

The time evolution of GRBs is such that there may be substantial evolution within the first day for some sources [495], but the typical light curve evolves over a timescale of roughly 1 week. This sets a requirement on cadence that a wide sky area is surveyed multiple times per week. More sky area simply increases the probability of catching a GRB within the field.

Theoretical estimates [468] are that there is at least one observable mm-wave GRB afterglow on the sky at any given moment at a flux density around 1 mJy at 150 GHz. The counts are expected to be in the Euclidean regime, where the number of sources above a flux limit $N(> S)$ scales as $S^{-3/2}$. At a sensitivity higher than 10 mJy the rate would be such that there would be at best a handful of events detectable per year.

Instrument Overview

3.1 Requirements and design drivers

The measurements outlined in the previous chapter require about an order of magnitude improvement in measurement accuracy compared with Stage-3 experiments. Such a large step in performance brings significant challenges, the most important of which is sensitivity. Detectors for ground-based CMB experiments are background-limited, so CMB-S4 must have an enormous number of detectors; roughly 500k detectors will be needed on telescopes with angular resolution ranging from tens of arcminutes for measuring r , to around 1.5 arcmin for measuring N_{eff} .

High raw sensitivity is only useful if systematic errors can be controlled. Beyond astrophysical foregrounds, which will be managed with broad frequency coverage, the major sources of systematic errors are atmospheric brightness fluctuations, pickup of unwanted signals from the ground, and polarization errors due to gain and beam shape differences between the two detectors in a polarimeter. Systematic errors are the principal driver for observing from sites with a stable atmosphere, and for using small telescopes, which are easier to shield, for measuring large angular scales. Only small telescopes have demonstrated control of systematic errors at a level that is needed to constrain r , so the reference design uses small telescopes for the r survey.

Many systematic errors can be measured and corrected, so their impact is primarily through unmodeled residuals. Simulations for the CDT report demonstrated that residual, noise-like, additive, systematic errors, both correlated and uncorrelated across different frequencies, with various angular power spectra, should be $\lesssim 10\%$ of the noise to ensure $\lesssim 10^{-4}$ bias in r . Residual systematic errors at this percentage level are consistent with what has been achieved in Stage-3 experiments [5]. A similar percentage error is expected at CMB-S4 sensitivity levels, because the various techniques that are used to model or filter the low-order effects of systematic errors are limited by the noise. In this case, CMB-S4 will not require vastly smaller absolute systematic errors than Stage-3 experiments, as long as the errors are stable. Much of the detailed design effort for CMB-S4 will focus on demonstrating that systematic errors can be controlled at an acceptable level, but our working assumption for the reference design is that Stage-3 technology is sufficient; we just need more of it.

3.2 Instrument configuration

CMB-S4 is a single experiment with a mix of large and small telescopes at the South Pole and in the Chilean Andes. Both sites offer excellent observing conditions at millimeter wavelengths, and both have well-established infrastructure that has supported many CMB experiments [496, 497, 498, 499]. Chile is higher, so the dry component of the atmosphere is more transparent, but the South Pole has consistently low precipitable water vapor and a more stable atmosphere, with substantially lower $1/f$ noise due to blobs of water vapor drifting by (e.g., see Ref. [500]).

A measurement of N_{eff} requires arcminute resolution over about 70% of the sky, but only about 10% of the sky is accessible from the South Pole, while about 80% is accessible from Chile, so the two large telescopes for

the N_{eff} survey must be in Chile. There is no strong driver for opening up a third site to increase the survey area, but we would likely take advantage of such a site if it were developed outside the scope of CMB-S4.

A measurement of primordial gravitational waves requires a survey using small telescopes for the primordial signal and one large telescope for delensing. The optimal siting of these telescopes will depend on the value of r , the nature of polarized galactic foreground emission, and the actual performance of the instruments in the field. For very low values of r , the smallest, deepest, sky patch is required, and the ability to observe a single patch at all times from the South Pole favors siting all the SATs there. For larger values of r we can take advantage both of the larger sky area available from the Atacama to observe multiple patches, and of the differences between the two sites to help control systematics, so a split distribution would be preferable. For the reference design we baseline siting all the SATs at the South Pole to ensure that we meet our L1 science requirements for any value of r , but we cost and schedule siting the SATs both at the South Pole only and equally split between the South Pole and Atacama to ensure that a wide range of options remain viable. The delensing LAT is sited at the South Pole in either case. We know that additional information on all of the issues informing this decision will become available during the CMB-S4 construction project from both existing experiments and from the planned South Pole and Simons Observatories (with the latter pair being particularly informative about comparable instrument performance), and we will continue to update the siting plan based on these new data. Finally we include the option within the baseline plan of moving SATs from the South Pole to the Atacama during operations should results from early operations indicate that doing so would be beneficial—for example, if an early detection were made of a moderate value of r and the priority became to control foreground and/or site-specific systematics and to reduce sample variance.

The small telescopes are based on designs that have been used for many years by the BICEP/Keck experiments, but with dilution refrigerators, for continuous operation, and three telescopes per cryostat to reduce the number of telescope mounts. The large telescopes are essentially copies of what is being implemented for Simons Observatory, but with more sub-field cameras to completely fill the telescope field of view with detectors. All of the telescopes have horn-coupled transition-edge-sensor detectors, with time-domain-multiplexed readout, because these technologies are scalable and have a strong heritage in Stage-3 experiments. Small residual systematic errors require well-controlled beams, so horns are an obvious choice. Windows, cold optics, and thermal filters are copies of Stage-3 designs.

The raw data rate for CMB-S4 will be approximately 20 TB/day, about the same as LSST, so transmission, storage, and analysis will be challenging, but manageable. Real-time monitoring, flagging, and calibration of time-domain data will all have to be more automated than in existing experiments, and reduction to well-characterized maps will have to be efficient, so that systematic errors can be identified quickly and corrected.

The flow down of measurement requirements to instrument requirements is still at an early stage, especially for requirements associated with instrument systematic errors, e.g., ground pickup and polarization errors, where it is the residual error, after modeling and filtering, that matters. Quantitative estimates of residual systematics require fairly detailed simulations; in some cases, errors must be applied in the detector timestreams and propagated all the way through to science analysis. Detailed simulations are needed to develop requirements for sidelobe level, which drives the design of baffles and shields for the telescopes, and for instrumental and cross polarization errors, which mainly affect the design of the horn to detector coupling scheme. We do not have to freeze all the instrument requirements before detailed design begins, because we already have some understanding of how to allocate systematic errors to subsystems from Stage-3 experiments, but developing simulation tools that accurately capture the effects of instrument errors is a high priority for CMB-S4.

3.3 Reference design summary

The reference design for CMB-S4 is a point design that establishes the feasibility of the experiment and provides a basis for preliminary costing and project planning; it is not the final design. Many details will change, based on the results of technology development and design activities, but the overall scope and general characteristics of the experiment will not change.

The reference design has 18 0.5-m-class refractors (small-aperture telescopes, or SATs) for measuring large angular scales in the r survey. The basic characteristics of the SATs are shown in Table 3-1. It also has three 6-m crossed-Dragone telescopes (large-aperture telescopes, or LATs) for measuring arcminute scales for the N_{eff} survey, and for delensing the r survey. The basic characteristics of the LATs are shown in Table 3-2.

The full system includes over 500k horn-fed TES detectors with time-division multiplexing (TDM) readout, with band centers ranging from 20 to 270 GHz. To keep the number of telescopes reasonable, most of the camera pixels are dichroic; each pixel has two frequency bands and two polarizations, for a total of four detectors per pixel. SATs deployed in Chile have fast polarization modulators to reduce the effects of atmospheric brightness fluctuations, but the SATs at the South Pole do not.

Property	LF		CF High		CF Low		HF	
Center frequency (GHz)	30	40	85	145	95	155	220	270
Primary lens diameter (cm)	55	55	55	55	55	55	44	44
FWHM (arcmin)	72.8	72.8	25.5	25.5	22.7	22.7	13	13
Fractional bandwidth	0.3	0.3	0.24	0.22	0.24	0.22	0.22	0.22
NET ($\mu\text{K}\sqrt{\text{s}}$) per detector	177	224	270	238	309	331	747	1281
N_{det} per optics tube	288	288	3524	3524	3524	3524	8438	8438
N_{tubes}	2		6		6		4	
N_{wafers}	24		72		72		36	
$N_{\text{wafers total}}$	204							
$N_{\text{detectors}}$	576	576	21144	21144	21144	21144	33752	33752
$N_{\text{detectors total}}$	153232							
Data rate (18 optics tubes)	1.7 TB/day							

Table 3-1. *Small-aperture telescope (SAT) receiver properties.*

All the raw data from the Chile site (approximately 1 Gbps or 4 PB/yr) will be returned to the US promptly for analysis and permanent storage. Only a small subset of the South Pole data can be returned via satellite, with the full data set being stored on disk and shipped to the US every austral summer. This difference in data link bandwidth means that while Chile needs only modest local computing resources, the South Pole site may have to support reduction to maps.

3.4 Manufacturability and reliability

The CMB-S4 reference design is based on a subset of Stage-3 technologies that worked well, were relatively straightforward to manufacture, and appear to be scalable, so we are confident that the reference design can

Property	ULF	LF		MF		HF	
Center frequency (GHz)	20	27	39	93	145	225	278
FWHM (arcmin)	10.0	7.4	5.1	2.2	1.4	1.0	0.9
Fractional bandwidth	0.25	0.22	0.46	0.38	0.28	0.27	0.16
NET ($\mu\text{K}\sqrt{\text{s}}$) per detector	438	383	250	302	356	737	1840
$N_{\text{detectors}}$ per tube	160	320	320	3460	3460	3744	3744
N_{wafers} per tube	4	4		4		4	

Chile (Wide Field Survey – 2 LATs)

N_{tubes} per LAT	0	2	12	5
Data rate (2 LATs)	10.8 TB/day			

South Pole (Delensing Survey – 1 LAT)

N_{tubes}	1	2	12	4
Data rate (1 LAT)	5.0 TB/day			

Total (3 LATs)

$N_{\text{detectors}}$	160	1920	1920	124560	124560	52416	52416
$N_{\text{detectors}}$ total	357952						
N_{wafers}	4	24	144	56			
N_{wafers} total	228						

Table 3-2. Large-aperture telescope (LAT) receiver properties.

be built, and that it will meet the performance requirements. What has to change for CMB-S4 is the level of quality control, in particular a better understanding of the requirements for subsystems, thorough testing, and testing on time. CMB-S4 cameras contain an enormous number of parts which must be tested cold, and at varying levels of integration, so a large part of the project planning is focused on testing.

The scale of CMB-S4 will require a more organized approach to manufacturing, most obviously for the detectors and readout electronics. Anti-reflection-coated lenses, cryostat windows, and thermal filters, will also be needed in quantities that will be a challenge for suppliers of Stage-3 experiment parts. All of these items are included in pre-project technical development or in the early stages of on-project development.

3.5 Design approach

CMB-S4 is a large experiment that involves complex technologies ranging from large precision structures to micro fabrication of superconducting detectors. Matching experts with critical design tasks will require a mix of in-house designers, in universities and national laboratories, and industry partners. Stage-3 experiments followed a similar approach, so we have substantial experience in choosing what to do in-house and what to contract out. The large telescopes, and mounts for the small telescopes, are probably best designed by industry, e.g., through design/build contracts. Cameras, data acquisition systems, and analysis tools involve more specialized technologies and will probably be designed in-house, though many parts will be manufactured under industry contracts.

In-house design work will be done by a mix of university researchers and national laboratory staff. National laboratories will focus on cryostats, detectors and readout, and data management, drawing on expertise in universities through subcontracts. DOE project funds will go to a lead national laboratory, and then to other national laboratories and universities via subcontracts. NSF-funded universities will focus on telescopes, sites, simulations, and analysis. NSF project funds will go to a lead institution, and then to universities via subcontracts.

Pre-project technical development for CMB-S4 is focused on reducing the highest risks (see Sect. 6.4). Our top priority is demonstrating detector noise performance appropriate for CMB-S4, at full wafer scale, with two different readout technologies. We also have some development work on technology options and simulations that will be needed to support the detailed design of CMB-S4. Development activities must be complete before final design starts, which occurs at different times for different subsystems in the reference design's technically limited schedule. Technical development is currently funded through LDRDs and DOE awards to national laboratories, and we have submitted an NSF proposal for project development funds.

3.6 Options

Viable options are an important part of any risk mitigation plan. The highest risks for CMB-S4 are in detectors and readout, so our options include two readout technologies and an alternative approach for detector fabrication. In some cases, options can reduce the overall cost of the project, though a more likely outcome is improved performance for a given cost. Collaborations tend to grow, and this often brings new options because new partners typically want to contribute their own designs and technologies. Exploring new possibilities helps to advance technology in the field, which expands the broader impact of the project. While options are generally a good thing, they can derail the overall effort; a clear, timely, selection process is critical to avoid wasting resources on bad ideas.

Options are a central component of the pre-project development for CMB-S4. We are actively working on: re-use of existing telescopes and site infrastructure; fmux versus μ mux readout; kinetic-inductance detectors; lenslet coupling; planar antennas; and commercial detector fabrication. Details are given in Appendix B. In a technically limited schedule, the time for development is short, so it is critically important to meet goals for making down-select decisions. A simple, open process is needed to develop selection criteria, review the work that has been done, and recommend a path forward. The process must also include an a priori plan for dealing with the very likely event of no clear preference. The best solution is sometimes to implement multiple options, but this should be decided well in advance to minimize disruption.

Reference Design

The reference design for CMB-S4 is a point design that establishes the feasibility of the experiment and provides a basis for preliminary costing and project planning. This design and the technologies used for its components are in advanced states of readiness. However, given the ongoing development of the present generation of CMB experiments and the heritage of technical innovation from the US CMB community we expect to capitalize on technology advances for the final experiment.

In this section we describe the technical details of the reference design. We organize this according to the major subsystems which are: (1) Sites; (2) Large-aperture telescopes; (3) Small-aperture telescopes; (4) Detectors and readout; (5) Data acquisition and control; (6) Data management; and (7) Integration and commissioning. Each subsection provides a discussion of the design drivers, a description of the design, and a presentation of implementation details.

4.1 Sites

The CMB-S4 scientific requirements necessitate a site with access to wide area sky coverage and a site that allows for the small area ultra-deep survey. Chile and the South Pole are the only mature sites that can meet these two needs. In this section we describe the flow down from the science requirements through the instrument design to the site infrastructure requirements needed to support the CMB-S4 experiment.

4.1.1 Site locations: design drivers and overview

State-of-the-art CMB observations require the highest and driest sites on Earth to achieve high atmospheric transmittance and low atmospheric emission within the spectral bands required. The two best sites developed for CMB are South Pole in Antarctica and Cerro Toco in the Chilean Andes (their current states being shown in Figs. 29 and 30). We baseline using the Cerro Toco site in Chile, but will conduct an early trade study to confirm this choice, as compared with using other nearby Chilean sites. The atmospheric transmittance at the relevant frequencies has been studied extensively (e.g., Refs. [501, 502]). Both of these sites have been in use for several decades and have hosted many CMB telescopes (see, for example, [496, 497, 503, 504, 505]). Known information about the atmospheric conditions at these sites was used in the sensitivity forecasts described in Appendix A.

The Chilean Atacama site, at a latitude of 23° South, has access to more than 80% of the sky, enabling observations of the 70% of the sky required by many of the science goals discussed above. The South Pole site, on the other hand, is better suited to deep integration on a small portion of the sky. As described in Chapters 2 and 3, this makes the South Pole site particularly well-suited to making observations on the ultra-deep field. The reference design presented here achieves the CMB-S4 science goals with the configuration of telescopes shown in Table 4-1. If r is measured between now and the Preliminary Baseline review (DOE CD-1/3a, currently planned for Q3 2021), we will re-evaluate this configuration in the context of the measured value of r . Up until that point, both sites will be designed to be capable of hosting all the SATs.




Figure 29. South Pole mm-wave instrumentation site as it currently exists. CMB-S4 would expand on the existing infrastructure at this site.



Figure 30. CMB observatories in Chile are all at this Cerro Toco site. The white arrows indicate possible LAT locations that would not conflict with the preliminary Simons Observatory (SO) instrumentation layout.

Site	Number of LATs	Number of SAT cryostats	Number of individual SATs
Chile	2	0	0
South Pole	1	6	18

Table 4-1. Number of LATs and SATs at each site.

Equipment	Average power	
	Chile	SP
	Total [kW]	Total [kW]
LATs (including receiver and readout but not DAQ)	210	105
SATs (including receiver and readout but not DAQ)	0	264
Cooling system	30	14
DAQ, compute, and office	44	34
Data management	10	50
Site power	30	20
Total	304	487

Table 4-2. Power budget for each site. Note that requirements will be re-evaluated in the context of any redistribution of SATs.

The long history of CMB observations from these sites has led to significant infrastructure buildup at each, and a solid understanding of the additional infrastructure required to enable CMB-S4. Large-aperture and small-aperture telescopes of similar size to those required for CMB-S4 have already been deployed to each site; the primary new features of CMB-S4 which drive project and site requirements are the number of detectors per telescope, and the number of small-aperture telescopes that will be deployed. These demand a larger-scale integration effort, higher power requirements, and higher data bandwidths than previously fielded experiments at each site. Options to re-use infrastructure are discussed in Sect. 3.6. We assume here that we will not re-use any existing CMB instrumentation infrastructure at the site, and cost the site development appropriately. The infrastructure required at each site is described below. Two critical inputs from the instrument to the infrastructure required are the overall electrical power that will be used and data rate that will be generated. Data rate calculations are shown in Chapter 3. A summary of the electrical power requirements is in Table 4-2. Implementation of the power generation necessary to supply this power will be an important long-lead time item to consider in site construction.

It will be critical to begin planning for long lead-time site infrastructure, including foundations, buildings and power, at both sites to ensure they are ready in time for telescope deliveries.

4.1.2 Chilean site

A possible layout of the Chilean site is shown in Fig. 30. This layout uses the area that has been studied by existing CMB instruments. The maximum horizon blockage seen within this area (presented by Cerro Toco) is 15° . Rock in this area is generally appropriate for construction of the necessary foundations, and the instruments will not present significant horizon blockage to existing instruments.

The critical schedule path for the site work in Chile is either the telescope foundation fabrication or the political negotiations necessary to establish a legal presence in Chile and grant that legal entity the right to build on the site. If either these tasks or the inputs necessary to complete these tasks are delayed, the readiness of the Chilean site will be delayed. Note, however, that site readiness is not currently on the critical path of the project.

The following subsections are organized according to the work breakdown structure (WBS) for site work in Chile. While most of this WBS organization is similar to that for the South Pole (Sect. 4.1.3), there are important differences. The low-elevation facility necessary in Chile is functionally equivalent to the South Pole station where people live and eat while deployed. The safety requirements at both sites are critical, with Chile having the added issue of high elevation, and the project having to supply its own safety infrastructure.

4.1.2.1 Management, personnel, and safety

Management, Personnel and Safety includes some of the most important parts of the site work package. These are critical to support the successful integration and commissioning of the instruments, and will lay the groundwork for the operations phase of the project.

Political and legal issues: The major tasks that need to be done to have formal approval of the project's siting in Chile are as follows: (1) determine who will represent CMB-S4 legally in Chile; (2) establish that legal presence, and ensure they can represent CMB-S4; (3) determine desired location of site; (4) determine if our agreement will be with Parque Astronómico or CONICYT; (5) negotiate agreement with correct government agency; (6) determine outreach plan; and (7) sign off on agreement for CMB-S4 in Chile. Items (1)–(5) could each take many months. Item (3) can be in parallel with others, but otherwise they are somewhat required to be in series. The CMB-S4 team aims to mitigate the schedule risk by aiming to retire the first four steps in this negotiation as part of the interim project office (IPO) work between CD-0 and CD-1.

Outreach to the Chilean community: Public outreach within Chile by CMB-S4 is critical to satisfy the requirements of the Chilean government and to mitigate the risk of Chilean community members objecting to our use of their resources. We will develop an outreach program that is visible and has the backing of the Chilean government.

Personnel: Personnel here can be separated into US management and engineering personnel, who are needed early in the project to shepherd the political and design process, and Chilean on-site management, engineering, and technical personnel who will implement plans on-site. It is key that the relevant civil and mechanical engineering personnel be identified on-schedule so that design schedule risk can be retired as soon as possible after receiving the relevant input from: (1) the vendors supplying the LAT and SAT mount; and (2) the geotechnical study of the area once the site layout is established. This engineering effort will be coordinated with Chilean contractors who have important expertise.

Safety: The high-elevation safety requirements set by Chilean law will be complied with; this requires a safety officer knowledgeable about those requirements, and the presence on-site of a safety officer during major construction, integration, and commissioning activities. There is significant local expertise (because of previous astronomy and mining projects) to fill this position. A safety policy will be developed consistent with Chilean law and constraints coming from the relevant US institutions, and that safety policy will be reviewed by an external panel and the safety officer.

4.1.2.2 Site infrastructure and logistics

This work package contains a large number of separate tasks. None of them are schedule-critical provided site permission (described in the section on political and legal issues above) is obtained in a timely manner. The major tasks that need to be done to have formal approval of the project's siting in Chile include plans for: (1) accessing the site with people and equipment; (2) a site layout; (3) power generation; (4) electrical distribution; (5) data/communication to and within the site; (6) non-telescope foundations, buildings and interconnects; (7) cooling system; (8) site monitoring for data quality; (9) tools and supplies; (10) telescope infrastructure; and (11) the low-elevation facility.

Item (4) is discussed in more detail in Chapter 3. We note here simply that there is a path to full-bandwidth communication from the Chilean site to North America and beyond with ms-level latency. Power generation in the baseline design is achieved through diesel generators, though we note that there are risks associated with this that could be mitigated by implementation of solar and wind power. Some other important points are discussed in more detail below.

Low-elevation facility: Personnel working at the Chilean site will commute daily to the site from the town of San Pedro de Atacama, or somewhere in the vicinity. Because of the tourist economy in San Pedro, there are a number of places capable of housing and providing food for a deployment team. Requirements for the low-elevation facility should be developed early so that we can reach an agreement with one of these establishments and retire cost risk associated with this activity.

Telescope infrastructure and non-telescope foundations, buildings, and interconnects: The telescope infrastructure itself, for the purpose of this work package, is simply the mechanical interface to the telescope. This is the foundation that provides the required stability to meet the long-term tilting requirement of the telescope, and to prevent the telescope from falling over in wind, snow load, or earthquake conditions. The foundation design will be significantly different at the two sites, since at the South Pole ice pads are used, while in Chile the foundation will be cement anchored to rock. For this reason, foundation design is part of the site work package. The metal anchoring ring (foundation insert) that inserts in the foundation and provides the attachment points for the telescope is part of the LAT and SAT work packages for their respective foundations. The site team is responsible for providing a summary of the possible environmental conditions to the telescope vendors. The telescope vendors will use that input to calculate possible loads on the foundation created by the telescope in all environmental conditions. The LAT and SAT vendors will then be expected to provide an interface specification with the design of that metal foundation insert and the loads that they have calculated. The site team is then responsible for contracting the design of a foundation that provides the necessary anchoring.

The other buildings (such as labs for receiver assembly, office to house personnel, and structures to house utilities), the foundations associated with them, and the interconnects between them, all need to be laid out and specified based on the instrument requirements. Establishing these instrument requirements in time for the necessary procurement will be an important task for each instrument team and will require close interfacing between the instrument teams and the site design team.

4.1.3 South Pole

At the South Pole, site infrastructure for previous experiments has been designed and installed collaboratively with NSF and their primary contractor for Antarctic logistics, currently ASC. We expect to follow the same path for CMB-S4's site infrastructure requirements, as detailed below. Outdoor construction is done in the austral summer (November–February), while interior work can be done in the austral winter (February–November). Cargo arrives at the Pole in the austral summer, either by C-130 or via a month-long land traverse from McMurdo. Cargo gets to McMurdo either via air (C-17 or C-130) throughout the season, or by ship late in the austral summer. The slower methods (ship and traverse) are preferable from a cost standpoint, but the air routes are quite capable; SPT, for example, was shipped entirely by air transport to Pole. We currently baseline a mixture of air transport and traverse for CMB-S4.

4.1.3.1 Laboratory building

CMB-S4 will require a central space (used for the LAT and SATs) with a variety of lab facilities, such as clean room space for assembling focal planes, a small shop for mechanical work, an electronics space for assembling, debugging and repairing electronics, and a room to house a data analysis cluster and data storage. The scale of this building is expected to be similar to the Dark Sector Lab (DSL) that currently houses BICEP and is attached to SPT, or the Martin Pomerantz Observatory (MAPO) which is attached to the Keck Array.

4.1.3.2 Telescope buildings

A longstanding practice for winterover telescopes at the South Pole is to have all drive components and receivers accessible from a heated space. BICEP is built into the top of DSL. SPT is built on a separate tower, with a heated equipment room beneath it from which the receiver can be accessed when the telescope is docked. The Keck Array is on a tower with a room below the telescope, similar to SPT, but connected to a different building (MAPO).

All the new CMB-S4 telescopes will be placed on towers, with heated equipment rooms below them, attached to the new laboratory building mentioned above by enclosed, heated walkways. This is critical to enabling ongoing summer and winter maintenance and operations. The towers will be designed to meet specifications spelled out by the telescope contractor, but will not be part of the telescope test build in the United States. They will be installed on site prior to the telescopes' arrival. We have budgeted for these towers and equipment rooms by scaling the cost of the SPT versions of these to today's dollars.

4.1.3.3 South Pole Station support

NSF's Admundsen-Scott South Pole Station supports all the science projects located there. The elevated station can house approximately 155 people; overflow buildings have been used in busy construction seasons, boosting the available population above 200. Winter population has recently been between 40 and 50, but has been much higher in the past during periods of major construction.

NSF provides housing, meals, cargo services, and data transmission for all projects. Large projects such as CMB-S4 are charged a per-person rate for flights, housing and meals, and a rate per pound for cargo

depending on the transport type. We have used the NSF-provided rates for these in our budget. Data transmission is currently provided as a service at no cost to the projects, but is limited to existing channels.

As shown in Table 4-2, the CMB-S4 average power requirement is nearly 500 kW. This will likely require new power generation equipment at the South Pole. We are working with NSF to better understand the currently available power limit, and to ensure CMB-S4's future power requirements can be met.

4.2 Large telescopes

4.2.1 Requirements and design drivers

The wide field survey drives the key measurement requirements for the large telescopes (see Chapter 3):

1. angular resolution < 1.5 arcmin at 150 GHz;
2. noise level of a few μK -arcmin at 95 and 150 GHz over approximately 70% of the sky;
3. measurements at frequencies up to at least 270 GHz to control foregrounds and enable astrophysics science goals.

The angular resolution and frequency range, combined with typical characteristics of the atmosphere at a good millimeter-wavelength site, and a typical CMB experiment scan strategy, lead to the telescope design requirements in Table 4-3. The delivered half wavefront error (HWFE), or equivalent surface error, ensures diffraction-limited performance at $\lambda = 1$ mm: the scan pointing knowledge ensures that mixing of E to B modes due to pointing errors does not limit measurements of r [506]; the scan following error is small enough to ensure uniform coverage of the target field; the scan speed is fast enough to freeze atmospheric brightness fluctuations (typically at a height of 1 km, and drifting along at 10 m s^{-1} , so the speed on the sky is 0.6 deg s^{-1}); the scan turn around time allows efficient observations (scans are typically a few tens of seconds, so any wasted time must be just a few seconds); and the emissivity is small enough to ensure that loading on the detectors is dominated by the CMB and atmosphere, not the telescope.

Achieving the low noise requires $\approx 250,000$ detectors, but the large telescopes in Stage-3 experiments only support $\approx 10,000$ detectors. Large telescopes are expensive, so a design with many more detectors per telescope is critical for controlling the overall cost. A wide-field telescope design can support $\approx 100,000$ detectors; just three of these telescopes will be needed for both the deep and wide field survey and for the ultra-deep survey.

Low noise is only useful if systematic errors are also small. Systematic errors are dominated by unwanted pickup through the telescope sidelobes. Ground pickup is the main effect, but pickup from the galaxy, Sun, and Moon are also a concern. Minimizing the telescope sidelobes is therefore a key design driver. Sidelobe calculations are notoriously difficult and inaccurate, so the design has to rely heavily on practical experience from Stage-3 experiments. Small sidelobes require: (i) an unblocked (i.e., offset) optical configuration; (ii) low illumination of the mirror edges (i.e., small spillover); (iii) small scattering due to gaps between mirror panels; (iv) large clearances to avoid clipping of the beam by telescope structures; and (v) absorbing baffles and reflective shields to terminate unwanted sidelobes on stable surfaces, or on the cold sky. The large telescopes in Stage-3 experiments all have reflective, comoving shields that wrap around the lower part of the telescope; CMB-S4 telescopes will have more extensive shielding around the entire telescope.

Parameter	Value	Notes
Aperture diameter ^a	5.5 m	1.5 arcmin beamwidth at 150 GHz
HWFE	< 37 μm rms	80% Strehl ratio at $\lambda = 1$ mm
Scan pointing knowledge	< 5 arcsec rms	< 1/10th beamwidth at $\lambda = 1$ mm
Scan following error	< 10 arcsec rms	1/4 beamwidth at $\lambda = 1$ mm
Scan speed AZ(EL)	3(1.5) deg s ⁻¹	≈ 1 deg s ⁻¹ on the sky to freeze atmosphere
Scan turn around time	2 s	requires 6 deg s ⁻² acceleration
Emissivity	< 0.01	cf. a few % atmospheric loss at mm wavelengths

Table 4-3. Large telescope design requirements. Note: (a) the physical aperture should be a little larger, e.g., 6 m, in order to reduce spillover and scattering.

Spillover onto warm surfaces increases loading on the detectors, which in turn increases the noise. Spillover is always a challenge, but the requirements for CMB-S4 are similar to those for Stage-3 experiments, so the same general approach can be used: a cold stop to limit the size of the beam on the mirrors; and oversized mirrors to direct any spillover onto the cold sky.

Polarization errors are not a strong design driver for the large-telescope mirror configuration, because those polarization errors are generally much smaller than those associated with the planar antennas or horns that feed the detectors.

If CMB-S4 is to achieve its science goals in a reasonable time, the observing efficiency must be high, so the telescopes must be reliable and easy to maintain. In practice, this means locating the drive mechanisms and electronics in protected areas, with easy access, providing lifting equipment to support the removal and installation of heavy motors, gearboxes, and brakes, and planning for on-site replacement of the axis bearings.

4.2.2 Reference design summary

The large telescopes are a 6-m crossed-Dragone design on a conventional elevation over azimuth mount, with the instrument boresight aligned with the elevation axis. This alignment facilitates instrument rotation and enables partial telescope boresight rotation, because the elevation axis can rotate a full 180°, from horizon to horizon, enabling two distinct telescope orientations at every azimuth position. The design copies the SO-LAT and CCAT-prime, and was adopted because this is the only high-throughput design that will be demonstrated on the timescale of CMB-S4.

The reference design has two new large telescopes in Chile to support the deep and wide survey. An option is to re-use the SO-LAT and CCAT-prime, in which case no new large telescopes would be built for Chile. A third large telescope is needed at the South Pole to support the ultra-deep survey. The South Pole large telescope has the same optical configuration as the SO-LAT, but with some modifications to the mount to accommodate logistics constraints and operation in the cold.

Since the telescopes are copies of the SO-LAT and CCAT-prime, no prototyping is needed, and procurement contracts can be awarded early, so telescope testing can be completed before cameras are delivered.




Figure 31. The CMB-S4 Large Aperture Telescope reference design is based on the design of the SO-LAT and CCAT-prime telescopes [507, 508]. The mirrors (M1 and M2) are completely enclosed in the co-moving elevation structure (left), which will improve sidelobe mitigation compared to existing telescopes. The Large Aperture Telescope receiver (LATR, right) is aligned with the telescope elevation axis. This enables two new and valuable features for suppressing systematics compared to existing large telescopes: (1) the LATR can either track the elevation angle of the telescope or to be rotated independently of the telescope angle; and (2) partial boresight rotation of the telescope can be achieved because the telescope boresight can be rotated from one horizon to the other, between 0 and 180°.

4.2.3 Optics

The high-throughput off-axis crossed-Dragone (CD) telescope design was first proposed by Corrado Dragone [509]. This configuration has previously been used for CMB measurements by the QUIET [510] and ABS [511] projects. Larger-aperture CD designs were recently shown to be capable of achieving many times larger throughput than existing 5–10 m microwave telescopes, particularly when combined with refractive optics tubes [512]. In 2017 a 6 m CD design that includes coma corrections [513] was adopted for CCAT-prime and the Simons Observatory (Fig. 31) [514, 508, 507], and both telescopes are scheduled to begin observations in Chile in 2021.

The CD layout offers several advantages compared to other telescope designs. For CMB-S4 the most important advantages are the large diffraction-limited field of view (FOV) combined with the compactness of the two mirror configuration (Figs. 31 and 32). While modestly larger diffraction-limited FOVs have been achieved with other telescope configurations, including three mirror anastigmat designs, the alternative designs were not nearly as compact and use of a third mirror will degrade the sensitivity of the detectors due to the extra mirror emission.

The polarization properties of the CD configuration have been shown to be superior to Gregorian and Cassegrain designs, which is in part due to the smaller angles of incidence on the reflecting surfaces [515]. Systematics associated with the telescope mirrors are much smaller than systematics caused by the lenses, filters, and feedhorns inside the cameras [516].

The mirrors are each approximately 6 m in diameter and will be segmented into many panels (Fig. 33) [507]. All mirrors on CMB telescopes near this diameter have been segmented thus far. The sidelobe pickup




Figure 32. Mirror optics definition (left) for the CCAT-prime and SO-LAT [514]. The pink rays indicate the geometric keepout zones for a 7.8° diameter FOV illuminating the full 6-m aperture primary mirror. The primary and secondary instrument spaces are also shown. The mirrors alone provide diffraction-limited image quality across much of the usable FOV as shown by the contours (right). The contour colors (blue through red) show the regions where the Strehl ratio is $> 80\%$ at 870, 490, 345, 230, 150, 100, and 75 GHz [507]. Refractive optics can improve the optical performance beyond what is shown here within individual optics tubes.




Figure 33. *Left: Mirror panel views, reflecting surface (top), backside (bottom). The panel is 700 mm on a side and significantly light-weighted from a solid block of aluminum. Eight adjusters locate the panel. The five z-axis adjusters allow for some compensation of low order distortions. Middle: Three views of a preliminary carbon fiber bus structure for the primary mirror panels. The red elements are tuned CTE beams that minimize the half wavefront error and pointing error due to imperfect mirror alignment from thermal effects. Right: Similar views of the carbon fiber bus structure for the secondary mirror panels.*

associated with the gaps between the mirrors is sufficiently small that it will not contaminate the high multipole E -mode polarization, lensing B modes, or temperature signals that will be the foci of measurements with these telescopes. The gaps between the mirrors for the CCAT-prime/SO-LAT design are required to be $<0.5\%$ per mirror at the nominal observing temperatures. This requirement both minimizes loss in the gaps and minimizes the amplitude of the sidelobes due to the panel gaps. The mirrors are supported by carbon fiber bus structures as shown in Fig. 33 to minimize variations in HWFE and pointing error due to thermal variations. Table 4-4 shows the HWFE budget.

The reference design provides a CD FOV that is >2 m diameter. An FOV this large can make coupling to arrays of detectors challenging. A natural approach to take advantage of such as large FOV is to split it into multiple independent optical paths or optics tubes [512, 517]. These optics tubes are designed to be modular and easily replaceable, which facilitates deployment of different frequencies in each optics tube. Each tube has an independent image of the primary, or Lyot stop, to control illumination of the telescope mirrors. The stop is cooled to 1 K to minimize excess background loading on the detectors. The reference Lyot stop design would geometrically illuminate 5.5 m of the 6.0-m aperture primary mirror. Thus, the mirror will provide a 0.25-m baffle in radius to control diffracted spillover, and designs for low precision baffles to extended the mirrors are being studied now [516].

Several different sizes of optics tubes have been studied for the SO-LAT/CCAT-prime telescope design, ranging from 0.2-m to 2-m diameter optics. There are a variety of practical constraints on the detector array size (detectors are typically fabricated on 150-mm wafers), availability of optics materials (e.g., the maximum single crystal silicon diameter is 450 mm), vacuum windows, filters, and mechanical structure between tubes that influence the optimal optics tube size. A few of the optics tube options that were studied and compared are discussed in [517, 518]. These and related cost versus sensitivity analyses led to the selection of optics tubes with 450-mm spacing for SO-LAT/CCAT-prime. A preliminary design for a 19-optics-tube camera for these telescopes is shown in Fig. 34.

The detector arrays within each optics tube will be feedhorn-coupled transition-edge sensor detectors. Feedhorns are one of the most mature technologies for controlling polarization systematics. They have circular beams and small polarization errors. By coupling the feedhorns through a cryogenic Lyot stop,




Figure 34. *Left: Raytrace of the CD telescope design with 19 optics tubes that are each designed to illuminate 5.5 m of the 6-m aperture telescope. Middle: Side view showing five of the 19 optics tubes. Right: Preliminary design for a 19-optics-tube cryostat developed by the Simons Observatory collaboration.*

Subsystem	HWFE ($\mu\text{m rms}$)
Primary	18
Secondary	15
Telescope alignment	21
Camera filters	10
Camera lenses	10
Camera alignment	10
Total	36

Table 4-4. *HWFE budget. Telescope errors include gravity, 15-K ambient temperature change, 1.4-K temperature gradients (in direct sunlight), 9 m s^{-1} wind (3rd quartile in Chile), and 0 scan acceleration.*

they can be packed much more closely together, which improves the overall instrument sensitivity, but can significantly increase detector count. Roughly half of the beam from each feedhorn falls on the Lyot stop (see Table 4-5), but the sensitivity degradation is small because the stop is very cold.

4.2.4 Mount

The mount is a fork-style elevation over azimuth design, sized to accommodate the crossed-Dragone optical layout. The mount has three main parts: a fixed support cone at the base; the yoke, which rotates in azimuth and contains all the science equipment; and the elevation housing that supports the mirrors. The design allows for elevation angles past zenith, which is useful for measuring systematic errors, and can even be driven to -90° to facilitate maintenance. The chief ray between the secondary and the focal plane is along the elevation axis, so the camera can be attached to the yoke structure. Gravity acts mostly in-plane for the secondary mirror, minimizing gravitational deformation. The overall size of the structure is approximately 23 m long \times 8 m wide \times 16 m tall, with the elevation axis \approx 11 m above ground for the Chile version. The total weight is 210 T (see Table 4-6).

The elevation housing contains the primary and secondary mirrors. The beam from the sky enters through a shutter that can be closed during bad weather, reflects off of the primary, over to the secondary, and then

Subsystem	Loss 30/40/95/150/220/270 GHz	Temperature (K)	Notes
Telescope mirrors			
Spillover	0.02	300	2 mirrors
Ohmic loss	0/0/0.004/0.01/0.02/0.03	300	2 mirrors
Camera window	0.001/0.002/0.005/0.01/0.015/0.02	300	1/2 inch UHMWPE
Filters			
Alumina	0.0002/0.0003/0.0006/0.001/0.0015/0.002	80	2 mm thick
IR blockers	0.004	300, 80, 40, 4	4 blockers
Low-pass	0.03	4, 1, 0.1	3 filters
Lenses	0.006/0.009/0.02/0.03/0.05/0.06	4, 1, 1	3 lenses
Feedhorn	0.72/0.51/0.75/0.57/0.35/0.20	1	Spillover on cold stop
Total	0.77/0.57/0.82/0.67/0.48/0.36		

Table 4-5. Optical loss budget. See Table 4.3.4.4 for filter details.

exits at a right angle to the entrance beam through the 4-m ID elevation bearing. The focal plane is ≈ 3 m beyond the bearing.

The mount is a welded steel structure, split up into transportable sections. The mirrors have machined aluminum panels, ≈ 0.7 m on a side, mounted on a CFRP back up structure. Each panel has three adjusters to set the in-plane position, and five adjusters to set the height of the reflecting surface. The five height adjusters allow some correction of low-order machining errors. The back up structure is kinematically attached to the elevation housing via a 6-point connection with three vertical and three horizontal constraints.

The outside of the structure is covered with aluminum-clad insulated panels to reduce pointing errors and HWFE due to thermal gradients, which are driven mainly by solar heating. In order to provide a reasonable operating environment for equipment and personnel, the base of the yoke is temperature controlled to 15 ± 5 C using heaters and waste heat from the science equipment with outside air for cooling. The elevation housing and yoke arms, including the camera space, are actively vented to keep the structures at outside temperature; any electronics attached directly to the camera are also at outside temperature.

Telescope servo cabinets, instrument electronics, and helium compressors are all located inside the base of the yoke, so if the camera does not have a rotator to track parallactic angle, camera connections do not have to pass through a cable wrap. The camera space has a 6-T gantry crane for loading/unloading, positioning, and assembly/disassembly work while the camera is installed.

The scanning requirements of the telescope drive the choice of bearings and drives. The drive has two motor/gearbox units for the elevation axis, and 4 motor/gearbox units for the azimuth axis. The azimuth drives are on the base of the yoke, around the support cone, and the elevation drives are in the yoke arm opposite the camera. Preliminary modal analyses of the structure suggest a first mode with a locked azimuth rotor at 2.6 Hz and a locked rotor elevation mode at 3.0 Hz [514]; both of these frequencies have increased as the design has matured.

The pointing performance of the telescope (see Table 4-7) is limited mainly by deformation due to wind forces and insolation. The effects of slowly varying thermal deformation are partly corrected by continuously measuring the locations of the many point sources in each survey patch.

Component	Mass (T)
Mirrors ^a	5
Camera	5
EL structure	50
Yoke	135
AZ cone	15
Total	210

Table 4-6. Telescope mass. Note: (a) total for primary and secondary, including adjusters and back up structure.

Contribution	Pointing error (arcsec rms)
1.4-K temperature gradients ^a	2.5
9 m s ⁻¹ wind ^b , 15-K ambient temperature change	0.5
Random errors ^c	0.8
Radio pointing ^d	1
Total	2.9

Table 4-7. Scan pointing knowledge error budget. Notes: (a) In direct sunlight; (b) 3rd quartile in Chile; (c) bearing wobble, tiltmeter error, encoder coupling; (d) the pointing offset is measured every $\approx 1/10$ hr using radio observations of bright point sources.

The telescope design meets survival requirements for Chile (see Table 4-8). Survival requirements for the South Pole are less stringent, except for the lower ambient temperature. Design modifications for the South Pole are discussed in Sect. 4.2.9.

4.2.5 Baffles and shields

Errors in the shape of the beam convert temperature fluctuations on the sky to polarization, so control of systematic errors at a level appropriate for CMB-S4 requires accurate knowledge of the beam. Main-beam errors are a tractable problem for the large telescopes because the beam shape can be measured directly, using a bright source, e.g., a planet, or by filtering the CMB observations with templates that capture the effects of beam errors. In addition, CMB fluctuations on arcminute scales are small, so systematics due to beam errors are also small.

Load case	Chile	South Pole
Wind speed (ms ⁻¹)	69	35
Seismic acceleration (g)	1/3	0
Ambient temperature (C)	-30	-90
Ice (cm)	1	1
Snow ^a (cm)	120	0

Table 4-8. Survival requirements. Note: (a) 100 kg m⁻³.

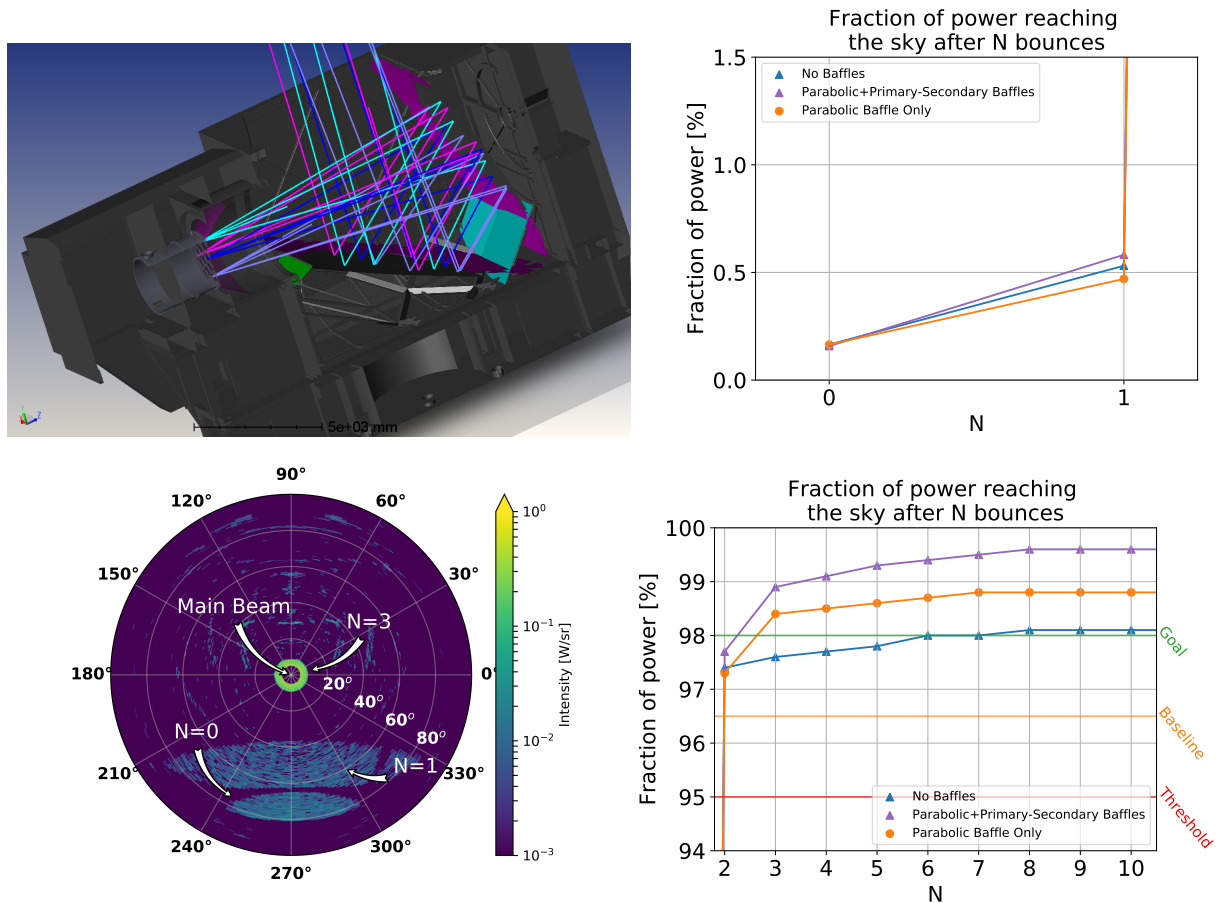


Figure 35. Preliminary baffle and shielding studies for the Large Aperture Telescope using non-sequential ray tracing [516]. The upper left shows a baffle configuration for the Large Aperture Telescope that includes a parabolic baffle near the receiver plus large secondary and primary guard ring baffles. Current analyses suggest that the parabolic baffle (pink, left) is important, but the large guard ring baffles (green, teal, and also pink, right) will not be needed. The position of 4 fields being launched from the focal plane of the telescope is also shown. We note that the large co-moving ground shield from the elevation structure will provide more baffling than in any previous large CMB telescope. Upper right (lower right) shows the fractional power that reaches the sky from the center field of the center camera versus number of bounces for a perfectly reflective telescope solid model for N in the interval 0–1 (2–10) bounces. Lower left shows the power density at the sky for $N=3$ bounces for a configuration with a parabolic baffle alone. Power density at the sky is normalized to have an injected power equal to 1 W. In this configuration, one sidelobe is expected between 60 and 80 degrees, corresponding to the direct line of sight from the receiver camera to the sky. The parabolic baffle forms a ring around the main beam at $N=3$, which contains most of the difference in power between the parabolic baffle and the no-baffle model. Additional analyses that build on these will inform the desired finish (generally absorbing vs. reflective) and design for each of the baffles.

Unwanted pickup through far sidelobes is more complicated, because far sidelobes are difficult to measure or model, especially in polarization. Far sidelobes are generated by scattering from mirror panel gaps, clipping of the beam as it passes by telescope structures, and unwanted lines of sight, e.g., directly from the camera to the sky. The far sidelobes cause unwanted pickup from the ground, Sun, Moon, and galaxy. Recent work using a combination of measurements and ray-tracing has been successfully used to model far sidelobes for ACT [519], and the same techniques are being applied to CCAT-prime and the SO-LAT [516].

Scattering from panel gaps is the dominant source of far sidelobes in telescopes with segmented mirrors. The sidelobes are sharp, arc-like features, that typically account for 1% of the total response [520]. If the sidelobes fall on the ground, they cause enormous pickup, so large CMB telescopes always have comoving ground shields. The shields are usually reflective to direct sidelobes to the cold sky; an absorptive shield would significantly suppress far sidelobes, giving a cleaner beam, but at the expense of increased loading on the detectors, which would result in lower sensitivity. The mirrors in the CCAT-prime and SO-LAT designs are completely enclosed in a metal box, which naturally acts as a comoving shield. In this respect, the design has better shielding than existing Stage-3 large telescopes.

Low spillover is critical for reducing the loading on the detectors. The part of the beam that does spill beyond the mirror edges is usually managed with reflecting baffles that direct the spillover to the sky, rather than to warm absorbers, essentially trading detector loading for sidelobes. Fig. 35 shows a spillover simulation in which rays are transmitted from the camera windows. With no baffles, the spillover on warm surfaces is $\approx 2\%$, but adding baffles around the camera, and behind the mirrors, reduces the warm spillover to $< 1\%$.

An issue that has not yet been fully resolved is whether or not the large telescopes need a fixed ground shield surrounding the entire telescope, as was done for ACT [498]. Simulations of ground pickup during pre-project technical development will inform this decision; our expectation is that fixed ground shields will not be needed, but they are included in the cost estimate for the reference design. The ground shields are simple static structures, so the cost impact is small.

4.2.6 Control and monitoring

The large telescopes have a stand-alone control system that follows a path commanded by the observatory control system, with offsets based on a generic set of scan patterns. The telescope control system works in (AZ, EL) coordinates; conversion from celestial coordinates is handled at higher level in the observatory control system using an open-source astrometry package. The telescope control system commands the motor control loops to move within safe limits, and monitors the actual positions of the telescope axes and the status of the drive systems.

Encoder positions, with accurate timestamps, and a summary of the drive status and faults, are reported at 200 Hz, and are available to the observatory control system within 1 s, which is fast enough to support decisions that are needed to control observations, e.g., waiting for the telescope to acquire a specific position before starting a scan. Encoder positions, and information about the status of the drive, are recorded continuously, even when observations are not running.

The telescope control system provides tools for debugging drive problems, and for acceptance testing, e.g., plots of axis position and following error vs. time, plots of following error spectral density, and measurements of the drive transfer functions.

4.2.7 Assembly and test

The large telescopes will be assembled on-site by a contractor. The contractor could be the telescope designer, a construction company, or some combination of these.

The telescope mirrors need to be assembled and aligned before they are installed in the telescope structure. Assembly and alignment could be done at the factory, as for SO-LAT and CCAT-prime, but if we decide to fly mirrors to the South Pole, they will have to be broken down to fit in a C-130 aircraft, in which case some assembly and alignment will be needed on site. Even if mirrors are aligned at the factory, the contractor will measure the surface profile immediately before installation using photogrammetry or a laser tracker. Photogrammetry measurements of the ALMA antennas achieved $30\ \mu\text{m}/6\ \text{m}$ rms error [521]; a laser tracker is a factor of a few better, with $20\ \mu\text{m}$ peak error on a 6-m part. Mirror surface measurements will be made at the typical observing temperature in Chile, or at ambient austral summer temperature at the South Pole. Following installation in the telescope, the mirror positions will be aligned, again using photogrammetry or a laser tracker.

After assembly, the drive system functions will be tested, and the following error will be measured at the axis encoders to demonstrate that the drive meets the requirements for fast scanning. The pointing stability will be demonstrated using observations of bright stars with a small optical telescope, which is fast, easy, tests many telescope and observatory systems, and quickly reveals problems that are difficult to see with slower radio measurements of the pointing. Drive system functional testing and following error testing are tasks for the contractor, but pointing measurements will require support from CMB-S4.

Final alignment of the telescope optics will be based on out-of-focus holography measurements using a planet as the source. With a signal to noise ratio of 100 at $\lambda = 3\ \text{mm}$, out-of-focus holography can achieve a few $\times 10\ \mu\text{m}$ rms for the lowest-order modes of the wavefront error [522]. Holography measurements will be done by CMB-S4 using either a commissioning camera or the full science camera.

4.2.8 Safety

Telescope safety is managed by a safety-rated programmable logic controller that monitors axis limit switches and speed sensors, e-stop switches, lockouts, and access interlocks. Lights on the outside of the telescope indicate when it is safe to approach. Enclosed spaces have fire and carbon-monoxide sensors, and there are cameras for remote monitoring. The telescope design follows fire and electrical codes (NFPA 70 and similar IEC standards) and structural safety standards (DIN 18800).

4.2.9 Modifications for the South Pole

The SO-LAT and CCAT-prime were designed for operation in Chile, but constraints for the South Pole are different, so some design changes are needed.

1. Transportation. The telescope structure and mirrors may have to be broken up into small pieces that will fit in a C-130 aircraft. Adding bolted connections to the structure generally increases the mass and may also result in lower stiffness. It may be possible to send larger pieces by traverse, but this will involve a 1-year delay because the traverse arrives at the South Pole fairly late in the austral summer. The reference design plan assumes that the large telescope parts will be transported by air.

2. Foundation. Telescopes at the South Pole must be mounted on towers to give a reasonable lifetime before snow starts to bury the structure. Existing telescopes are on 5–8 m towers which sit on wood rafts. Addition of a tower reduces the stiffness of the telescope structure, resulting in worse pointing and scanning performance. The CMB-S4 large telescopes have roughly the same mass and envelope as the existing 10-m SPT, so the reference design assumes a similar tower and raft.
3. Operation in the cold. Temperatures at the South Pole can be as low as -80 C (cf. -20 C for Chile), so bearings, drive motors, electronics, and wiring that are appropriate for Chile may not work at the South Pole. These elements can be heated to address the lower temperatures after enclosing the bearings and isolating the instrument space from ambient temperature using a large conical baffle between the instrument front plate and the bearing at the elevation structure. Heating in general results in thermal deformations, which can affect pointing and will be addressed through a combination of modeling and measurements during commissioning.
4. De-icing. At the South Pole, telescope mirrors generally have to be heated about 1 K above ambient to prevent icing during the long polar night. The temperature rise is small, so thermal deformation is not an issue, but provision must be made to apply heat uniformly over the back of the mirror. For the machined aluminum panels on the 10-m SPT, we typically apply 50 Wm^{-2} to keep the surface clean.
5. Power. Electrical power for the SO-LAT/CCAT-prime is Chilean standard 400V, 3-phase, 50Hz, but the South Pole is US standard 480V, 3-phase, 60Hz. Some changes will be required in the electrical distribution and in equipment with large motors, e.g., helium compressors.

4.3 Cameras for large telescopes

4.3.1 Requirements and design drivers

The key measurement requirement for the large telescope cameras is the map noise, because this determines the number of detectors needed at each observing frequency (see Chapter 3). The field of view of the telescope described in Sect. 4.2 is large enough to accommodate the required detector counts if two-color pixels are used, e.g., 220/270, 95/150, and 30/40 GHz. In practice, some field of view is lost to dead space between the camera optics tubes and between detector wafers in each optics tube, so it becomes challenging to accommodate enough pixels. Achieving the required mapping speed (i.e., enough detectors with sufficiently low noise), without adding an additional, expensive, large telescope, is the principal design driver for the cameras.

An obvious approach for maximizing the number of detectors in the cameras is to pack the optics tubes and detector wafers as close as possible. A single hexagonal detector wafer per optics tube is attractive in this respect, but it requires a large number of small tubes. For the reference design, we have adopted a scheme with three hexagonal wafers and three smaller rhomboid-shaped pieces per tube, but this choice will be revisited during detailed design.

The mapping speed can be increased by making the pixels smaller, so more pixels can be squeezed into the telescope field of view, but the noise in adjacent pixels becomes correlated; there is essentially no improvement in mapping speed for pixel spacings smaller than $\approx 1F\lambda$ [518], so the reference design has $1F\lambda$ pixel spacing.

Achieving sufficiently low detector NET values requires careful control of losses in the camera optics. Small losses require thin lenses (which generally means smaller diameter lenses) made of low-loss material, e.g., Si, with good anti-reflection coatings, e.g., moderate-bandwidth machined structures.

Stage-3 experiments use either 250-mK sorption fridges or 100-mK dilution fridges. Sorption fridges must be cycled, which reduces observing efficiency, but dilution fridges operate continuously. Observing efficiency drives mapping speed, which is critical for CMB-S4, so the reference design uses dilution fridges. The lower base temperature for a dilution fridge reduces phonon noise in the detectors, but the improvement in overall noise is small if the detectors are background limited. Fridges dominate the site power requirements, so the number of fridges impacts the cost of operations.

Polarization errors are usually dominated by the detectors (e.g., differential gain errors between Q and U detectors, ortho-mode transducer alignment errors, and leakage at the horn to detector wafer interface). Small, stable polarization errors can be corrected by filtering the CMB measurements, so absolute polarization errors in CMB-S4 can be similar to those in Stage-3 experiments.

Cool down time is a serious concern, because the cameras for CMB-S4 large telescopes are much bigger than Stage-3 cameras. A cooling time longer than a couple of weeks makes installation and commissioning difficult, and it can severely impact the overall efficiency of the experiment, because cameras do have to be warmed up occasionally for maintenance. Handling is also a concern for large cameras. Minimizing the mass and size of the cameras makes for easier operations and a simpler telescope interface.

4.3.2 General configuration

The camera has a 2.6-m diameter cryostat containing 19 close-packed optics tubes with the detector distribution shown in Table 3-2. The design is based on the SO 13-tube camera shown in Fig. 36, but with a slightly larger cryostat and a full hexagonal close-packed array of optics tubes. Most of the details presented in this section are from the SO 13-tube camera design.

The design is modular, in that all the optics tubes have similar mechanical, thermal, and electrical interfaces. Replacement of an optics tube is relatively straightforward, so the assignment of bands to telescopes can be adjusted to deal with unexpected problems, e.g., complicated spectral behavior in foregrounds. The modular design also allows several teams to work independently on parts for one camera, so fabrication and testing can proceed quickly.

The 19-tube version of the camera will require some structural changes to the front plate, to handle the larger force on the plate, and a custom forging may be needed, because of the size, so the cost will be higher. Internal components, e.g., radiation shields, will have the same general configuration as in Fig. 36, but some parts will require design changes to accommodate 19 optics tubes. The cooling capacity may also have to be increased by adding additional pulse tube refrigerators, more heat straps, and possibly changing to a higher capacity dilution refrigerator. The modifications will not be trivial, and they will result in higher mass and cost.

4.3.3 Cold optics

The cold optics are comprised of silicon lenses, IR-blocking and band-defining filters, and spline-profiled feedhorns.

Silicon lenses were selected due to the combination of a high index of refraction ($n = 3.4$) as required by the optical design, low loss tangent ($\tan \delta \sim 10^{-5}$), and ability to mitigate reflections using the proven approach of machined metamaterial antireflection coatings. These AR coatings have demonstrated control




Figure 36. *SO 13-tube camera design.*

of reflections to 3×10^{-3} , minimizing systematics effects and nearly eliminating sensitivity losses due to reflectance in silicon lenses. The LATR requires lenses with diameters up to 40 cm, while silicon is available up to 46 cm diameter. Fabrication of lenses at the production rate required for CMB-S4 is on track to be demonstrated in early 2019 for Simons Observatory.

The IR-blocking filters are composed of metal mesh shaders which reflect IR radiation used in conjunction with alumina filters which absorb the light. The alumina filters will be AR coated using a laminate of plastic layers as was demonstrated on SPT-3G. The metal mesh filters are fabricated using standard lithographic techniques on dielectric substrates. Three successive low-pass edge filters will be used to block radiation outside the passband of each dichroic detector array. This eliminates susceptibility to blue leaks in the detectors. The low pass edge filters are fabricated by combining a number of patterned mesh filters into a laminated stack.

Finally, radiation arriving at the focal plane is coupled to the detectors through a monolithic array of spline-profiled feedhorns fabricated by machining aluminum. These feeds offer tight control over beam systematic effects, ability to optimize the coupling efficiency, and ability to produce these arrays at a relatively low cost per horn. This approach has been used on SPTPol, MUSTANG-2, and is being scaled up for Simons Observatory. Production at the required rate for CMB-S4 is a low risk proposition.

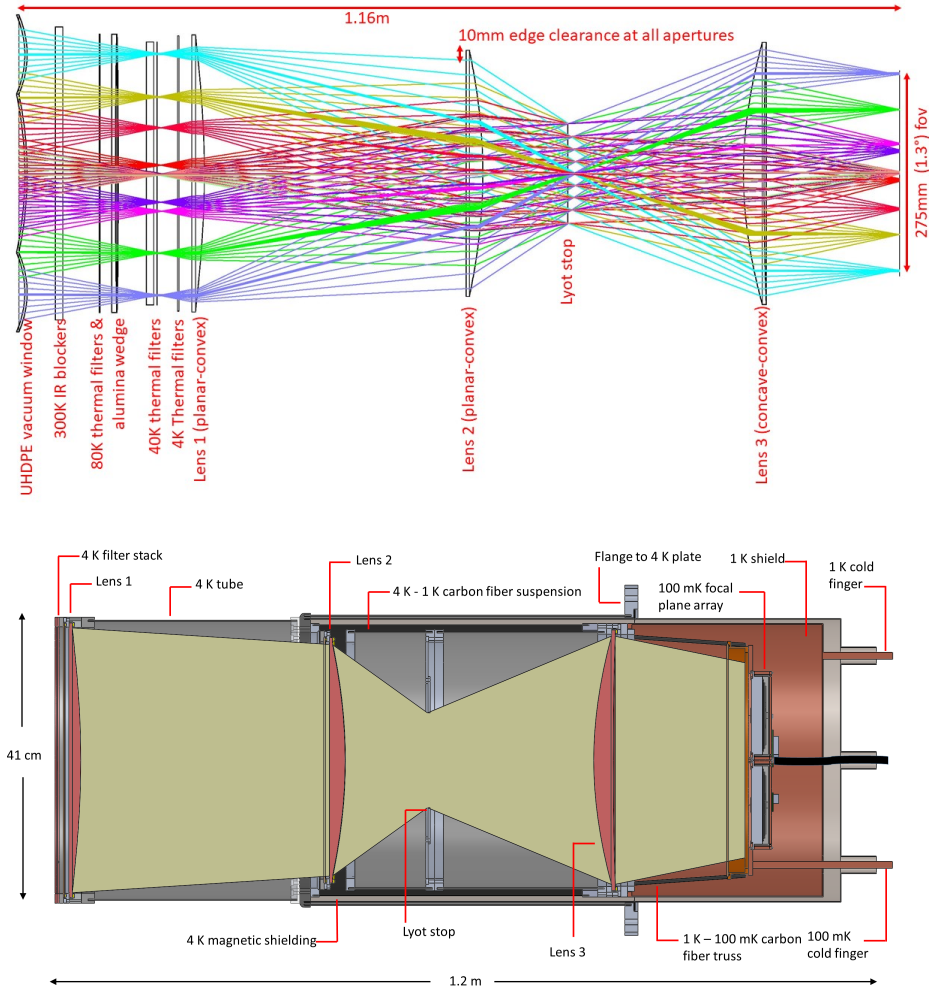


Figure 37. Preliminary optics tube designs for the LAT receiver [517]. The top shows a ray trace and includes labels for all the optical elements. The bottom shows much of the internal structure of the mechanical design that will be used to support, thermally isolate, and cool the optics and detector arrays.

Fridge	Type	Quantity	Cooling capacity per fridge				
			80 K	40 K	4 K	1 K	100 mK
PT90	Pulse tube	2	90 W				
PT420	Pulse tube	2		55 W	2 W		
LD400	Dilution	1				25 mW	400 μ W

Table 4-9. *Refrigerators.*

4.3.4 Cryogenics

4.3.4.1 Refrigerators

The detectors are cooled to ≈ 100 mK, with intermediate stages at 1, 4, 40, and 80 K for radiation shields, filters, and wiring thermal intercepts [523]. The 80-K stage has two single-stage Cryomech PT90 pulse tubes, and the 40-K and 4-K stages have two PT420 pulse tubes (see Table 4.3.4.1). Two additional pulse tubes can be installed in the camera if a faster cooling time is needed. The 100-mK stage is cooled by a BluFors LD400 dilution refrigerator with an intermediate stage at 1 K. The dilution fridge has its own PT420, which provides additional cooling power for the camera 40-K stage. Flexible copper braids are used to connect the refrigerators to the camera cold stages. All the refrigerators must be kept within $\approx 60^\circ$ of vertical to maintain cooling capacity.

4.3.4.2 Heat pipes and switches

Roughly 1500 kg needs to be cooled by the PT420 4-K stages, which have limited cooling power, especially at temperatures $\gtrsim 100$ K. During cooldown, when the camera's 40-K stage is near its base temperature, the 4-K stage will still be at ≈ 200 K, so we would like to transfer heat efficiently between the two stages at high temperature and switch off the transfer at temperatures $\lesssim 40$ K. Possible solutions include a heat pipe filled with high pressure nitrogen gas [524] and pre-cooling using liquid nitrogen. Pre-cooling is used in many laboratories, but is awkward at a telescope site, because it requires large quantities of liquid nitrogen, so we will use heat pipes charged with 1000 psi of nitrogen at room temperature.

Heat switches can be used to significantly reduce the cooling time by temporarily creating strong thermal links between the 1-K and 4-K stages, and between the 100-mK and 1-K stages. The heat switches must have small conductance in the open state and high conductance in the closed state. Gas-gap and mechanical heat switches are viable options; both are available commercially, and samples are being tested.

Gas-gap heat switches have a small helium adsorption pump that can actively fill or evacuate a small cavity between two copper rods inside a thin-walled stainless steel tube. The open conductance is reduced to the minimal conduction through the small cross-sectional area of the stainless steel tube. The closed conductance is several orders of magnitude larger, since the small gap between the cold and warm copper rods is filled with helium gas, which conducts heat across the gap efficiently.

Mechanical heat switches have an actuator to make or break a physical connection. The open conductance can be essentially zero, and the closed conductance can be very high as long as enough pressure can be developed at the interface. Thermal boundary resistance is significant issue at lower temperatures [525], so extensive testing of several mechanical switches is currently underway.




Figure 38. Thermal model of an optics tube showing all the filter elements and tube walls at each temperature stage. The light lines do not represent the actual ray trace; they merely highlight the location of the Lyot stop.

4.3.4.3 Optics tubes

The optics tubes contain all optical and detector components between 4 K and 100 mK. Each tube is self contained, so it can be installed as a single unit (see Fig. 37). An optics tube has three silicon lenses [517]. The first lens is at 4 K, and is supported from the cryostat 4-K plate by an ≈ 1 -m long aluminum tube which has a thin wall to keep the mass low. The tube is fabricated from 1100-H14 or 6063-O aluminum, as a compromise between strength and thermal conductivity, and it is lined with Vacuumschmelze A-4K magnetic shielding material. The second lens, Lyot stop, and third lens are all at 1 K, and are supported from the 4-K tube by a thermally-isolating carbon fiber tube. The inside surfaces of all the tubes that surround the beam are blackened or have annular baffles. The detector assembly is supported by a carbon fiber truss attached to the 1-K structure, with a second layer of magnetic shielding around the detectors. Thermometer wires, cold finger connections to the camera cold stages, and detector readout cables run through the back of the optics tube.

4.3.4.4 Heat loads

The total power emitted and absorbed at each temperature stage has been estimated using a radiative transfer simulation based on numerical ray optics [523]. The simulation captures the optics tube geometry (see Fig. 38) and spectral properties of the tube walls and filters. Calculated heat loads for an optics tube are shown in Table 4.3.4.4 and filter temperatures are shown in Table 4.3.4.4. Heat loads for the complete camera are given in Table 4.3.4.4 [526, 525].

Temperature stage	80 K	40 K	4 K	1 K	100 mK
Optical load	3.90 W	2.0 mW	27.6 mW	29 μ W	312 nW

Table 4-10. Heat loads for a single optics tube.

Stage	300 K	80 K	80 K	40 K	4 K	4 K	1 K	100 mK
Filter type ^a	IRB	IRB	AF	IRB	IRB	LPF	LPF	LPF
Cutoff ^b (μ m)	27	40		300	660	800	1470	1610
Filter center temperature (K)	297	252	81	74	62	12	1.1	0.105

Table 4-11. Optics tube filter temperatures. Notes: (a) IRB=metal-mesh infrared blocking filter, AF=alumina filter, and LPF=low-pass filter; (b) LPF cutoffs are for 90/150-GHz optics tubes.

4.3.5 Mechanics

4.3.5.1 Vacuum shell

The camera vacuum shell presents a significant challenge, because the front plate is large and flat, it has many large holes for the windows, and deflections must be $\lesssim 2$ cm to avoid a touch to the 80-K stage behind. Stresses are highest at the center of the plate, so this region is ≈ 1 cm thicker. The holes for the windows follow the hexagonal shape of the beam, which allows more material to be left in the front plate (see Fig. 39). The window holes also have tapered walls that match the beam divergence, again to allow more material between windows (see Fig. 40) [526]. The plate is machined from a single billet of 6061-T6 aluminum, with aggressive lightweighting on the outside face to keep the overall mass of the camera reasonable (see Fig. 36). The front plate is 7 cm thick, has a total mass of 350 kg, and deforms 1 cm. The complete vacuum shell conforms to the ASME VIII-2 pressure vessel code.

4.3.5.2 Windows

The cryostat windows are made of 3-mm thick ultra high molecular weight polyethylene, anti-reflection coated with porous teflon sheets that are hot-pressed onto the window surfaces. The windows are clamped to the vacuum shell front plate with hexagonal metal rings.

Stage	Support	Radiative	Optical	Readout	Total	Cooling capacity
80 K	1.3 W	7.2 W	50.7 W	N/A	59.2 W	180 W
40 K	6.0 W	4.6 W	13.2 W	22.3 W	46.1 W	110 W
4 K	0.84 W	0.01 W	0.36 W	0.73 W	1.94 W	2.0 W
1 K	5.01 mW	0.01 mW	6.46 mW	6.33 mW	19.82 mW	25.0 mW
100 mK	68.6 μ W	0.1 μ W	0.5 μ W	45.7 μ W	115.0 μ W	400 μ W

Table 4-12. Heat loads for the complete camera.

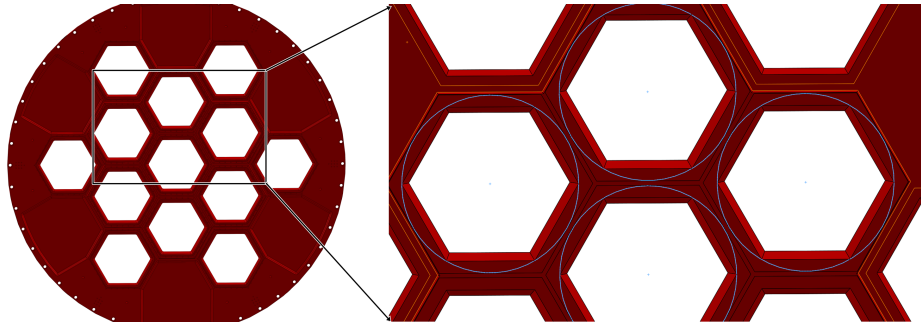


Figure 39. Front plate detail showing the material saved by using hexagonal windows. The gray circles are the size of the hole of equivalent minimum beam clearance at the outside surface of the front plate. Making the windows inscribed hexagons instead of the circles which circumscribe them adds a significant amount of material at the weakest point in the fplate.




Figure 40. Cross-section of the front plate showing tapering of the window holes to match the beam divergence. Light is entering the cryostat from the right, and the detector arrays are on the left.

Component	Mass (kg)	Notes
Vacuum shell	1310	
Refrigerators	≈ 200	
80-K stage	248	
40-K stage	451	
4-K stage	225	
Optics tube	77	times 13 (19) optics tubes
Total	3435 (3897)	with 13 (19) optics tubes

Table 4-13. Camera mass.

4.3.5.3 Stage supports

The 40-K and 4-K stages (labelled shields and plates in Fig. 36) are nested cylindrical cans, supported by rings of G10-CR tabs that provide high axial and lateral stiffness, with low thermal conductivity. The tabs accommodate differential thermal contraction between stages by flexing inwards [526]. The 1-K and 100-mK stages are OFHC copper bus bars supported by carbon fiber trusses [523]. The complete assembly of cold stages and optics tubes has a total mass of ≈ 2 T, which is much larger than in existing CMB cameras, but the large size of the stages results in a stiff support. The lowest natural frequency of the cold assembly, with 13 optics tubes installed, is ≈ 19 Hz. Stiffening the 4-K plate has the potential to significantly increase this frequency [526]. A lower natural frequency makes the camera more sensitive to vibrations, which typically couple to the detectors and readout electronics through magnetic fields, so the optics tubes have two layers of magnetic shielding.

4.3.5.4 Assembly procedure

Camera assembly requires a 4-m high bay, so the optics tubes can be inserted vertically. The assembly procedure begins with bolting together the two halves of the vacuum shell and mounting the shell in its cart. The 40-K radiation shield is assembled, wrapped in multi-layer insulation, and inserted vertically into the shell. The 4-K plate is lowered into place, the PT-420 coolers are installed, and then the 4-K radiation shield is installed. The 40-K and 80-K filter plates are bolted in place, the PT-90 coolers are installed, the optics tubes are inserted, and then the vacuum shell front and back plates are attached.

4.3.5.5 Vacuum system

The camera has a turbo pump attached to an ISO-150 gate valve on the back of the vacuum shell. The gate valve closes automatically in the event of a power loss. There are KF-50 ports on the front and back, for gauges and additional pumps, and pressure relief valves to protect the windows in case the vacuum shell becomes pressurized.

4.3.6 Control and monitoring

The cameras have a stand-alone control system that manages the vacuum and fridge systems, heaters, and temperature monitors. The control system is essentially a wrapper for various commercial units; it provides

local control and monitoring for testing purposes, and it accepts high-level commands from the observatory control system for automatic and remote operations. Temperatures, vacuum sensors, and the status of the various fridge and compressor controllers are sampled every few seconds and reported continuously, even when the camera is not observing.

4.3.7 Prototypes, assembly, and testing

The cameras are complicated systems, and while we can rely heavily on the work done by SO, the CMB-S4 and SO cameras are not identical, so prototypes will be needed. The reference design includes the following two levels of prototyping

1. A single optics tube, with detectors, in a small test cryostat that can be cycled quickly. The goal here is to check fit, function, and performance of everything in an optics tube. Tests will include operating temperatures and heat loads (using thermometers on key components), passbands (using a Fourier transform spectrometer), beam shapes (using a chopped thermal source on an xy stage), polarization errors (using a chopped thermal source with a rotating wire grid), and noise. Optical tests will require a cold neutral density filter, e.g., at the cold stop, or bolometers with a high- T_c TES that can handle room temperature loading.
2. A prototype of the final camera design, including a full-size vacuum shell with fridges and radiation shields, at least one optics tube, and heaters to simulate heat loads from missing wiring and windows. In this case, testing will focus on temperatures and heat loads.

The single-tube prototype cryostat will become a commissioning camera that will be used for checking telescope mechanical and software interfaces, and for measuring optical loading and telescope beam shapes. For this purpose, the single-tube prototype cryostat will be mounted in an adapter that simulates the mechanical interfaces of the full-size camera. Since the tests that will be done with the commissioning camera address telescope design issues, it will be sufficient to perform the tests on just one of the large telescopes.

The cooldown time for a complete camera is several weeks, so the Stage-3 practice of using the camera to test subsystems is impractical; the camera must be assembled with a working cryostat and fully tested optics tubes and detectors. Optics tubes will be tested in a suite of four small test cryostats, like the one used for testing the prototype optics tube. Each test cryostat will be supported by a small team that will run tests, analyze the results, and fix problems. Testing will likely be done at several different locations, but will be part of a single quality assurance process with common procedures, software, and documentation.

Each camera vacuum shell will initially be assembled with port cover plates, and leak checked. Radiation shields, fridges, thermometers and heaters, monitor and readout wiring, and windows will be installed, and the camera will be cooled to check cooling rates and base temperatures. After fixing any cryogenics issues, any and problems with thermometry and wiring, the optics tubes will be installed (with cold neutral density filters if the bolometers do not have a 2nd TES with high T_c) and the complete camera will be cooled. The tests that were done on each optics tube will be repeated, but for a subset of the detectors, with the goal of finding problems that appear only in the complete camera (e.g., excess noise when all the electronics are operating). Testing at this level will be part of the same quality assurance process used for the individual optics tubes. After testing, the camera will be warmed up and the parts will be prepared for shipping. Our intent is to ship the cameras essentially complete, to minimize re-assembly work on site.

The main constraint for assembling and testing the camera is the long cooldown time. A full thermal cycle will take a few weeks, so it is impractical to include many cycles in the plan.

4.4 Small telescopes and cameras

4.4.1 Design drivers

The ultra-deep survey, which targets the signature of primordial gravitational waves, drives the key measurement requirements for the small telescopes.

The measurement challenges for this survey, outlined in Sect. 2.3, differ from those of the deep and wide survey in ways that set unique design drivers for the small telescopes. To achieve the science goal of measuring r to the required precision ($\sigma(r) = 0.0005$ for a non-detection) ultimately demands measuring primordial B -mode polarization patterns with uncertainties of < 10 nK at angular scales of a few degrees on the sky. To do this requires an instrumental design that enables sufficient control of statistical noise fluctuations, galactic foregrounds, and instrumental systematics, all of which are made harder at degree scales by $1/f$ contributions to instrumental and atmospheric noise fluctuations and by the red spectrum of many foreground and systematic effects.

Small aperture telescopes offer intrinsic advantages for meeting the unique requirements of measuring r , which include:

1. **efficiency** to test, integrate, and deploy large arrays of detectors, due to their small size and large optical throughput;
2. **stability** of fully cryogenic telescopes, including their beams and internal systematics;
3. **calibrators** can be aperture-filling or easily placed in the far-field;
4. **modulators** can be aperture-filling and fully cryogenic;
5. **Mounts** can be compact and allow full boresight rotation;
6. **sidelobes** can be suppressed with unobstructed on-axis optics and superior baffling and shielding.

Additionally, small-aperture telescopes from the BICEP/Keck series at the South Pole have produced all the leading r constraints so far from ground-based B -mode measurements through Stages 1, 2, and 3. Small aperture telescopes, for example the ABS experiment, have also had the greatest success among telescopes in Chile in achieving sensitivity on degree scales, by using half-wave plate modulators to mitigate $1/f$ noise from the less stable atmosphere there. The intrinsic advantages of the small-aperture approach, combined with the experience that over the past decade, only small-aperture telescopes have demonstrated sufficient performance in degree-scale B -mode measurements to allow direct scaling that can meet the measurement requirements of CMB-S4, clearly dictate the use of small-aperture telescopes for the reference design.

The detailed design choices for the small-aperture telescopes are motivated by the requirements flowdown illustrated in Fig. 41. In considering changes for the reference design compared to the small-aperture telescopes that have achieved previous deep r measurements, we have incorporated new technologies—for example dichroic detectors, dilution refrigerators, and (if small-aperture telescopes are deployed to Chile) cryogenic half-wave plate modulators—where there is a consensus that they promise improved performance

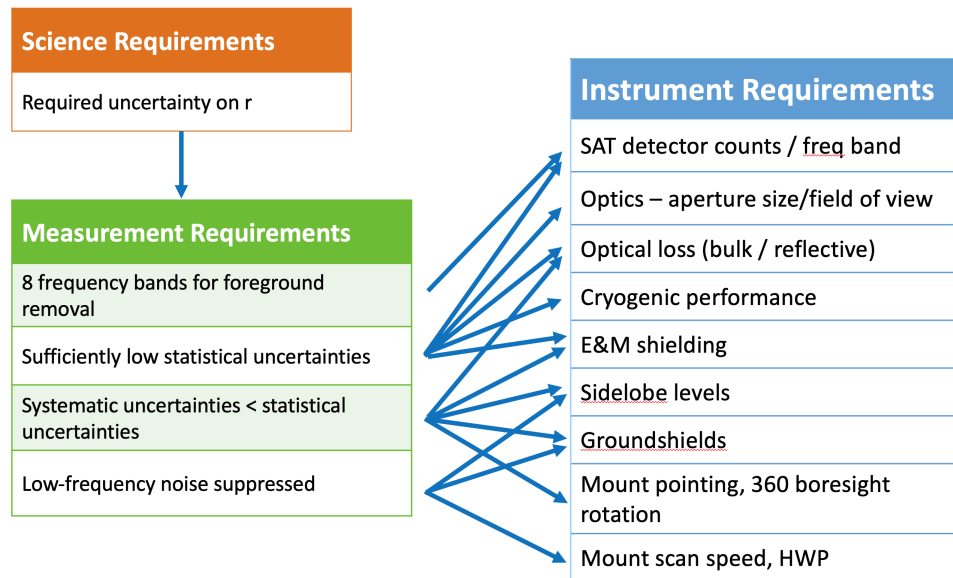


Figure 41. The small-telescopes instrument requirements, detailed in this subsection, are motivated by the unique measurement challenges set by measuring r to the required uncertainty using ultra-deep B -mode measurements at degree scales.

while adding little technical risk. In making design choices we have distinguished between *engineering issues*, those that can be fully developed and demonstrated in the lab to retire risk, and *science issues*, those whose impact on successfully meeting the measurement and science requirements must be judged with comparison to direct experience of making deep B -mode maps. For example, cryostat design is primarily an engineering issue because we are confident our design choices can be fully validated in the lab. Examples of science issues include beam and sidelobe optical performance, polarization modulation approach, ground pickup and shielding, and other systematic effects, and for design choices that impact these issues we have endeavored to stay close to and to build upon proven experience.

4.4.2 Reference design summary

Table 4-14 summarizes the optical design features and distribution of detectors between optics tubes. The lower frequency tubes use alumina lenses, while the highest frequency tubes use silicon lenses for their known low loss. We rely on proven experience from South Pole and Chile to guide our choice of polarization modulation technique for the small-aperture telescopes. For South Pole, the polarization signal is measured by differencing pairs of orthogonal detectors, modulated by scanning the telescope in azimuth. Any telescopes deployed to Chile at 85 GHz and above will be outfitted with continuously-rotating half-wave plates (HWP), described in Sect. 4.4.4, to suppress atmospheric $1/f$ noise and instrumental systematics. The atmosphere is sufficiently stable in the lowest frequency bands that a HWP is not needed.

Bands	Lenses	Field of view	Min. edge taper	Modulation (Pole/Chile)	Detectors / tube	Tubes
30 / 40	2× 55 cm Al	29°	−9.3 dB	scan	576	2
85 / 145	2× 55 cm Al	29°	−6.2 dB	scan / HWP	7048	6
95 / 155	2× 55 cm Al	29°	−8.4 dB	scan / HWP	7048	6
220 / 270	3× 44 cm Si	35°	−13.4 dB	scan / HWP	16876	4
total:					153,232 detectors, 18 tubes	

Table 4-14. Summary of small-aperture telescopes for the reference design.

4.4.3 Optics

The measurement requirements described above can be economically met with simple on-axis refractors. The 2-lens refractive optics of the BICEP series of experiments—BICEP1, BICEP2, Keck Array, BICEP3, and now BICEP Array—changed little over the past 14 years between the generations. The only notable upgrade was the increase in aperture/throughput that drives the change of the lens material from high-density polyethylene (HDPE) to alumina from BICEP3 onwards. The choice of on-axis refractive optics was re-affirmed by the SO collaboration who, after extensive comparisons between refractors with the leading alternative, crossed-Dragone reflectors, has also selected refractors for their compactness, lower cost, and easier sidelobe control.

The lenses must be anti-reflection (AR) coated to prevent reflections at percent or better levels. The SAT reference design baselines dichroic detectors which require similarly broadband cryogenic AR coatings on the optical elements. For silicon lenses, subwavelength features can be cut into the lens surface with a custom three axis silicon dicing saw. These features constitute a metamaterial that behaves as a simple dielectric coating with excellent optical and mechanical properties. For the alumina lenses, this can be done by gluing or laminating appropriate dielectric layers onto the lens surfaces. In particular, laser-treated epoxy coating has been applied successfully in BICEP3 (95 GHz) and POLARBEAR2 (95/150 GHz).

The SAT reference design will have 55-cm alumina optics at low frequencies and 44-cm silicon optics at high frequencies. This distribution was driven by the resolution requirements at low frequency, and uncertainties in loss and anti-reflection coating of alumina at high frequencies. These designs and aperture sizes are based on optics that have been fully demonstrated in BICEP3 and/or realized for the SO design.

An aperture size larger than $\approx 20(\lambda/2\text{mm})\text{cm}$ is required to resolve the degree-scale B modes. Beside the resolution requirements, a larger aperture offers larger optical throughput which proportionally reduces the cost of cryogenics and telescopes given the same number of detectors. On the other hand, the size of the vacuum window that can be reliably built and the availability of the lens material present practical limitations. Low-loss alumina lenses are available up to 80 cm in diameter. The Stage-3 experiment BICEP Array will adopt 55-cm alumina optics at all frequencies. For the CMB-S4 SAT reference design, silicon lenses are chosen for high frequency bands for two main reasons. First of all, high purity-silicon has superior transmittance compared to alumina. Additionally, broadband AR coating of silicon has been demonstrated and verified up to 240 GHz.

In addition to the lenses, each of these optics tubes requires a vacuum window and several thermal filters. The vacuum window consists of HDPE anti-reflection coated with layers of expanded polytetrafluoroethylene (PTFE) or HDPE. They can now be made with > 60 cm diameter and < 3 cm of thickness, which contribute only a small fraction of in-band optical loading. For a 55-cm aperture, more than 100W of thermal radiation




Figure 42. *The CMB-S4 small-telescope optics are simple cryogenic refractors of approximately half-meter aperture, a design based on heritage from the cryogenic refractors of the BICEP series of telescopes which have proven their performance in deep r measurements through Stages 1, 2, and 3. Shown here is the BICEP3 telescope, which has been observing since 2015, as an example of an existing instrument that illustrates the essential optics design elements (including lenses, cold baffles, filters, and vacuum window) that all the CMB-S4 small telescopes will contain.*

is entering the cryostat through each of the vacuum windows. This must be significantly reduced by thermal filtering. In Stage-3 experiments such as BICEP3, SPT-3G, and POLARBEAR2, adequate thermal filtering is achieved by a combination of scattering and absorptive filters. The scattering filters consist of stacks (≈ 10 layers) of thinly sliced foam sheets, sometimes known as the RT-MLI (radio-transparent multi-layer insulation). The remaining infrared loading can be absorbed by IR-opaque alumina filters, whose high thermal conductivity also makes them very effective heat sinks. In BICEP3, a foam stack reduces the incoming thermal loading by 90 W and an AR-coated alumina filter absorbs and dumps 12 W of the remaining power into the first stage (40 K) of the pulse tube cooler. In the BICEP series of experiments, additional nylon filters were used on the 4-K stage to reduce loading onto the sub-Kelvin stage, and low-pass edge filters were used directly above the focal plane to reduce out-of-band direct illumination of bolometer islands. Given the much larger cooling power of dilution refrigerators, and the bolometer island shielding offered by a feedhorn architecture, both of these final layers of filtering can likely be demonstrated in lab testing to be unnecessary, but to keep the reference design maximally conservative it retains these filters for now. Overall, the vacuum window and thermal filtering required by CMB-S4 SAT do not go beyond the level already demonstrated in BICEP3 and therefore present very low risks.

Another key decision made for the SAT reference design is related to the pixel density. As is well known, when the spillover illumination of the feeds was terminated at sufficiently low temperature (< 2 K), formal sensitivity optimization tends to favor feeds with diameters smaller than or comparable to the size of a diffractive spot ($f\lambda$), given a fixed optical throughput. However, the corresponding larger field strength at the edge of the aperture and the integrated spillover power terminated inside the optics tubes tend to generate near- and far-sidelobes with non-trivial polarization asymmetries through diffraction and/or small-

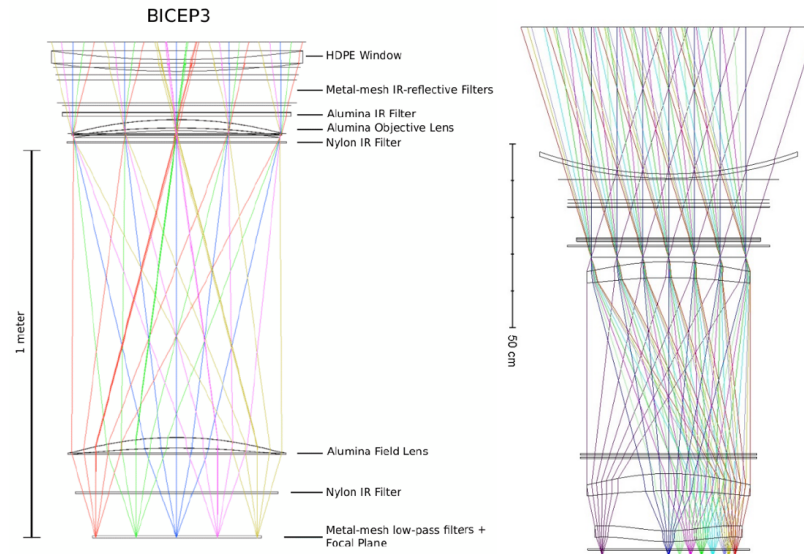


Figure 43. The detailed optical designs for the small telescopes are based on the realized 2-lens (BICEP, left panel), and 3-lens (SO, right panel) refractors that offer large optical throughput (*etendue*), symmetric main beams, and excellent polarization properties. Refractors are significantly more compact compared to crossed-Dragone telescopes with the same primary aperture, and offer advantages of symmetry and sidelobe control.

grazing angle reflections/scattering. This corresponds to more demanding requirements for the cold aperture stop and the absorptive baffles, and for the beam measurement and control of temperature to polarization leakage in analysis. The BICEP series of experiments have used conservative edge tapers that are in the -8 to -23 dB range. For the SAT reference design, we have relaxed this slightly to allow edge tapers as high as -5.7 dB (see Table 4-14), which are still relatively conservative. There are ongoing efforts to make the cold absorbers more mechanically robust under cryogenic conditions, and better matched in polarization response. However, with this conservative choice of moderate packing density for the feeds, the SAT cold stop, baffling, and systematics control and analysis strategies likely would not require anything beyond that has been demonstrated in Stage-3 experiments.

4.4.4 Half-wave plates

In the event that SATs are deployed to Chile, those at 85 GHz and above will be equipped with continuously-rotating cryogenic half-wave plates (HWP) to reduce the impact of the higher level of atmospheric noise fluctuations at that site. The HWP can also be used to eliminate the effects of instrumental polarization by optical elements between the HWP and the detectors. The HWP is located near the cold aperture, on the sky side of all the lenses, therefore eliminating the optical systematics due to the lenses and detectors. Reduction of instrumental polarization reduces temperature-to-polarization leakage, and the effects of polarized sidelobes generated by effects on the detector side of the HWP.

Our baseline HWP design is similar to that of the SO small-aperture telescopes, which is derived from the HWP implementation in the Simons Array receiver (Fig. 44). The optical element consists of three-layer stack of A-cut birefringent sapphire, with an anti-reflection coating similar to that of Alumina lenses. The




Figure 44. *Left: drawing of the SO cryogenic half-wave plate section. The HWP system consists of the sapphire HWP and rotation mechanism adopting superconducting mag-lev bearing. The optical clear aperture is 50 cm in diameter. Right: a photo of the Simons Array cryogenic half-wave plate system, from which the SO system is derived.*

three-layer stack allows a wide-enough modulation bandwidth for the dichroic detectors. The HWP is located on the 45-K stage. Since the sapphire loss-tangent rapidly decreases as a function of temperature, the thermal emission from the HWP is negligible below 100 K. The rotation mechanism consists of superconducting mag-lev bearing and electromagnetic drive, achieving non-contact rotation. This allows smooth rotation and long-lifetime operation in the vacuum and cryogenic environment.

4.4.5 Cryostat

4.4.5.1 Cryostat design drivers

The cryostat consists of lenses and an aperture stop that are cooled to ≈ 1 K, and a focal plane maintained at ≈ 100 mK by means of a dilution refrigerator. Maintaining these temperatures rather than higher ones offers sensitivity gains, and fabrication margin for the superconducting detectors.

We considered two alternative designs for the cryostat: (1) a cryostat with a single optics tube and focal plane assembly; and (2) a cryostat that contains three optics tubes and focal plane assemblies. We refer to these as the “single-tube” and “3-tube” systems. In both cases, three tubes share a common telescope mount. Our choice was based on the following design drivers.

Electrical Power: The system of pulse-tube and dilution refrigerator coolers is the largest power consumer during operation, and supply of power is among the most expensive logistics requirements at either site. In particular, limitation in power at the South Pole may be a driving constraint. The power consumption per unit of sensitivity is among the primary figures of merit. The 3-tube cryostat allows more efficient use of power, and thus has a better power consumption figure of merit.

Modularity: The modularity of the system is a key element during the integration and testing. While a single-tube system allows more modular integration and deployment, the total number of units to deploy is smaller with the 3-tube design, and they still offer a large degree of modularity compared with the LAT cameras.

Risk: Commissioning the cryogenic systems is one of the largest schedule risks. When optimizing the system for power consumption, the 3-tube system has larger margin in cooling capacity, and thus has less risk




Figure 45. On the left, the SO cryostat, illustrating a design that uses a dilution refrigerator to cool focal plane and 1-K optics. On the right, the reference design “3-tube” cryostat. It accommodates three cryogenic optics tubes each fitting within a 70-cm diameter \times 150-cm long envelope, cooled to 4 K by the three pulsetubes located around the lower perimeter of the cryostat and to 1 K and (for the focal planes) 100 mK by the dilution refrigerator. The optics tube shown is the low-frequency, alumina lens version, derived from the BICEP Array design. The high-frequency, silicon lens version will fit in the same volume, with bolt-on extensions between the cryostat and each vacuum window to accommodate the HWPs if deployed to Chile.

in cryogenics. On the other hand, all the pulse-tube-cooled small-aperture systems deployed in the past (BICEP series, ABS, and CLASS) are single-tube cryostats, and this heritage translates to reduced risk.

Baffling: A single-tube system must have larger separation between the three cryostats’ vacuum windows (on a single telescope mount), allowing each tube to have a dedicated forebaffle centered on that tube. These centered, symmetric forebaffles are advantageous for the single-tube system. A 3-tube system would have a common and proportionally larger forebaffle, leading to higher cost and a construction challenge. On the other hand, since the separation among the vacuum windows in a 3-tube system is smaller, the ground screen would also be smaller.

Cost: A 3-tube cryostat will be operated with a single dilution refrigerator offering cost-savings relative to the single-tube option, which requires three times more refrigerators.

Based on these design drivers, and specifically because of lower power consumption and cryogenic risk the reference design uses the 3-tube system. However, we maintain the single-tube system as an option. Among the near-term development items is the investigation of the cryogenics significantly contributing to this *engineering* decision in the future.

Among the various multi-tube cryostat options, the *three-tube* was judged optimal because the 1-K stage and 4-K stage cooling power of the BlueFors SD400 and Cryomech PT420 are well-matched to support three optics tubes. Increasing the number of close-packed optics tubes per cryostat beyond three likely requires additional cryocooling, and will also lead to a larger common forebaffle. For these reasons we limited the number of tubes per cryostat to three.

Fridge	Type	Quantity	Power	Cooling capacity per fridge			
				45 K	4 K	1 K	100 mK
PT60	Pulse tube	1	3.3 kW	20 W			
PT420	Pulse tube	1	12.5 kW	55 W	2 W		
PT410	Pulse tube	1	8.4 kW	40 W	1 W		
SD400	Dilution	1	2.0 kW			25 mW	400 μ W

Table 4-15. Refrigerators. The PT410 is only used for the dilution refrigerator system.

4.4.5.2 General configuration

The general configuration of a SAT receiver will consist of a 2-meter cryostat with three optical tubes. The main shell will have three windows and a common 40 K radiative shield. Inside this shield, we will place three receivers that will have separated 4 K shields and optics. The size of the single optical tube will be around 70 cm, comparable with the present Stage-3 small-aperture instruments. This design will allow different teams to work in parallel on optics tubes that fill a single cryostat design, which would be identical for the South Pole and Chile small telescopes. A key factor of the optical tube design will be the modularity of the mechanical, optical and readout parts. That will allow a fast and straightforward replacement of parts and also the capability to reconfigure the instrument observation bands to enhance multi-frequency coverage and map depth. Standardization of the parts that involve cabling, mechanical supports, and detector module interfaces has demonstrated advantages of cost savings and efficiency and is standard practice within in Stage 2 and Stage 3 experiments (e.g., Keck Array, BICEP Array, SO) that have deployed or are planning to deploy small-aperture telescopes at multiple frequencies.

4.4.5.3 Cryogenics

The detectors are cooled to ≈ 100 mK, with intermediate stages at 1, 4, and 45 K for radiation shields, optical filters, and wiring thermal intercepts. The 45-K and 4-K stages have one PT420 pulse tube supplemented (to add margin) by one single-stage PT60 pulse tube (see Table 4.4.5.3). The 100-mK stage is cooled by a BlueFors SD400 dilution refrigerator with an intermediate stage at 1 K. The dilution fridge has its own PT410, dedicated for the dilution fridge system. Flexible copper braids are used to connect the refrigerators to the cryostat cold stages. All the refrigerators must be kept within $\approx 45^\circ$ of vertical to maintain cooling capacity.

The reference design cryogenic architecture for a SAT receiver includes three optical tubes inside a single 40 K shield volume, cooled by the PT420 and PT60 pulse tube first stages. The PT420 pulse tube second stage will cool the three tubes down to 4 K. The PT410 will provide the base temperature to run the dilution refrigerator for cooling the insert sub-kelvin stages and the focal plane to an operating temperature of 1 K and 100 mK, respectively. The required cooling power for the SAT can be extrapolated considering similar configuration experiment design and the BICEP Array receiver. We are expecting to have roughly 50 W at 40 K and 1.5 W at 4 K that can easily be managed by the combination of Cryomech PT420 and PT60. The dilution refrigerator will have stages at 1 K and 100 mK respectively. Again comparing with similar instrument designs we are expecting around 0.5 mW at 1 K and 45 μ W at 100 mK, well within the capabilities of a BlueFors SD400 Dilution Refrigerator.

In order to mitigate input thermal loads, the 40-K and 4-K radiative shields have to be wrapped with Multi-Layer Insulation aluminized mylar. Moreover, the mechanical supports between the stages have stringent requirements in terms of their thermal conductivity. The current generation of cryogenic instruments for

CMB uses G10 fiberglass supports (sometimes supplemented by titanium tension members) above 30 K and the composite carbon fiber ones below 30 K. These solutions will be implemented in the SAT design. The choice of materials for the shields and the thermal straps is crucial. Based on past experience, we will largely use high-conductivity 1100 Aluminum for the radiative shields and 101 Oxygen-free Copper for cold plates and thermal straps. For the design of the heat straps, depending on the temperature stage, we will use copper braids or foil shims. Their design maximizes the thermal conductivity allowing the differential thermal contraction due to the cooling process to take place without damaging the cryogenic system.

Superconducting devices such as TES bolometers and KIDs are sensitive to the presence of external magnetic fields. Also, the readout architecture for TES bolometers is based on the use of Superconducting Quantum Interference Devices (SQUIDS), which are very sensitive magnetometers. These millimeter-wave sensors and the cold readout electronics are therefore quite sensitive to electromagnetic pickup induced by even very weak magnetic fields. The Earth's magnetic field (modulated by the telescope scan), as well as fields produced by the instrumentation surrounding the experiment, are potential sources. This pickup can result in the production of artifacts in the CMB temperature and polarization maps that can be very hard to disentangle and remove. A common strategy for mitigating these effects is to implement magnetic shields in the cryostat mechanical design, a technique that is currently adopted in all the Stage-3 CMB experiments such as BICEP Array, SPT-3G, and Advanced ACT. The shielding strategy varies from experiment to experiment but it is based on the interleaving of high permeability and superconductive material layers. The baseline design for the SAT magnetic shield is assumed to copy that of BICEP Array, which uses a high-permeability (Amunéal A4K) cylinder around the 40-K tube combined with a Niobium superconductive cup at 300 mK (or at 1 K for S4), to reduce magnetic fields by a factor of about 200 at the focal plane. Additional shielding layers are usually added locally in the detectors/SQUID areas. Looking again at the BICEP Array architecture, its FPU detector module design provides an additional factor of 500 shielding with the use of an A4K planar sheet and a Niobium enclosure; a similar shielding factor at the detector modules is assumed to be achievable for the S4 reference design. The use of finite-element model simulation software (e.g., COMSOL Multiphysics) can help optimize the shielding configuration.

4.4.5.4 Mechanics

Survival requirements on the vacuum jacket and vacuum window include: safe storage to temperatures of -90 C and maintenance of vacuum to temperatures of -40 C . In past experiments the latter has been straightforward to achieve with silicone elastomeric o-rings. In operation the cryostats will be contained in an environmental enclosure which surrounds the telescope mount, with heated airflow directed around the vacuum window to eliminate frost/snow accumulation, so exposure of cryostat components to colder ambient temperatures while in operation is not expected.

Within the cryostat, pointing rigidity of the optical tubes is maintained by arranging for kinematic constraint of the thermal standoffs, with rigid triangulated G10/carbon fiber truss elements supporting bottom end of the 50-K, 4-K, and 1-K stages and Ti-Al-4V tensile members providing radial constraint at the top (open) end of each tube [527]. Lenses are supported from optics tubes using tangential-blade standoffs which provide thermal conductance while maintaining concentricity and relieving radial strain from differential thermal contraction. Based on pointing performance of Keck Array and BICEP3, we expect pointing effects from gravitational deflections of the optical elements within the cryostat to be easily limited to < 15 arcsec with such a mechanical support system, and to be repeatable. The cryostats themselves are supported in the telescope mount with rigid triangulated space frame members which connect hard points near the top and bottom vacuum flanges to mounting points on the inner diameter of the boresight rotation stage that surround the cryostat's center of gravity.




Figure 46. *The CMB-S4 small-aperture telescope reference design telescope mount is based on the existing BICEP Array mount design (left). Although originally designed for four individual small-aperture telescope cryostats (each with a single pulsetube), its size is well-suited to mount the single 3-tube cryostat with three pulsetubes and dilution refrigerator (right). It allows three axis motion with infinite rotation (using rotary feedthroughs) on the azimuth and boresight axes, and a design that allows for full environmental enclosure with space for interior access to all critical components.*

The total cryostat mass is estimated to be 1800 kg (approximately three times that of a BICEP Array cryostat). SAT cryostat assembly (both in North America and in the field) will take place on dedicated wheeled lab stands designed to distribute their mass over a 2×2 m footprint, minimizing requirements for specially-reinforced lab floors. Assembly of the cryostats requires a well-controlled 2-ton hoist with a hook height of ≥ 12.5 feet.

4.4.5.5 Control and monitoring

The small-aperture telescope cryostats share the same control and monitoring interface as the LAT cameras (described briefly in Sect. 4.2.6). We monitor temperature sensors, vacuum gauge and valve, and the status of the refrigerator and compressor. Control interfaces are used for heaters, the vacuum valve, refrigerator, and compressor.

Control and monitoring for each SAT integrates into the overall observatory control software in the form of one or more subsystems, as outline in Sect. 4.6.1.2. The requirements for timing synchronization are modest and the monitoring data rate is insignificant in comparison to the bolometer data.

4.4.6 Mount

The CMB-S4 reference design mount for the small-aperture telescopes is a 3-axis fast-scanning mount, based on the existing design of the (Stage 3) BICEP Array mount (see Fig. 46 and Ref. [527]).

Parameter	Value	Notes
Mass of instrument	up to 4500 kg	includes cryostat, DR system, electronics, forebaffle
Motion	3 axis	full boresight rotation of instrument and forebaffle
Scan pointing knowledge	< 15 arcsec rms	< 1/20th beamwidth at $\lambda = 1$ mm
Scan speed AZ/EL/TH	5/1/1 deg s ⁻¹	≈ 3 deg s ⁻¹ on the sky for fast diff. measurements
Scan accel. AZ/EL/TH	3/1/1 deg s ⁻²	turnaround efficiency
Range AZ/EL/TH	$\infty/45 \dots 110/\infty$	continuous AZ desirable
Shipping envelope	standard double pallet	deployment via C-130 / standard vehicles
Mount mass	< 25 tons	includes instrument, comoving forebaffle and scoop
Survival: wind	70 m s ⁻¹	Chile dominates
Survival: seismic	0.3 g	Chile dominates
Survival: temperature	-90 C	Pole dominates

Table 4-16. *Small-telescope mount design requirements.*

The CMB-S4 mount must be capable of carrying the SAT cryostats and their associated electronics and cryogenic equipment and of scanning them on the sky rapidly enough to make differential degree-scale measurements in the presence of $1/f$ instrumental noise and changing atmosphere and ground signals. It must achieve an accuracy of pointing knowledge while scanning that ensures sufficient control of systematic uncertainties, both in mixing of E to B modes due to pointing reconstruction errors [506] and in cross-referencing maps across SAT frequencies and with the LAT data for accurate galactic foreground and lensing separation. It must provide 3-axis rotation, including full rotation around the boresight axis to allow both for suppression of instrumental systematics in polarization measurement and for consistency checks offered by a full set of null tests [528]. It must provide an environmental enclosure to protect the cryostats, electronics, and other elements from extremes of temperature or snow exposure and to offer a stable operating environment with reliable interior access to all critical components. The design requirements are summarized in Table 4-16

The reference design mount illustrated in Fig. 46 meets the design requirements using the existing BICEP Array mount design. Although this mount was originally designed for four individual small-aperture telescope cryostats, each with a single pulsetube, its size is well-suited to mount the single 3-tube cryostat with three pulsetubes and dilution refrigerator. This mount includes two separate rotary unions which allow continuous rotation about the azimuth axis and the array boresight axis without the need for cable wraps. These rotary unions each contain 10 helium channels. Eight of these connect the pulse tubes and their compressors, while two channels serve as pressure guards. An additional nitrogen channel provides a pressurized environment on front end of the cryostats which prevents water absorption into the window material. Slip rings at the ends of the unions additionally provide data and power connections to electronics across separately rotating stages of the mount. These rotary unions allow the helium compressors required to operate the pulse tube coolers to sit below the mount structure in the stationary equipment room. Helium lines route upwards into the lower (ground fixed) half of the first rotary union and then out through the upper half which rotates in azimuth along with the receivers. The hoses from the upper half are then routed through a short cable chain that provides flexure when rotating in elevation. The second rotary union is then similarly connected between the elevation and boresight stages. Modifications to the current mount design will be required mount the dilution refrigerator gas handling system inside the envelope of the instrument package which co-rotates on the boresight stage, and engineering work will be required to ensure this system can operate while tilting up to 45° from vertical.




Figure 47. Exterior view of the small-telescope mount and shielding. The 3-tube cryostat is surrounded by a co-moving absorptive forebaffle. The mount sits within a large reflective ground shield. An additional co-rotating “scoop” is used to keep the shield dimensions reasonable. A flexible environmental seal surrounds the mount structure.

The mount provides a flexible environmental seal that fully encloses the components seen in Fig. 46, while exposing only the vacuum window, forebaffle, and co-rotating shields (see Fig. 47). Access to the telescope for service is from an equipment room below which houses helium compressors and control computers, similar to existing BICEP and Keck Array facilities currently in use at the South Pole. The access passage through the azimuth bearing is large in diameter, and is designed to accommodate installation of cryostats and other components as well as personnel access, all within the interior environmental space of the observatory complex. Alternatively, cryostats may be lifted into the mount from above using a crane.

4.4.6.1 Control and monitoring

The large and small telescopes share a common control and monitoring interface (see Sect. 4.2). Drive amplifiers, brushless servo motors, reducers and drives, encoders, emergency stop and safety interlock systems are specific to each telescope design. Based on the systems developed for the Keck Array and BICEP Array mounts, we expect these drives to consume 10 kW peak power.

4.4.7 Ground shields and exterior baffles

Extended beam response can couple to the warm ground or the Galaxy/Sun/Moon during CMB scans, potentially adding structure at the degree scales relevant for the r measurement. A key advantage of the small-aperture approach is that the entire telescope can be enclosed in multiple levels of shielding, significantly reducing the effect of sidelobes. Indeed, all of the leading r constraints to date have used similar shielding strategies to absorb sidelobes and/or reflect them to the cold sky.

The first level of baffling is a co-moving absorptive cylinder which extends around the telescope aperture. Sidelobes—usually at $\gtrsim 10^\circ$ scales—that terminate on this surface see a constant-temperature load, and

therefore do not contribute to degree-scale structure in the map (although the extra loading slightly decreases sensitivity). Like in the BICEP experiments, we will also use an additional fixed, reflective ground shield that surrounds the telescope mount. While no rays from the aperture couple directly to the ground shield, those that diffract over the lip of the co-moving forebaffle do; the ground shield ensures that rays must diffract twice before terminating on the warm ground. A ray that couples to the ground shield can add azimuth-fixed structure to the map, which must then be filtered from the timestreams. Finally, to keep the ground shield dimensions reasonable, we use a tertiary shield that co-rotates in azimuth with the mount (but is fixed in elevation), which partially surrounds the telescope to couple to sidelobes at the lowest elevations. Figure 47 illustrates the three levels of shielding around the 3-tube cryostat.

The dominant sources of sidelobes in small-aperture telescopes will be different from those in the large-aperture case. Instead of mirror scattering and panel gaps, the most important effects for the SATs will be: diffraction at the cold stop and primary aperture; scattering from the window, filters, and other optical elements near the aperture; and non-sequential reflections between optical elements. We expect to model the beams using Zemax and GRASP simulations to demonstrate that we can meet the stringent systematics requirements. These models will then be confronted with *in situ* measurements, both of the sidelobes themselves and of the amplitude and temporal stability of the ground pickup signals on the relevant angular scales, taken with and without various levels of shielding.

4.4.8 Assembly and test plan

Because integrating a single Small Aperture Telescope optics tube, focal plane (FP) and cryostat is expected to take on the order of 6–9 months and because there are six SAT 3-tube Systems totaling 18 tubes, the plan is to create three independent integration and test (I&T) facilities. In addition to documenting and reviewing assembly/test plans for the focal planes and optics/cryostat assemblies, we will be defining the necessary infrastructure requirements for these I&T sites.

After the sites are chosen, SAT specific modifications will have to be completed. It is expected that clean rooms will be required for focal-plane assembly. In addition, vacuum and cryogenic related systems will be required for cold testing. Finally, high bay space with adequate crane coverage and a hook height of at least 12.5 feet is required for handling the large cryostats and assembled optics tubes. Each of these facilities will be occupied for a period of approximately 2 years including time for testing and crating parts and assemblies.

The first of these three SAT I&T sites will be utilized to initiate the assembly/test process. The first focal plane (FP) integration and test process will be used to capture necessary improvements and refine the FP I&T procedures prior to passing them on to the other two sites. The optics/cryostat integration procedures will be tested and refined during a careful first integration of the optics and cryostat. Again, test procedures will be exercised and refined during pumping, cool-down and trial operation of the first cryostat. All these refined procedures and lessons learned will be utilized in integrating the remaining SAT systems.

Should throughput for the FPs be sufficient at one or two facilities, it would be possible to reduce the number of facilities for this portion of the I&T scope. However, to match the fabrication schedule for the remaining parts, we believe three sites are necessary. This plan allows all SAT systems to be integrated, tested and crated within somewhat less than 2 years. It also allows for contingency time in the event of problems arising during testing of the FPs or cryostat/optics assemblies.

To summarize the integration process, the FP hardware and detectors are assembled in a clean room and cold tested to ensure a high percentage of detectors are working. The tested FP assemblies are then delivered to the optics tube assembly area where they are temporarily installed on the back of the tube. Three-

tube structures are then installed in the integrated cryostats (with pulse tubes and vacuum equipment). The integrated cryostat is then pumped and cooled to operating temperature which is expected to take approximately 3 weeks time. A quick functionality test (to be defined) is performed to ensure proper operation. If problems arise, the system must be allowed time to come back to room temp so that the cryostat can be opened and the necessary repairs made. This could extend the test time by 2–3 months.

Once all systems are verified to be operating correctly, the cryostat will be opened to allow the FP to be removed and proper shipping supports installed before resealing and installing in the shipping crate. Once at Site (Pole or Chile), time will be required to remove the cryostat assemblies from their crates, open and remove shipping blocks, install cryogenic supports, re-install the FPs and button everything back up for a final cool-down and functionality test. This is expected to take up to 2 months time at site.

4.5 Detectors and readout

4.5.1 General description

The detectors and readout for CMB-S4 include the systems that couple RF from telescope at the focal plane, define the frequency bands, detect the RF power, and readout the resulting signals to room temperature.

CMB-S4 will use feedhorn-coupled transition-edge sensor (TES) bolometers as the focal plane detector technology. The pixels will employ a dichroic architecture with good performance over a 2.3:1 optical bandwidth. This dichroic scheme allows a single pixel to simultaneously observe in two distinct optical passbands, efficiently utilizing focal plane real estate without substantially burdening readout or optical coating technologies.

The detectors will couple to the telescopes through monolithic arrays of machined spline-profiled feedhorns. Feedhorns are a well developed technology providing high-efficiency telescope coupling with circular beams and low side-lobes. The spline-profiled shape provides the required 2.3:1 optical bandwidth with minimal crosspolar leakage. The waveguide output of the horn couples to a broadband planar orthomode transducer (OMT), which separates the orthogonal polarization components of the incoming radiation onto two pairs of transmission lines. The superconducting transmission lines feed diplexer filter banks that channelize the signal into the two optical passbands. For CMB-S4, the bandwidth of each optical passband may be as low as 15–20%. Though potentially smaller than currently fielded systems, this bandwidth can still be achieved with demonstrated technologies. The diplexer output feeds a 180° hybrid tee, which eliminates coupling to higher order modes in the waveguide providing single-moded performance over the full pixel optical bandwidth. The difference output of the hybrid tee is transmitted to a TES bolometer where the optical signal is thermalized and measured. Summing the outputs of a bolometer pair within each pixel measures the total intensity of the incoming radiation, and differencing the two measures the Stokes Q parameter. The full linear polarization of a given location on the sky is measured by combining measurements of multiple polarimeters rotated relative to each other or by a continuously rotating half-wave plate.

The optical signal is measured using a transition-edge sensor (TES) bolometer, a highly sensitive thermometer consisting of a thin superconducting film weakly heat-sunk to a bath temperature much lower than the superconductor T_c . TES dynamics are well understood, providing good descriptions of the noise and response for real devices. For CMB-S4, the TES detectors are designed so that the sensitivity is dominated by the (irreducible) statistical fluctuations of the absorbed photons. For ground-based experiments, this noise is typically $O(10)$ aW/ $\sqrt{\text{Hz}}$, though values vary depending on platform/site, observation frequency/bandwidth, and the instrumental throughput/efficiency. Apart from photon noise, TES bolometers also measure




Figure 48. Schematic diagram of a horn-coupled focal plane [529].

fluctuations in the thermal carriers of the TES weak thermal link. The CMB-S4 TES bolometers will be designed and fabricated with parameters tailored for individual passbands at each site. These parameters will be chosen to provide the detectors with sufficient dynamic range to accommodate variations in weather and to provide sufficient thermal isolation so that the thermal fluctuation noise is much less than the photon noise. This optimization process is well developed from Stage-2 and Stage-3 experiments. Together with sufficiently low noise readout electronics, the TES detectors for CMB-S4 will operate with nearly “background limited” sensitivities. Despite the nearly background limited TES detector performance, over 500,000 detectors are required to achieve sufficient sensitivity to realize the CMB-S4 science targets. The CMB-S4 detector count is 5–10× larger than focal planes currently being developed, which will make it the most sensitive ground-based CMB experiment in the world.

The TES bolometers for CMB-S4 will have two superconducting transition sensors in series, each with different superconducting transition temperatures. This design allows the TES bolometers to have two modes. One mode uses a sensor with higher superconducting transition, allowing the TES bolometers to operate under a high optical load from calibration equipment. The second mode uses a sensor with a lower superconducting transition temperature. In this mode, the dynamic range of the detectors are tuned to achieve background limited sensitivity for astronomical observation. In this observation mode, a sensor used for the calibration source does not interfere with observation as that sensor is superconducting with zero resistance. This technique has been implemented successfully across multiple deployed CMB experiments. Aluminum is used for sensors with higher superconducting transition temperature (approximately 1.2 K) and manganese doped aluminum will be used for observation sensors (0.16 K).

The TESs will be read out using superconducting quantum interference devices (SQUIDs). SQUIDs have a large noise margin over the detector noise, a critical feature for operating large arrays of detectors at sub-kelvin temperatures. In this Reference Design, we use time-division multiplexing (TDM), where a group of detectors is arranged into a two-dimensional logical array. Each column of detectors shares a dedicated readout amplifier chain, and only one row of the array is routed to the amplifiers at any given time. The various rows are addressed cyclically in rapid succession to record the entire array. In TDM, the first-stage SQUID is wired in parallel with a Josephson junction switch, and the series voltage sum of all such units in the column is amplified by a series SQUID array (SSA) for transmission to the warm electronics. During

multiplexing all but one of the switches are closed to short out the inactive SQUIDs, so that only a single first-stage SQUID feeds the SSA at any given time. The bandwidth of each pixel is limited to below the Nyquist frequency of the sampling, so that the signal in each pixel can be faithfully demultiplexed.

The focal plane multiplexing readout will be composed of three components: a DC wafer, multiplexing chips, and a critical line (CL) wafer. The DC wafer will have inductors and resistors for biasing the detectors, but also function as a wire-routing wafer to connect the detector signal lines through the inductors/resistors to the SQUIDs and to connect the addressing lines. The multiplexing chips will be mounted to the DC wafer and connected by superconducting aluminum wires. The CL wafer will have pockets to get the bias lines out from the DC wafer and the addressing lines from the mux chips. It will also serve as a platform for the interconnects to get the signals from the focal plane to the warmer temperature stages.

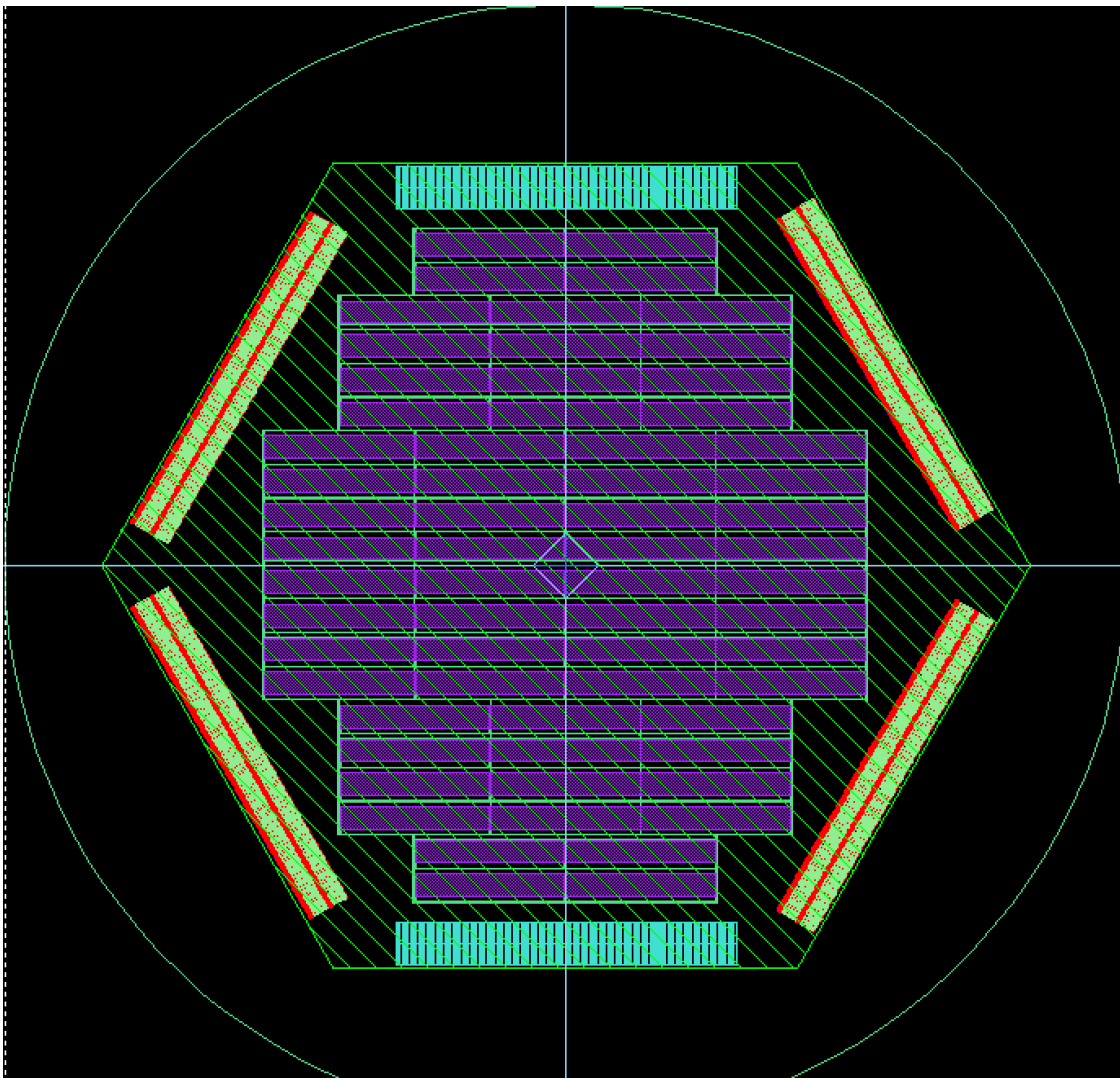


Figure 49. The figure represents a DC wafer with 32 multiplexing chips for allowing 2048-detector readout capability. The addressing lines will be routed out from top or bottom of this wafer and the bias lines will be routed out from the left/right vertices.

The CMB-S4 detector arrays will be assembled into closely packed universal focal-plane modules (UFMs). The UFMs will provide a universal cryomechanical interfaces so they can be used in both the large and small-aperture CMB-S4 telescopes, but there will be different versions of the UFM for different detector wiring at different frequencies, etc. Each UFM will carry a 150-mm detector array, a feedhorn array and associated optical coupling parts, and the cryogenic components for multiplexing the 1000+ detectors, in a 20–40 mm-thick hexagonal package, depending on the frequency. Each UFM will mount to a focal plane base (FPB) plate via a flange which provides an optical reference surface (i.e., the flange is at a well-defined distance below the phase centers of the feedhorn array). Flanges from neighboring UFMs interlock (reference SPT-3G and PB2) to allow closer packing between arrays, given edge-to-edge distance of 125 mm.

The number of detectors on one wafer varies depending on the frequency, and on the horn size which is different for the LAT and SAT. Thus a variety of different design layouts are needed for the cryogenic multiplexing components. Even more designs are needed for detector wafers, as the thermal conductances will also be a function of the size of the telescope, as well as the site.

Tables that list the band center frequencies, the number of wafers per optics tube, the number of TESes per wafer, the number of SATs and LATs, the number of wafers, and the total number of detectors are shown in Tables 3-1 and 3-2. Note that the focal planes as designed include partial-wafer sub-arrays, some shaped as half-hexagonal arrays, others as rhomb-shape sub-arrays. The TESs on three rhombs add to the same number as a hexagonal wafer, and, similarly, the TESs on two half-hexagons add to the same number as a hexagonal wafer. The wafer counts in the tables referenced above add the partial wafer detector counts into an equivalent full wafer count.

4.5.2 Detectors

Detectors for CMB-S4 will be fabricated on silicon wafers using well tested micro-fabrication techniques built on the foundation of processes that were established for Stage-2 and Stage-3 experiments. These fabrication processes will be optimized so that > 90% of the detectors on an array are operational (i.e. can be electrically connected, optically coupled, and biased).

Detector array layout: The CMB-S4 detector arrays will be fabricated on 150-mm diameter wafers. A combination of hexagonal, half-hexagon, and rhomboidal tiles that comprise 1/3 of a hexagonal wafer will be cut from the 150-mm diameter wafers such that they can be assembled in a close-packed form to make a large focal plane as shown in Fig. 50.

Pixels are arranged on a hexagonal tile to maximize packing density and yield. The pixel count per hexagonal tile will range from approximately fifty to five hundred, depending on the observation frequency of the detector array. Higher frequency bands will have more pixels per array. Each pixel consists of an OMT to optically couple to the pixel, radio-frequency (RF) circuits to control the RF signal (band pass filters, mode-selector and cross-under), TES bolometers to measure the incident signal, and microwave transmission lines that connect these elements as shown in Fig. 51. Each pixel on the wafer will be sensitive to two frequency bands and two orthogonal linear polarizations. Thus, each pixel will contain four TES bolometers. Each wafer will have approximately 10% additional TES bolometers that are not connected to optical coupling element to monitor detector sensitivity to the environment. In total, there will be approximately a few hundred to several thousand TES bolometers on a hexagonal tile. The electrical connection to the detector wafer is made with rows of wire bond pads placed at the perimeter of the hexagonal detector tile.




Figure 50. (Left) Photograph of an Advanced ACTPol detector array fabricated by NIST on a 150-mm diameter silicon wafer. (Right) Photograph of a POLARBEAR-2 focal plane. The diameter of the focal plane is approximately 400 mm. The focal planes for the small aperture telescopes will be constructed from combinations of hexagonal, half-hexagon, and rhomb sub-arrays that vary with the frequency band. A hexagonal focal plane for the large aperture telescopes will be constructed from three close-packed hexagonal wafers with three smaller rhombuses at the edges. The three smaller rhombuses will be fabricated on a single wafer.




Figure 51. Optical microscope image of an Advanced ACTPol detector pixel highlighting several of the key components. Magnified images of the major pixel components include: (a) the planar orthomode transducer; (b) the coplanar waveguide to microstrip transmission line; (c) the band-defining in-line stub filters; (d) the 180 degree hybrid tee; and (e) one of the AlMn TESs.

Fabrication plan: The detector fabrication processes will be built on the foundation of detector fabrication processes established during Stage-3 experiments. The baseline plan utilizes teams of two fabrication engineers carrying batches of 150-mm detector wafers through the fabrication process with multiple, standardized quality-control measurements.

The detectors are fabricated using several cycles of material deposition, patterning, and etching to build up a complete device. Superconducting materials (niobium, aluminum and manganese-doped aluminum) and normal metals (gold and palladium) will be sputter- or e-beam evaporation deposited. Dielectric films will be deposited with plasma enhanced chemical vapor deposition. Features will be defined using stepped lithography for frontside patterning and contact-alignment lithography for backside patterning. Chemical wet etch, plasma etching and ion-milling will be used to define lithographed features. A deep reactive silicon etch will be used to release membranes for the OMT and the TES bolometer thermal isolation.

The detectors will be produced in nanofabrication facilities with class 100 clean rooms based at several national laboratories. These facilities have an extensive heritage in superconducting transition-edge sensor fabrication, having fabricated multiple Stage-3 CMB instruments, SuperCDMS, X-ray spectrometer instruments, etc. Photos of a few of these facilities and examples of fabricated detectors are shown in Fig. 52.




Figure 52. Photographs of nanofabrication facilities and devices fabricated at: (Left) ANL (Center) LBNL/ U.C. Berkeley with (Right) SLAC

Fabrication quality control: One critical challenge for detector fabrication is to develop robust fabrication procedures that guard against process deviation. To address this, CMB fabrication facilities have been tracking processes throughout fabrication and ensure that the fabrication tools receive regular maintenance. All fabrication and metrology steps will be tracked in a computer-based logging system. The status of the fabrication tools and clean room environment, including chamber base pressure, plate temperature, clean room temperature, and humidity will be recorded each time a tool is used. The data will be continually compared to historical norms to identify deviations from process stability. This monitoring process has been implemented successfully in detector fabrication for Stage-3 experiments. While the tooling may differ between fabrication facilities, all detector arrays will be required to meet the same performance specifications, ensuring uniformity across all detector arrays.

Each detector wafer and detector pixel design includes test structures that can be used to monitor the fabrication processes and material characteristics. Sacrificial test wafers will also be carried through some process steps to monitor process steps. These methods have been proven to be useful in producing consistent detector arrays for Stage-3 fabrication. Fabrication engineers will use inspection tools including profilometers, reflectometers, ellipsometers, sheet resistance, and film stress monitors on these test structures to monitor fabrication processes attributes such as film uniformity in deposition, etch rate, surface roughness, dielectric constant of the dielectric film, sheet resistance, and film stress. Additionally, materials will be tested at cryogenic temperatures. Variable temperature four-wire resistance measurements will be used to measure the transition temperature of superconducting materials. Low-frequency ($\approx 2\text{--}8$ GHz) diagnostic microwave resonators will be used to evaluate the loss in the dielectrics and superconducting materials.

Detector characterization: Detector characterization is an essential part of the fabrication process. Characterizing test structures, single pixels, and full arrays provides critical feedback to the fabrication process and ensures that the performance and uniformity targets are met. While finished detector arrays are tested at room temperature for connectivity, much of the detector characterization requires cryogenic temperatures to verify detector performance at operating temperatures. Here we discuss tests during the fabrication process, single pixel testing before array fabrication, and how characterized data are used to monitor detector fabrication performance. Further details on full the detector assembly and its testing are discussed in Sect. 4.5.6.

First, single pixels will be tested in a “dark” environment where there is no optical signal on the detectors. Dark tests provide the fabrication team with information about detectors such as: superconducting transition temperature of the TESs, superconducting transition shape of TESs, resistance of the TESs, saturation power of the TES bolometers, time constants of the detectors, and noise performance. These single pixel and full array testing results will be continually fed back to fabrication teams, who will make adjustments to processes to ensure that detectors achieve the specified performance. For example, if the time constant

of the detectors is too fast, the thickness of the heat capacitor material (palladium) can be increased to add more heat capacity to the TES bolometers, slowing the time constant.

Single pixels will then be tested in an “optical” environment, where a blackbody signal illuminates the detectors. Optical tests will provide the fabrication team with information about the RF circuit such as optical efficiency, antenna pattern, polarization purity, and bandpass filter performance. As with dark testing, the fabrication team will make adjustments to the detector design and/or fabrication method as needed to ensure that the detectors are within specification. For example, if the frequency range of the bandpass filter is slightly high, it is possible to make a minor correction by changing the thickness of the dielectric material without changing the detector layout.

Single pixel cryogenic detector characterization is carried out using sub-Kelvin cryostats configured for specific measurements (e.g. a dark cryostat with SQUID readout for measuring the transition temperature, resistance, and saturation power of released TES.) Each of the candidate fabrication facilities has a suite of cryostats to allow for detailed characterization of their fabrication processes and the resulting detectors.

4.5.3 Detector RF coupling

The CMB-S4 reference design uses smooth-wall, wide-band, spline-profiled feedhorns for the detector RF coupling. The feedhorn array sits above the detector wafer, and couples the light onto the orthogonal fins of the detector OMT (Fig. 48). Feedhorns are a mature technology that has been demonstrated on many CMB experiments with published CMB polarization science results across a wide range of frequencies, including ACTPol, AdvACT, SPTPol, ABS, and CLASS.

Spline-profiled feedhorns have good performance across wide bandwidths, making them ideal for multichroic pixels. AdvACT has fielded spline-profiled feedhorns on multichroic arrays at 27/39 GHz, 90/150 GHz, and 150/230 GHz, which spans almost the entire planned frequency range for CMB-S4. Figure 53 shows several spline-profiled feedhorn designs across the CMB-S4 frequency bands and examples of beam maps showing their performance. Spline-profiled feedhorns maintain high coupling efficiency with small apertures, enabling denser pixel packing and thus increased sensitivity. Because their profile designs can be tuned between beam coupling efficiency and beam systematic control, spline-profiled horns allow for greater design control and balanced performance [530, 531, 532, 533].

Spline-profiled feedhorns can also be made monotonically increasing in diameter so that they can be direct-machined into metal, significantly reducing fabrication time and cost. We plan to use custom drill and reamer sets to direct-machine the feedhorn profiles into Al 6061. Al 6061 is cost-effective, low-density, and has high machinability [534]. A CMB-S4-like 90/150-GHz feedhorn array with 507 feedhorns would only take about 20 hours to machine. To efficiently machine a large number of arrays, the feedhorn profiles for a batch of several horn arrays are drilled into a single block of Al, and then individual arrays are cut from the block. The feedhorn array production can be readily performed with the tools available in most machine shops, so it can be spread across multiple machine shops. Because Al is superconducting at the operational temperature of the array, it requires the addition of a normal metal for heat sinking, so we will gold-plate the arrays after fabrication. This fabrication process is simple and as such has an expected yield of 100%.

We will use non-contact metrology on cross-sectioned feedhorns to verify that the machined profiles match the designed profiles and beam maps to check that the feedhorn performance matches simulations [534]. One large advantage of feedhorns is that they can be tested at ambient temperature prior to integration with the detector array. This enables rapid turnaround on array testing and reduces risk by only integrating fully vetted horn arrays with the detector arrays. Before production begins, the drill and reamer set is vetted

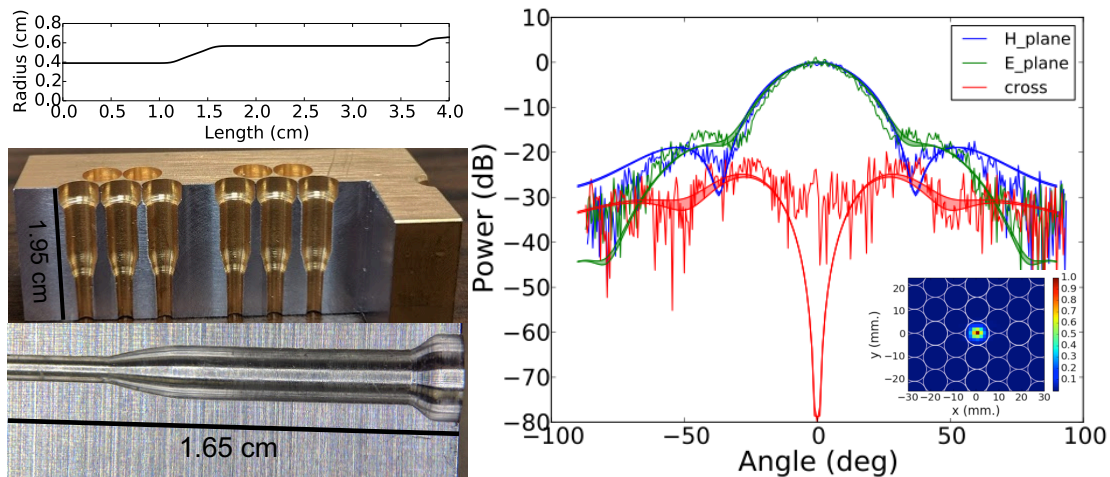


Figure 53. Spline-profiled feedhorn profiles across the planned CMB-S4 frequency bands are shown on the left. The top profile is the AdvACT 27/39-GHz design, which was direct-machined into SiAl alloy. The middle profile is a Au-coated Al feedhorn with the same design as the 90/150-GHz AdvACT feedhorn. The bottom profile is an Al SO 220/280-GHz design. The figure on the right is a measurement of the beams of the AdvACT 150/230-GHz feedhorn at 150 GHz with its simulations. The inset shows a 2D beam measurement of the Al 90/150-GHz design at 125 GHz. The feedhorn beams can be measured at room temperature for fast turnaround and match their simulations well.

with non-contact metrology measurements, and the performance of a single feedhorn is measured with beam maps. After the arrays are produced, we will measure several horns at several frequencies on each array to verify performance. Tool wear on the drill and reamer set can be assessed within batches with the beam maps and between batches with non-contact metrology.

The horn arrays are aligned with the detector stack by a pin and slot to account for the differential thermal contraction between the Al feedhorn array and the Si detector stack. The detector stack is comprised of a photonic choke to reduce leakage at the interface between the horn array and the detector stack [535], a waveguide interface plate, the detector wafer, and a \sim quarter-wave backshort. Behind the backshort is a Cu heat clamp plate for heat sinking that bolts directly to the metal feedhorn array. The feed array also has additional mounting features machined directly into its design for mounting onto the focal plane. This array package design is based on designs used for ACTPol, AdvACT, and SPTPol.

Wide-band, spline-profiled feedhorns are a mature technology. Smooth-walled Al feedhorns have been fielded on BLASTPol [536], and Simons Observatory (SO) will use Al feedhorns for the 220/280-GHz arrays and half of the 90/150-GHz arrays. The main challenges associated with Al feedhorn arrays are achieving the necessary machining tolerances in the feedhorn profiles and the alignment of the feedhorn array with the detector stack, especially at the highest frequencies (220/280 GHz). Simulations using the scatter measured in the a 220/280-GHz feedhorn profile from non-contact metrology measurements show that the current level of variation in the profile is negligible to the feed's performance [534]. Additional simulations show that the SO 220/280-GHz array must be aligned with the detector stack to $\approx 10 \mu\text{m}$. Differential thermal contraction can cause greater misalignment and add additional mechanical stress to the array, but these risks can be retired with cryogenic testing in the design phase. Two prototype SO feedhorn arrays are currently in production and will be tested with full detector arrays by the end of 2018. In 2019, SO will produce, test, and integrate approximately 30 Al feedhorn arrays, further demonstrating the necessary production capacity and performance of this technology for use in CMB-S4.

4.5.4 Cold readout electronics

Cold Readout Introduction The readout of the individual detectors will be multiplexed at the lowest temperature stage using time-division SQUID multiplexing (TDM). In TDM a group of detectors is arranged into a 2D logical array, in which each column of detectors shares a dedicated readout amplifier chain. The various rows are addressed cyclically in rapid succession to record the entire array. When a row of the array is actuated, a single TES bolometer in each column is read out by its associated column amplifier. The signal from each bolometer is low-pass filtered before multiplexing to a frequency below the Nyquist frequency of the sample rate, so that the signal can be reconstructed with acceptable degradation from aliasing. Time-division SQUID multiplexers have extensive on-sky heritage in the CMB community, including instruments such as the BICEP/Keck cameras and the MBAC and ACTpol cameras on ACT.

Figure 54 shows the schematic of one column of a time division SQUID multiplexer. The multiplexer chips outlined in blue in this figure are located at 100 mK, whereas SQUID Array Amplifiers (SAAs) are heatsunk to ≈ 4 K. A detector bias/filter chip (not shown), which sits between the multiplexer and the detector, contains shunt resistors that provide the TES voltage bias and inductors to limited the bandwidth to below the Nyquist frequency of the sampling. The multiplexer footprint in the focal plane is $1.8 \text{ mm}^2/\text{channel}$. The size of the bias/filter chip is driven by pixel bandwidth and TES sensor impedance; in practice, however, the silicon area required is roughly the same per pixel as the SQUID multiplexer chip. Superconducting wire bonds and superconducting circuit board traces pass signals between the detector wafers, filter chips, and multiplexing chips; links to the rest of the amplifier chain may be made with non-superconducting pin-socket interconnects.

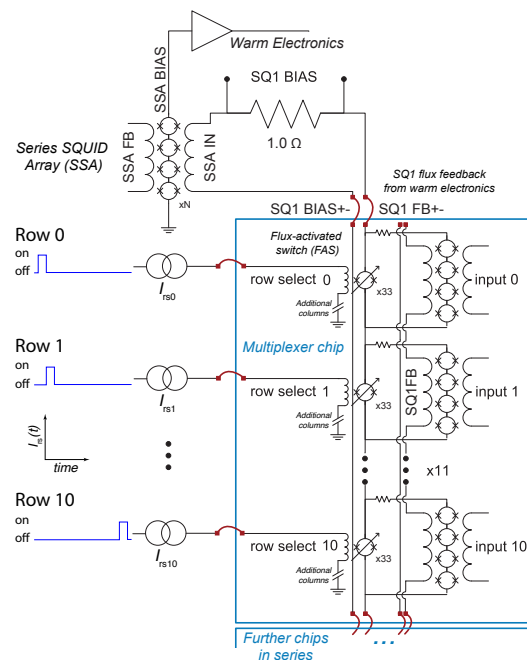


Figure 54. Schematic illustration of a single column of a time division SQUID multiplexer. Each TES is coupled inductively to a first stage SQUID (SQ1). All SQ1s in a column are wired in series to the input of a SQUID array amplifier (SAA), but at any given time all but one row of SQ1s is bypassed by a flux-activated switch (FAS). The various row-select lines are biased in sequence with low-duty-cycle square waves, as shown at left.

Packaging The small per-element footprint of TDM enables the readout to lay in-plane behind the detector wafers. This arrangement enables the close-packing of detectors to efficiently use the available focal plane. The approach has been used in BICEP3 [537] (see Fig. 55), and the same concept is implemented in the CMB-S4 reference design. Given the readout footprint, 150-mm wafers with up to 2,000 channels per wafer may be read out with one layer of multiplexing components, and this channel count is adequate for all the array types in the CMB-S4 reference design.




Figure 55. Behind-wafer TDM packaging implemented for BICEP3. Figure reproduced from Ref. [537].

Fabrication plan Many tens of thousands of TDM SQUID channels have been fabricated and successfully deployed in astronomical instruments. The multiplexer design and fabrication are mature. TDM fabrication will be done at class 100 clean room facilities at the National Laboratories. The baseline fabrication plan utilizes teams of two fabrication engineers carrying batches of 150-mm TDM wafers through the fabrication process, with multiple standardized quality-control measurements.

Superconducting metals (niobium, aluminum, and niobium trilayer tunnel junctions) and normal metals (gold) will be sputter deposited. Dielectric films will be deposited with plasma enhanced chemical vapor deposition. Features will be defined using stepper lithography. Both chemical wet etch and reactive ion etching will be used to define lithographed features.

Fabrication Quality Control All fabrication and metrology steps will be tracked in a computer-based logging system. Status of the fabrication tools and clean-room environment, including chamber pressure, platen temperature and clean-room temperature and humidity will be recorded each time a tool is used. Data will be continually compared to historical norms to identify deviations from process stability.

Each TDM wafer and 32-channel TDM chip includes test structures that allow monitoring of fabrication processes and material characteristics. Sacrificial test wafers will also be carried through some fabrication steps for process monitoring. Fabrication engineers will use inspection tools including profilometers, reflectometers, ellipsometers, resistance sheet and film stress monitors on these test structures to monitor fabrication processes such as etch rate, surface roughness, dielectric constant of dielectric film, sheet resistance and film stress.

One critical challenge for detector fabrication is to develop robust fabrication procedures that guard against process deviation. Fabrication tools will require regular maintenance by highly skilled engineers.

Cryogenic Multiplexer Screening After fabrication, a 150-mm SQUID multiplexer wafer is diced into individual time-division multiplexer chips. The reference design implements a multiplexing factor of 64, with

each readout column incorporated on two 32-channel TDM chips. Bias/filter chips will match this channel count and approximate form factor.

Every 32-channel TDM chip will be screened in a fast-turnaround 4-K test facility. Only chips with 100% working channels will be accepted. Conservatively assuming a 67% yield of 32-channel chips with all working channels, each successful 150-mm wafer will provide 4,096 channels of TDM readout (128×32 -channel chips). Bias/filter chips contain much simpler passive components, and are thus expected to be fabricated with higher yield and throughput. Occasional batch-screening of these chips in a similar fast-turnaround cryogenic system is adequate to guard against drifts in component values.

Based on previous experience, a production rate of 10 science-grade wafers/year/FTE may be assumed.

4.5.5 Warm readout electronics

Time-division SQUID multiplexing The time-division SQUID multiplexing (TDM) readout system includes room temperature electronics to control and bias the multiplexer, provide a bias to the bolometers, acquire and filter data from each bolometer in the array, and exchange commands and bolometer data with the observatory data acquisition (DAQ) system.

The warm electronics provide low-noise current and flux biases that configure the multi-stage SQUID amplifier into a working state. The SQUID series array (SSA), the last stage of the cryogenic amplifier shown in Fig. 54, is provided with a current bias to activate the SSA and a flux bias to place it in its linear regime. A second current bias is provided to activate each column's first-stage SQUIDs and flux-activated switches. The warm electronics also manage the TDM switching among the various rows in sequence. A fast-switching address card asserts a single row-select bias line at a time, which drives a single flux-activated switch (FAS) in each column normal. This serves to connect that row's first-stage SQUIDs to the readout chain while the remaining rows are shorted out by the other rows' inactive switches. The row select timing is chosen so that all rows are visited at a cadence of ≈ 20 kHz.

The warm electronics also operate a separate closed-loop servo to linearize and record data from each bolometer in the array. As each row is activated by its FAS, the warm electronics apply an appropriate nulling feedback signal (SQ1FB) to each column to keep the first-stage SQUID in that column in a linear regime. The SSA output voltage forms an error signal input to a PID loop which calculates the appropriate feedback signal for the next visit. The feedback and error signals are buffered, filtered, and reported to the observatory DAQ, typically at a few hundred Hertz. Mature software exists to tune and optimize the performance of this multi-stage multiplexer system.

In addition to providing the SQUID currents described above, low-noise bias cards in the warm electronics provide detector biases to the TES sensors in the array. These are absent in Fig. 54, but would be connected at the right in the diagram. Scripts running in the warm electronics optimize the bias levels in the SQUID-based amplifier system and tune the bias levels of the TES array in a matter of seconds.

TDM systems have been used to control and readout large TES arrays in sub-millimeter astronomy (SCUBA-2, Zeus2, and SOFIA), CMB experiments (ACT, BICEP/Keck, CLASS, SPIDER, ...) and also for "homeland security." These systems operate at sea level and high altitude sites, in the stratosphere and at the South Pole controlling bolometers in instruments as large as 10,000 bolometer elements.

Low-risk path to modernization The present generation of TDM warm electronics is already capable of handling the bolometer arrays proposed for S4 without a redesign. However, the system is based on field-

programmable gate arrays (FPGAs) which are near the end of their market cycle. These FPGAs should be replaced by modern units, a straightforward design effort. The optical fiber drivers in use for communication between the warm electronics and the computer directly connected to it are now obsolete and should also be replaced. This provides an opportunity to move to a more modern communication format. The design effort for this change is also modest.

The present electronics have been used to control arrays with as many as 64 rows, sufficient for the 2000-channel and smaller arrays in the S4 reference design. Moving to systems larger than 64×64 (4096) elements would require faster settling times in the cryogenic SQUID circuits than those presently in use. This would induce an undesirable performance coupling between the warm and cold electronics, and so is not incorporated into the reference design.

The present electronics are fabricated on ordinary multilayer circuit boards by commercial vendors and housed in an industry standard 6U crate, with minor mechanical modifications for direct electrical coupling to the cryostat. This system will be simple to manufacture on the necessary scale.

4.5.6 Modularity, packaging/test plan

Each CMB-S4 detector module is assembled in a vertical stack outlined in Fig. 57. The stack consists of the following three parts.

- Detector stack: the detector array, backshort wave guide and backshort end plate, with the latter including cryogenic magnetic shielding.
- Cold readout stack (CRS): the DI wafer, multiplexer chips, and IO wafer. This part houses all of the sub-kelvin multiplexer components. It is assembled first and screened at 4 K using a bank of test resistors. We require a CRS to have > 95% operational yield.
- Mechanical housing: feedhorn array and back plate, with the latter including micro-D connectors to connect wiring to the 4-K stage. This part provides the EMI shielding and mechanical structure for the module. It is mounted directly to the Au-plated Cu plates of the focal plane stage in the optics tubes.

The semi-hexagonal and rhombus shaped modules will have an identical vertical stack with appropriate overall form factor.

Modules are packaged as follows.

- The CRS is assembled from cryogenic components that have passed screening. The chips and wafers are glued together using cryogenic epoxy and the entire assembly is wirebonded using Al wirebond. As discussed above, the CRS will be tested at 4 K with a bank of calibrated resistors to verify high operational yield.
- The detector stack and CRS are assembled on top of an appropriate feedhorn array. Custom copper beryllium springs [538] will keep the stacks on the feedhorn and provide thermal contact to the temperature bath. Au wirebonds will be used to provide additional heat sinking of stack components.
- The detector array is wirebonded to the CRS using Al wirebond. The IO wafer is then wirebonded, with Al wirebond, to flexible circuit cables [539], which are glued by rubber cement to the IO wafer.




Figure 56. Block diagram of one bolometer array controlled and read out by one box of warm TDM electronics. Typically all boxes attached to a given cryostat are synchronized at the 50-MHz ADC clock or at the data strobe rate, which is often ≈ 1 kHz. No technical innovation is needed to use the present electronics to control a 64×64 element array. Command, data and clock connections to the Observatory DAQ are via optical fiber. The system requires three twisted pairs of wires per column to provide SSA bias and feedback and a SQ1 feedback signal. Twenty pairs are required to operate a 64-way multiplexer in the column direction, providing the Row Select signals. The number of wires required for bolometer bias depends on anticipated element uniformity. The resulting cryogenic harnesses of a few hundred wires are constructed of low thermal conductivity wire woven as twisted pairs. The system is comprised of several Readout and Bias cards, one Address card and one communication card in a rack. The Communication card is attached to a PC interposed between it and the Observatory DAQ.

- The entire module is enclosed by the backplate with the flexible circuit lines passing through appropriate slots. The flex lines are glued to the backplate and Al wirebonded to the PCB housing the micro-D connectors.

Packaging will take place at a number of facilities with required assembly tools including: custom mounting jigs; automatic wirebonders; and technical staff associated with maintaining and operating these tools. Facilities equipped with these resources exist at the national labs and some universities. Module packaging is expected to take one day per module. All detector and readout components will be tracked using a suitable database that collects all the information related to testing and device history. Each device will be issued its own tracking number and traveler document.

All assembled modules will be tested twice, once “dark” and once “optical.” Dark tests include measurements of IV characteristics versus stage temperature, operational yield, noise (including low frequency), and




Figure 57. Schematic of the CMB-S4 detector module stack consisting of a feedhorn array, detector array, backshort wafers, and the CRS. The CRS consists of the DI wafer with mux chips and IO wafer.

optical efficiency (via a cold thermal load). Optical tests will measure the detector optical bandpass for approximately 30% of the array. This approach exploits the fact that the polarized beam is determined by the feedhorn array and alignment of the OMT to the feedhorn assembly, which can be completely measured using room temperature measurements during feedhorn evaluation and module assembly. As a result, the only required cryogenic “optical” measurement is spectroscopy to evaluate the bandpass.

The full suite of tests will take approximately one month and requires 24 test cryostats, each operating three modules at a time. These cryostats will be primarily located at national labs, with a few potentially at universities. The test cryostats, including readout, will be procured, verified and cross calibrated prior to the start of production and all testing results will be stored in the database and reviewed by the project fabrication team.

4.6 Data acquisition and experiment control

The reference data acquisition and experiment control subsystem is closely based on that used in existing Stage-3 CMB experiments, with a clear technical path to handling the higher data rates from CMB-S4. This subsystem encompasses the acquisition of high-rate (7 Gbit/s across both sites, 3–4 Gbit/s per site) sample data from the detector arrays, the acquisition of low-rate data principally from thermal, position, and pointing sensors attached the telescope (broadly described as “housekeeping” hereafter), initial storage of data to on-site disk, the provision of timing and frequency reference signals, telescope pointing and cryogenic control, and real-time monitoring systems.

In the following, each CMB-S4 site is considered a completely independent entity; the logistical differences between a polar and a Chilean site both impose somewhat different requirements and make any unified control complicated with little return. The systems in both cases are either similar or identical, but are not planned to be connected.

The total expected data rate to disk from CMB-S4, summed across both sites, is approximately 8 Gbit/s, dominated by bolometer data.

4.6.1 Data acquisition

4.6.1.1 Bolometer data acquisition

Following the designs in existing projects, the warm readout electronics (Sect. 4.5.5) will be connected to Ethernet through several intermediate control computers (L1 aggregators) implementing the interface to the warm electronics (either Ethernet or a custom format implemented through a PCI-E card, depending on the modernization discussed in the previous section). These convert the data to a generic format and collate the responses from multiple boards, then forward them to another system (L2 aggregator) by Ethernet, which in turn assembles final data frames from an entire focal plane. As we do not record phase information, there is no clear benefit to taking data jointly from multiple telescopes and we plan to treat each telescope as an independent entity for the purposes of data acquisition.

The intermediate format used by the L1 and L2 aggregators is baselined as the common format currently used by SPT-3G, POLARBEAR2, and Simons Observatory, but may be adapted to CMB-S4 pursuant to the needs of the instrument. This format will, to the extent possible and to the extent to which the L2 aggregator and later data acquisition steps depend on it, be independent of the employed readout technology, limiting the impact of changes in readout design (Appendix B) to hardware and software changes on the L1 aggregator points.

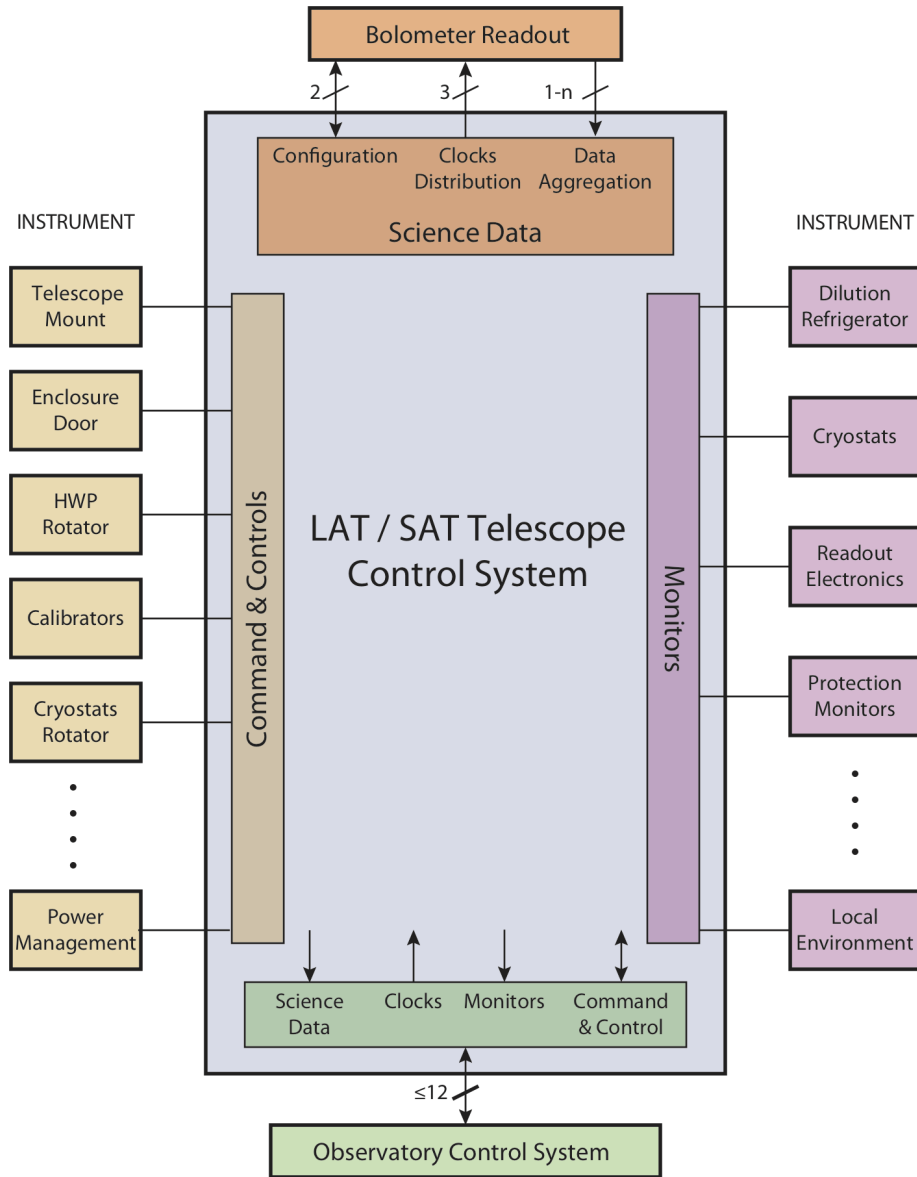
The L2 aggregator forms the only serial chokepoint in the data flow and needs to be able to aggregate at minimum one focal plane's worth of data in one place. This architecture has been demonstrated on SPT-3G with 16,000 channels. The highest-burdened serial stage of this data collation uses CPU time linearly with channel count and is a factor of three from the limit of a 2017-era CPU, allowing scaling to single receivers of 50,000 detectors, comparable to that planned for CMB-S4 (Table 3-2), without any further engineering effort.

Similarly, we are currently baselining the data acquisition code written for SPT-3G and POLARBEAR2 (and planned for use with Simons Observatory), which appears capable of handling the high data rates expected for CMB-S4 in limited testing. Extensive software research and development work for bolometer data acquisition is not expected.

After the L2 aggregator, data is forwarded to a first-level disk cache for later non-real-time unification with housekeeping data, compression, and archival. This follows the successful strategy employed by all Stage-3 instruments, which have demonstrated this to be a highly parallel process. It can thus be straightforwardly scaled to CMB-S4 data rates with current commercially available processing capabilities.

4.6.1.2 Aggregation of bolometer and housekeeping data

In addition to bolometer data, we will record data from a variety of housekeeping systems, which include telescope mount encoders and tilt meters, cryostat thermometry, a weather station, and other monitors of the observatory status. A computer running acquisition for subsystems will be responsible for packing its own data into timestamped containers that are sent via Ethernet to the first-level disk cache. The container data structure supports time-ordered data with different sample rates; most of the housekeeping systems will



Note: Interfaces to Facilities not shown

Figure 58. Telescope platform detail of the observatory control system, showing high-level control of a telescope platform, warm readout, and housekeeping systems.

be sampled much slower (≈ 1 Hz) than bolometer data. All subsystems will be synchronized via IRIG-B002, as discussed in Sect. 4.6.3, so timestamps can be used to identify and aggregate simultaneous data. It is not necessary for this process to occur in real time (though it will be nearly real time, in practice), so it is mostly insensitive to lags in delivery from particular subsystems. The bolometers dominate the overall data rate, so the additional housekeeping subsystems will not be a substantial burden on the network or the first-level disk cache.

After the first-level disk cache, we plan a third data aggregation step (L3) that happens asynchronously using an on-site computing cluster. At L3, the auxiliary data streams are merged into the same files as the high-rate bolometer data so that all the data at a particular time are colocated for correlation analyses during processing. A second copy of the auxiliary data will be stored to a database for retrieval by tasks that do not also require the bolometer data and can benefit from a more lightweight interface, such as site monitoring.

4.6.1.3 Live monitoring of bolometer data, auxiliary data, and alarms

The data monitoring subsystem is a server/client system that aggregates data from each of the subsystems that communicates with the OCS, and presents it over a web interface for monitoring of data quality and site status. The server provides a common API for each connected client to supply and receive monitoring information from other subsystems, and to issue alarms. Each telescope site maintains one monitoring server, and clients can connect from anywhere to request data, often from multiple sites. This system is separate from the main data management system described in Sect. 4.7, as it largely handles data provided in near-real-time, although it should also be able to retrieve archived data should a client request it.

For example, one such client would be a housekeeping consumer that monitors the state of each cryogenic system. This client provides a web interface to view housekeeping data retrieved from the server, and issues alarms if temperatures are out of scope. An alarm may be issued as an email to the operations team, a text message or a telephone call to the site, or a simple warning on the web interface.

A separate monitor/viewer is maintained for bolometer data at each telescope site, either streaming in near-real-time, or as snapshots of specific statistics (e.g., average current noise in a particular frequency range, depth into TES transition). These data can be used during calibration runs, and for monitoring data quality during normal observing periods. For example, one might issue an alarm if too many channels are outside of the normal operating range on the TES transition. Subsets of the data aggregated by this monitor are also sent to the central monitoring server.

4.6.2 Observatory control

4.6.2.1 Main observatory control system

The Observatory Control System (OCS) consists of software to facilitate both automated and user-interactive coordination of operations involving multiple telescope platforms, including all on-board subsystems (warm readout control, telescope encoder, half-wave-plate, etc). The full set of subsystems are likely to include a combination of commercial devices alongside specialized hardware designed for CMB-S4 or adapted from existing CMB projects.




Figure 59. Overview of the observatory control system, including the alarms, data acquisition control, and monitoring, including the interface to all telescope platforms.

While individual telescope subsystems within the observatory may produce large volumes of data, or may need to operate in hard real-time to meet performance requirements, there are no interactions between subsystems that require exchange or analysis of large amounts of data in order to maintain instrument performance. Furthermore, the synchronization requirements, when coordinating operations of various systems, are easily achieved assuming NTP-synchronized clocks and typical Ethernet latencies. As a result, we are free to adopt a simple OCS architecture wherein all devices accept commands and exchange low-speed data on the observatory network using a simple interface language defined by the OCS. All devices will produce status information, such as the success or failure of requested operations, in an OCS-defined format. Control software will operate by translating high-level requests (such as a request to observe a certain sky field for a certain amount of time) into a sequence of lower-level tasks, and then initiating and monitoring the execution of those tasks on the relevant subsystems. Any on-going device operations that are unrelated to specific observing requests will be initiated and monitored through the same system, but by different, persistent high-level control daemons within OCS.

In practice, the implementation of such an OCS involves creating, for each device, an associated piece of server software that can accept requests formatted according to the OCS protocol. For non-commercial, ad-hoc devices, such a server can be written as the primary means of control for the device. In the case of commercial devices, the associated server might consist of a simple translation of OCS requests into requests on the device's native high-level interface. Suitable libraries and example code will be provided by OCS maintainers for use in the development of subsystem servers. The OCS command and data exchange protocol will be implemented on top of remote procedure-call and publish-subscribe systems supplied by open source, community-maintained software libraries.

4.6.2.2 Telescope control

The telescope control system is comprised of an isolated computing environment that communicates directly with the telescope on-board computer (provided by the telescope manufacturer). This system receives pointing data and diagnostics from the telescope at a fixed rate, and schedules observations by providing

a sequence of scan profiles to control the telescope position and velocity. Commanding to the telescope is asynchronous with data acquisition to avoid reliance on strict real-time processing, while allowing the on-board control loop to run at a fast enough rate (typically 100 Hz) for smooth telescope operation. The control system interfaces with the OCS to receive high-level scheduling commands, send status messages, and process alarms.

4.6.3 Hardware—timing and synchronization

Synchronization for each observatory will be accomplished with IRIG-B002 timestamps distributed to the bolometer readout, telescope encoders, and other housekeeping subsystems. A timing card specified for use by most Stage-3 experiments will provide absolute GPS timestamps and can free-run without a GPS lock with minimal drift (5 parts in 10^{11}) over 24 hour time scales. It is configurable with outputs of IRIG-B timestamps, a 10-MHz oscillator, and a pulse-per-second (PPS). With an Allen variance on the 10-MHz oscillator of $< 10^{-11}$ s, these achieve the stringent requirements necessary for frequency-based readout as well as the $1 - \mu\text{s}$ absolute timing accuracy sufficient for all readout and subsystems baselined in the reference design. There are additional options we are likely to explore, in particular PTP networking protocol across each site. In particular, we are not baselining a synchronization pulse or word to be distributed across the site, however individual subsystems (e.g., synchronization between readout crates on a platform) may be required and specified by the subsystem, and that synchronization signal could be timestamped and recorded as well.




Figure 60. Distribution of timing signals (IRIG timestamps, 10-MHz clock line, optional pulse-per-second) for a given observatory.

4.6.4 Scaling and technical readiness from Stage-3 experiments

In large part, data acquisition technologies developed for Stage-3 experiments can be applied to CMB-S4 without fundamental changes in approach (and in some cases can be reused as-is), limiting technical risk in this sector to a low level. The only major uncertainty is the communications protocol between TDM electronics and L1 DAQ computing, which needs to be developed. As this communication does not approach any fundamental limits (bandwidth over fiber optic cables, hard latency requirements, etc.) this task is, like many of the others described here, a low-risk technical engineering task. The demonstrated performance of the technologies outlined here as part of Stage-3 experiments thus provides a high technical readiness level for CMB-S4.

4.7 Data management

The data-management subsystem deals with the transport, storage, and reduction of the instrument data, the generation and reduction of synthetic data sets, and the internal and external distribution of the derived data products. The remit of the data-management subsystem extends from the raw data recorded on disk at the two sites (where the remit of data acquisition is defined to end) to the production and distribution of well-characterized sky maps in the various CMB-S4 observing bands (where the various science analyses are defined to start). The subsystem includes dedicated on-site hardware (storage and analysis servers), computational resources in the U.S. (cycles, storage, and network bandwidth), software (including modules, pipelines and frameworks), data distribution (internal and external), and subsystem management.

The biggest project cost will be the personnel required to create, maintain, and execute the data analysis software stack. A smaller effort will be devoted to the systems for storing the data on site, transmitting it to the U.S., and distributing it to the collaboration and community. The smallest component will be subsystem management, including technical and project leadership. In addition to the managers, engineers, and scientists included as part of the project cost, it is assumed that additional scientists in the CMB community—at both national labs and universities—will provide significant uncosted resources in both construction and operations.

The computational resources required for the data-management subsystem are dominated by the unprecedented scale of the storage and cycles required for data analysis in the U.S. However, following the example of the *Planck* satellite mission (the closest analog to CMB-S4 in data management), our assumption is that these resources will be acquired through agreements with major national computing facilities, and they are therefore assumed not to be part of the project cost ¹. A much smaller on-project hardware component will provide the dedicated data storage, processing and transport resources required at each site.

4.7.1 Design drivers

The primary driver for the design of the data-management subsystem is the unprecedented volume of the CMB-S4 data set (Table 4.7.1).

This volume—comparable to that of the triggered LHC or LSST—is orders of magnitude greater than current CMB experiments. The US data center must have sufficient archival storage to hold the full data volume,

¹The one exception to this for *Planck* was the need to augment the standard NERSC disk quota; however under NERSC's storage plan for the next decade [540] this would not be necessary for CMB-S4.

Telescopes	Detectors	Sampling rate (GHz)	Data rate (Gbps)	Transmission (TB/day)	Storage (PB)	Samples ($\times 10^{15}$)
18 \times SAT	153360	100	0.46	1.69	4.31	3.4
1 \times LAT (delensing surv.)	114432	400	1.36	5.04	12.87	10.1
2 \times LAT (Legacy surv.)	243520	400	2.90	10.72	27.38	21.5
Total	511312	...	4.72	17.44	44.56	35.0

Table 4-17. *The reference design data rates and volumes for 7 years of observing by the 18 small and 3 large aperture telescopes (SATs and LATs respectively), assuming that the data can be compressed to 35% of their raw volume for transmission and storage, and only uncompressed on the fly for analysis.*

enough spinning disk to hold as much of the data as is required at any one time for analysis, and enough cycles to support the full data analysis program, all available for the duration of both the construction and operations phases. This translates to needing exa- to zetta-scale computing resources for the next 15 years, which can only realistically be provided by national computing resources such as the DOE high performance computing (HPC) and the NSF high throughput computing (HTC) facilities.

The data-management software stack must be robust, validated and verified, and able to run efficiently at appropriate scale on whatever computing architecture(s) are available to us at any epoch. Both the scale of the data to be processed, and the computational architecture on which that processing will occur, will evolve significantly over the lifetime of the project, and this evolution must be accounted for in both provisioning the computational resources and planning the software lifecycle.

Additional constraints are imposed by the bandwidth available from the remote observing sites and the need to deliver data products at the cadences required for monitoring data quality and supporting transient studies (taken to be daily here, but subject to further input from the scientific community). Both sites must have sufficient reliable local resources to store data pending transfer to the primary data center in the U.S. (including failsafe storage), and to perform any data reductions required prior to transfer.

4.7.2 Reference design

Figure 61 shows a schematic view of the reference design data-management subsystem, including on-site resources, data-transport systems, archival storage, computing systems, the overall software stack, and data distribution within and outside of the project.

4.7.2.1 On-site resources

The reference design assumes sufficient network bandwidth from Chile to return the data to the primary U.S. data center essentially in real time, following the models of both the Simons Observatory and LSST. As such, we only need to provide a sufficient local storage to hold a few days data in the event of a temporary network outage. However this is not assumed to be the case from the South Pole, and instead we follow the model of the South Pole Telescope and BICEP/Keck experiments, with a small subset of the raw data, together with some reduced data products, transferred over the network to the US each day, and the entire data set stored on site and shipped once a year. This means that the South Pole site must be provided with not only sufficient disk to store and transport a year's data, but also with sufficient computing resources




Figure 61. Schematic view of the data-management subsystem, spanning the range from Data Acquisition to Science Analyses, with on-project cost areas highlighted in cyan. Note that the named networking, storage and compute resources are indicative and anticipated, not yet confirmed.

to perform all time-critical data reductions—primarily to generate the data products that will be sent over the network each day, including live monitoring for data quality and transient studies. This piece of the data-management system requires close coordination with the site infrastructure systems to ensure that sufficient space and power are available for the on-site computing resources, especially at the South Pole. Taking conservative estimates, the reference design includes 100 TB of storage in Chile, and 5 PB of storage and 10 Tflop/s of computing at the South Pole.

4.7.2.2 Data transport

Raw data files will be registered on-site as they are gathered and this information will accompany the files to the primary U.S. data center. For the Chilean site, a fiber-optic network connection will be made to the nearby ALMA Array Operations Site as part of site development. From there, several possible paths are available to provide the necessary high-bandwidth connection to ESnet and hence to the data centers, including the current REUNA/AMpath networks and the possibility of sharing dedicated LSST fiber. From the South Pole, daily transfers of reduced data products will use the TDRS satellite link and MILnet. These transfer capabilities are controlled by the limited number of satellites in orbits visible from the South Pole and the limited fraction of the day during which they are above the horizon. As noted above, current network data-transfer capabilities (about 100 GB/day) are not sufficient for the full CMB-S4 data rate. For the reference design we will therefore plan to store each year’s data on site and transfer them to the U.S. at the start of each Austral summer, however we will also continue to work with NSF and DOE to investigate possibilities for implementing a connection to the US data center with sufficient bandwidth to stream the data. From the primary U.S. data center the data will be transferred over ESnet to the secondary center to provide a secure backup. The various reference design paths are illustrated in Fig. 62.

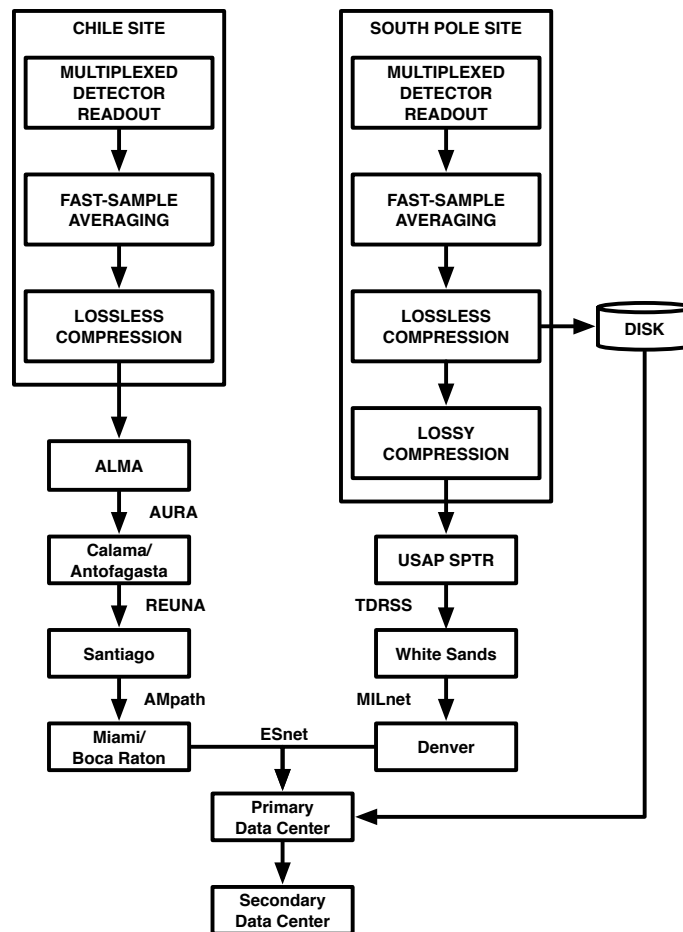


Figure 62. Schematic view of data-transport paths from acquisition to the U.S. data centers for the Chile and South Pole sites.

4.7.2.3 Data centers

The U.S. data centers must provide archival storage for the entire CMB-S4 data volume (including all derived data products and Monte Carlo simulation sets) at two distinct sites, sufficient spinning disk to hold the data volume required for the largest single analysis, and sufficient cycles for all of the data processing, including both the many re-processings of the real data and the generation of Monte Carlo simulation sets. Given the CMB-S4 data volume and project lifetime, these resources can only realistically be provided by national computing facilities. Such facilities have the additional advantages of being cost-free to the project, and of having long-term development plans that will see their resources steadily increase over the lifetime of the project—with new single-site DOE HPC systems being fielded approximately every 4 years, and more continuous growth of the distributed NSF HTC systems. However, given the variety of architectural responses to energy-constrained computing, and the increasing effort required to port to and optimize for them, we should also to limit the number of architectures we will use by adopting centers whose development paths are as similar as possible. The reference design therefore assembles the necessary resources from the DOE’s National Energy Research Scientific Computing Center (NERSC) and Argonne Leadership Class Facility (ALCF), and the NSF’s Extreme Science and Engineering Discovery Environment (XSEDE). Local computing resources may also be used by individual members of the collaboration for smaller-scale analyses, but their deployment is outside of the remit of the data-management subsystem except insofar as the data are made available from the data centers.

Given its ability to provide all the required types of resource, its co-location with ESnet, its 20 year history of supporting the global CMB data analysis community, and its open access policies, the reference design adopts NERSC as the primary U.S data center. Since so many existing CMB experiments already use NERSC resources, this will also allow us to leverage community-wide data and software. The secondary archive is then located at ALCF, with a very high bandwidth connection with NERSC over ESnet, and additional cycles are provided by both ALCF (with the most similar systems to NERSC) and XSEDE. Having cycles at multiple facilities mitigates against the risks posed by down-time or queue congestion at any one facility; having both HPC and HTC cycles allows us to map our processing to the most appropriate resource for a particular pipeline (specifically depending on whether the pipeline requires the inter-process communication supported by HPC but not HTC). Having cycles at a limited access facility devoted to “heroic” computations (ALCF) also provides additional resource security for our most intensive computations. ALCF also hosts many of the large-scale structure simulations being used to generate mock catalogs for experiments like LSST, *WFIRST*, and *Euclid*, which we hope to leverage in our sky modeling to support joint cross-correlation studies.

Since these are national facilities that allocate their finite resources to a large number of projects, we must mitigate the risk of insufficient allocations for CMB-S4. Following the example of the *Planck* satellite mission, we will seek a formal memorandum of understanding with DOE to guarantee long-term and sufficient NERSC access. With this in hand, we will then seek complementary agreements with DOE and NSF for long-term access to ALCF and XSEDE.

As noted above, both the maximum data volume and the computing architecture will evolve over the lifetime of the mission, so we must be confident that we will have sufficient resources at each epoch. One key milestone is to demonstrate that we can process the full CMB-S4 data volume before operations commence; in order to reach this goal we plan to double the data volume being synthesized and reduced every two years. Table 4-18 maps the resources needed for the 2020, 2022, 2024, and 2026 (i.e., construction project) data volume challenges to those known or anticipated on the various generations of NERSC flagship system that will be used, illustrating that the requirements can be met at NERSC (even while anticipating resources elsewhere too). Note that each of these runs require a significant fraction of a top 10 HPC system, illustrating the need for resources of this type. Since this will involve three generations of HPC system at NERSC, we will interleave the data volume challenges with architecture challenges, porting to and optimizing for a new

architecture in each of the intervening years. With this schedule, the architecture challenges would be in 2021 (Perlmutter), 2023 (Aurora at ALCF), and 2025 (NERSC-10).

Year	2020	2022	2024	2026
NERSC system	Cori (KNL)	Perlmutter	Perlmutter	NERSC-10
Cycles (P flop/s)	30	150	150	1000
Memory (PB)	1		4	8
Disk (PB)	30	150	150	300
Data fraction	1/8	1/4	1/2	1
Samples ($\times 10^{15}$)	0.9	1.8	3.5	7.0
Total cycles (E flop)	2.7	5.5	11	22
System fraction	3%	1%	2%	1%
Peak memory (PB)	0.8	1.6	3.2	6.4
System fraction	80%	40%	80%	80%

Table 4-18. *Estimated computational resources available at NERSC and required by CMB-S4 to scale the data-management software to the full data volume prior to first light, assuming characteristic observing efficiency, operations per sample (including both synthesis and reduction), computational efficiency, and memory high-water mark.*

Finally we note that this combination of network, storage and cycle resources, spread across multiple facilities, and connecting a major scientific instrument to a distributed collaboration, is conceptually what is sometimes called a “data superfacility.” We will therefore work with the superfacility experts at NERSC and elsewhere to see how we can leverage their work to make this a seamless computational ecosystem.

4.7.2.4 Software

The development and deployment of the CMB-S4 software stack comprises the bulk of the data-management subsystem. As with many aspects of CMB-S4, the primary challenge here is that of scaling existing technologies to unprecedented levels rather than developing entirely new ones, and hence more about efficient, scalable, implementations than algorithms per se. In addition, given the long-term nature of the project and the expected evolution of (and possible revolution in) computer architectures over its lifetime, porting and optimizing the software stack to new state-of-the-art systems will be an equally critical task. Note that even if existing system had the longevity to support CMB-S4 throughout its lifetime, they do not have the capacity or capability we will ultimately need, so that progression through multiple generations of system must be accepted as an inevitable feature of any data-management plan.

Another way in which CMB-S4 data management will necessarily differ from most (if not all) previous CMB experiments will be in the quality of our software engineering. Rather than the traditional patchwork of single-author/single-user tools developed on a best-effort basis and strung together into a loose sequence of analysis steps with untracked intermediate data products passed between the individuals running them, we will need a software development team producing a coherent suite of stable, tested, documented, modules that can be assembled into efficient, seamless, pipelines whose full provenance can be recorded at any run. While this approach is largely unfamiliar to the CMB community, it is widely used elsewhere and we will take advantage of the lessons learned and adopt established best practices.

The software can be divided into three broad categories—data synthesis, data reduction, and the overall infrastructure to support these.

Data synthesis: Synthetic data play key roles in CMB studies, including in the design and development of experiments, the validation and verification of their reduction and analysis pipelines, and the quantification of uncertainties in and removal of biases from their science results using Monte Carlo methods. Data synthesis includes modeling both the experiment and the sky it will observe, along with generating synthetic data sets both in the map- and the time-domains.

While the development of detailed models of individual sky components is outside of the realm of the data-management subsystem, it does include the assembly of such models into total skies to be as inputs for synthetic data sets. The total sky includes the temperature and polarization components of the scalar, tensor and non-Gaussian CMB; the extragalactic foregrounds (and their associated lensing of the CMB); and the galactic foregrounds. Realistic sky models will be particularly important for optimizing the frequency coverage of the various telescopes, and for the validation and verification of key science analysis algorithms and implementations including foreground cleaning and delensing. We will seek to support joint analyses of CMB-S4 with large-scale structure observations by adopting common extragalactic skies with experiments such as LSST, *WFIRST*, and *Euclid*.

Experiment modeling provides a parametric model of the entire experiment, comprising the instrument optics and electronics, and the observation environment and survey strategy. The parameterization should be rich enough to support the most detailed (and therefore computationally expensive) time-domain syntheses, while also providing standard reduced representations sufficient for approximate (but computationally cheaper) map-domain and Fisher forecasting approaches.

The data synthesis step applies the experiment model to the sky model to generate the data that would be gathered by the instrument making its observation of that sky. This includes signals from the sky, the atmosphere, and the ground, possibly modulated by a half-wave plate, convolved with each detector’s beam and bandpass, converted to a voltage by the detector, and fed through the readout system, and incorporating all of the detector-detector correlations, for example from multichroic pixels or highly multiplexed readout. The complexity of the processing and scale of the data necessarily introduce a trade-off between computational cost and realism, reflected in the prioritization of both map- and time-domain syntheses (with time-domain syntheses also requiring reduction to maps before analysis). The development and deployment of these two paths will be coupled to take advantage of the strengths of each approach and the ability to cross-check them against one another. Specifically we will ensure that each approach is able to synthesize data from the same combinations of experiment and sky model.

Data reduction: The data reduction pipelines form the core of the data-management subsystem. Comprising experiment characterization, live monitoring, and map-making, they must produce all of the data products required to monitor the performance of the instruments and to support all of the various science analyses, and be able to do so with the required cadence. Note that both the characterization and monitoring pipelines must run near-autonomously and, for the South Pole data, they must be able to be executed on site.

Experiment characterization covers both the measurement of instrument properties and the reconstruction of its observations. This category is particularly important during the deployment phase, and includes detector liveness checks; preliminary noise, calibration, and systematics measurements; and telescope pointing and beam checks. The software for these tasks must be flexible and easy to adapt on the fly in response to

unexpected features the data. Most of the development and implementation of this software is expected to be conducted by junior scientists at universities.

The live monitoring pipeline includes all time-critical processing, including both near-real-time checks on instrument performance and data quality and the generation of the daily data products needed to support transient searches, together with the alert systems associated with each of these. Performance/quality alerts will be passed directly to the instrument teams. The exact cadence and content of the transient searches will be determined in consultation with the wider community; for the reference design we baseline the production of daily maps of both the wide-area and the ultra-deep (delensing) LAT surveys, with the former generated at NERSC and the latter at the Pole and included in the daily satellite data-transfers. These will then be passed to a dedicated collaboration working group to identify transients and source alerts through standard astronomical alerts systems.

The map-making pipeline includes the pre-processing needed to make the raw data match the data model assumed by any particular map-making algorithm, the map-making itself, and the characterization of the resulting maps. Pre-processing includes such steps as finding and flagging spikes or jumps in the data, narrow-band filtering of spectral-line-like contamination from, e.g., detector sensitivity to the cryocooler or drive system, and calibration to convert the raw data voltages to temperatures. This part of the pipeline also needs to be flexible enough to support data exploration and the mitigation of unanticipated systematic effects. The map-making itself can use filter-and-bin, destriping, or maximum likelihood algorithms, each of which has a different data model and therefore requires somewhat different pre-processing. The characterization of the map can range from a simple weight map to a full covariance matrix to a suite of corresponding Monte Carlo realizations. This pipeline is also often iterative, with the map produced being used to refine the pre-processing steps. We anticipate producing maps at each observing frequency, for each survey (SAT, LAT delensing, LAT wide-area), for a range of data cuts in both time and across the focal plane(s); each map-making will necessarily include data from multiple telescopes, and possibly from multiple sites, though neither of these is anticipated to pose qualitatively different challenges to the single telescope, single site, analyses we are used to. The diversity of instrument and survey types in CMB-S4 (low- and high-resolution telescopes, small- and large-area surveys) dictates flexibility in map-making algorithms and implementations; the reference design therefore includes all types of map-making and allows for separate optimization for HPC and HTC architectures, with the goal of conducting direct comparison of the various approaches' products and performance.

Infrastructure: The overall data-management software stack will consist of the individual software modules described above, assembled into task-specific pipelines, which are then executed within architecture-specific frameworks, together with the databases needed to support them.

Pipelines will be built to execute standard sequences of data synthesis and reduction operations, and the critical features will be their overall efficiency and their adaptability. In order to avoid the prohibitive IO overhead of writing out a synthetic data set and reading it back in for reduction, when synthesizing and reducing mock data such pipelines will instead run on-the-fly. Since such data will not (in general) be saved to disk, such pipelines will also incorporate the ability to synthesize a data set once and reduce it multiple times (for example employing different algorithms and/or implementations and/or data cuts) within a single job execution, even when the reduction is destructive.

The frameworks will enable the pipelines to run within any HPC or HTC environment available to us (subject to basic sufficiency requirements), with the goal of being able to take advantage of all of the resources that we have access to in order to minimize any bottlenecks from overall resource limits or from congestion on shared systems. The baseline will be built on the existing TOAST HPC and SPT-3G HTC frameworks, originally developed for the *Planck* and SPT experiments respectively, and the goal will be to be able to

use the same pipelines on both HPC and HTC systems wherever possible. Note that this will only apply to pipelines that require no inter-process communication, for example filter-and-bin map-making. Enabling this hybridization is an open research and development project; for the reference design we assume that at very least any module that operates on a small subset of the data (as is often the case in pre-processing and some types of map-making) will be able to be run in either framework, and that we will have the tools needed to translate in situ between the frame-based (SPT-3G) and distributed (TOAST) data models.

Databases will be employed to track data provenance, support data selection for specific reductions, and enable data distribution to remote sites. They will also be used to record the iterations of the instrument model, which will continuously be refined with the analysis of laboratory data, commissioning test data, and ultimately the real data themselves, and to deliver these to the data synthesis codes. Finally they will be used to record the input parameters, intermediate and final data products, and the code module, pipeline and framework versions used for each analysis to support traceability and reproducibility.

4.7.2.5 Data distribution and publication

All CMB-S4 collaboration members will have access to the raw data and to intermediate data products up to and including the well-characterized single-frequency maps. The main mode for the distribution of data to the collaboration will be via the primary data center, where all collaboration members will also be eligible to have user accounts and have access to significant compute resources. Given its volume, typically we expect users to bring their code to the time-domain data, and only to transfer map-domain data to their local computing resources (though these can obviously be processed at NERSC too). Although they lie outside of the remit of the data-management subsystem, we anticipate that additional derived data products (e.g., component-separated maps, maps of the lensing potential, power spectra and likelihood codes) will be similarly distributed.

Public releases of CMB-S4 data products will occur at regular intervals through well-established CMB data archives. Maps of all flavors (single-frequency, foreground-cleaned CMB, lensing potential, etc), catalogs of galaxies and galaxy clusters, the CMB power spectra, and cosmological parameter likelihood codes will be archived at NASA's Legacy Archive for Microwave Background Data (LAMBDA), including both the real data and subsets of the supporting suites of Monte Carlo map realizations. Time-domain data and the full Monte Carlo suites, together with the code used to process them, will be made available through public directories at NERSC, where users will also have access to the computational resources needed to make practical use of such large data sets. We will also make the data-management software stack (complete with full documentation) public wherever possible both through a public repository on a version control system such as github and as a standard installation at NERSC.

4.7.2.6 Management

The management of the CMB-S4 Data Management team will include both the formal WBS elements of software control, resource management, and subsystem interfaces, and the informal responsibilities of team building and the synchronization of activities and products with the project and collaboration schedules.

The CMB-S4 DM software stack will be fully version controlled using standard repository management tools such as git. All code delivery will include validation and verification, unit tests, and documentation, and the full stack will be subjected to nightly builds for continuous integration. We will also adopt a limited set of programming languages (with formal criteria for expanding that set), and develop—or adopt—a common software style guide. We will ensure that there are regular training opportunities in the software standards

as part of the process for bringing in new team members, or as refreshers for existing members, as well as in coding for new architectures and tools such as debuggers and profilers.

The entire data-management plan rests on the allocation of sufficient computational resources to meet our needs across a range of facilities; the acquisition, allocation, and management of these resources will be a key activity. As noted above, we will pursue formal agreements with DOE and NSF to safeguard our access to the necessary resources across the various facilities, and for each facility and resource we will develop an annual burn-down plan centered on the major production runs scheduled for that year to ensure that these are not limited by insufficient resources. Where possible, the resources needed for the core activities will be allocated to a project production account (which a limited number of core developers will have access to) rather than to one or more individual user accounts in order both to ring-fence the resources and avoid being dependent on single users for production runs.

The data-management subsystem has critical interfaces both within the project and with the overall collaboration. Within the project, these interfaces are with data acquisition (corresponding to the hand-off of data from the telescope to the data centers), and with site infrastructure (ensuring that sufficient space and power are available for the on-site computing resources, and coordinating the transfer of disks to/from the Pole). These interfaces will be mediated within the overall WBS structure through the L2 leads and the project office. There are also equally critical interfaces with the collaboration's technical working groups, to develop a detailed, accurate, experiment model, and analysis working groups, to ensure that the overall suite of data products are necessary and sufficient. During the construction project, it will also be critical for the data-management subsystem to deliver simulated data sufficient for the working groups—both technical and analysis—to support instrument and observation design, and to validate and verify the software stacks for the various science analyses. These interfaces will be mediated through the Technical and Science Councils, as well as through a dedicated and cross-cutting simulations working group that will coordinate data science challenges alongside the data volume and architecture challenges.

The overall data-management team will be built on the convergence of the separate Stage 3 teams, leveraging both pre-project research and development funds and off-project cross-cutting collaborations. As the project matures these will coalesce around the various Level 3 activities (broadly synthesis, reduction, and infrastructure), while maintaining their interconnectedness through cross-cutting activities such as the annual production runs. We will make use of a variety of communication channels (from wikis to email to slack channel and telecons), as well as face-to-face and virtual hackathons and training events to develop the team, and, where appropriate, adopt an agile software development approach.

The project will face numerous review gates, and the data-management subsystem schedule must not only be synchronized to its own reviews, but also to the range of data needs and requirements imposed by those of all of the other subsystems. Here we will make the annual scaling and architecture challenges do double-duty as also producing state-of-the-art data sets for the project and collaboration as a whole, as well as responding to more limited but higher cadence requirements through the data or software needed for smaller and/or simpler simulations, including directly in the map domain.

4.8 Integration and commissioning

Integration and commissioning are the final activities before the completion of the project and the start of science observations. We will integrate and commission each large telescope (LAT) and small telescope (SAT) separately, allowing for a phased start of science observations. Integration and commissioning is done separately, by largely separate teams, at each site. To the extent that I&C requires support personnel from within the collaboration, careful coordination will be required to ensure that support personnel are available.

4.8.1 Scope

Integration and commissioning (I&C), as described here, refers only to on-site I&C activities. We assume here that the following have occurred off-site prior to the start of I&C on-site.

- The LAT has been installed on-site and validated with the vendor.
- The LAT receiver (LATR) has been integrated and tested off-site, and minimally disassembled for shipping. This disassembly is assumed to include removal of the focal planes and placement of shipping braces so that the optics tubes can be shipped inside the receivers.
- The SAT receiver (SATR) and telescope mount have been already tested together off-site, then minimally disassembled for shipping. This disassembly is assumed to include separation of the SATR from its mount and removal of the focal plane arrays from the SATR.

Integration of an SAT involves installing the telescope mount, installing the three FPUs into the SATR, and then installing the SATR on the mount. Integration of an LAT begins with the installation and use of a commissioning receiver to test the LAT; this is followed by the assembly of the LATR (similar to the SATR), and installation of the LATR on the LAT.

Commissioning begins after the end of integration. This includes the execution of validation tests to demonstrate that the receiver and telescopes are functional, and the execution of initial calibration measurements. Commissioning ends when functionality is demonstrated and science observations can begin. Note that functionality being demonstrated does **not** mean that all calibrations are complete; ongoing calibrations are required throughout the observing period.

4.8.2 LAT integration

LAT Integration includes two distinct categories of activities: (1) using a commissioning receiver to validate the integrated performance of the LAT; and (2) integrating the LATR, cooling it down in the LAT receiver cabin, and connecting it to all utility, control and data acquisition connections.

4.8.2.1 Commissioning camera

Prior to integration, the LAT telescope performance will be demonstrated in cooperation with the vendor following the requirements in the LAT contract. This should include demonstration of pointing using a star camera. Then, integration using a commissioning receiver begins. The commissioning receiver will arrive with its focal plane removed. It will be re-assembled and installed and cooled down on the LAT. This will provide a platform for testing mm-wave performance of the LAT and the interface between the LAT and all of the cryogenic and detector readout systems before installing the LATR. A list of tests to be done with the commissioning receiver will be developed by the LAT and LATR teams prior to the start of construction of the commissioning receiver to ensure it has the needed functionality.

Duration days	Receiver tasks	Detector tasks	Other tasks	Number of people
5	Receiver unpacking	Setup clean-room	Prep LAT space	12
3	Remove shipping braces	Check detector modules	Continue	12
17	Install FPUs, close receiver	Install detectors in FPUs	Prep LAT, DAQ, DM	17
2	Install LATR receiver cabin	Setup, test det readout	DAQ, DM	13
2	Pump out receiver	Continue	Continue	13
40	Cooldown, cryo validation	Continue	Continue	17
14	Receiver basic validation	Detector basic validation	Commissioning prep	13
78	Total			

Table 4-19. *Task, durations, and staffing levels during LAT integration.*

4.8.2.2 Installing the LATR

Integration of the LATR is then an intense period that requires 12–17 people on-site for about 80 days. Upon arrival, the LATR will first be re-assembled, then installed in the LAT receiver cabin to be cooled down. The detailed list of work to be done during that time is shown in Table 4-19. At the South Pole, it is critical that this integration work be completed during the Austral summer. Note that since the cooldown takes 40 days, if a problem is found that requires warming up and cooling back down again, it will delay the completion of integration into the winter season, to be completed by winter crew. In Chile, the constraints on total duration are less critical. However, a failed cooldown would still be a significant problem in terms of schedule delay.

4.8.3 SAT integration

Prior to SAT integration, the entire SAT system will have been tested in the United States. On-site, the tasks to be completed for integration, and their approximate resource requirements, are shown in Table 4-20.

4.8.4 Commissioning

The initial validation of individual components (e.g., detectors) will be completed before components are shipped to the sites. This section only covers the validation of the full integrated instrument. There are six main phases of validation during the commissioning stage prior to the beginning of scientific observations, which are listed below.

Integrated telescope and instrument validation: During the first phase, we will confirm that the DAQ, data archiving, and communication systems integrated with the receiver are operational and test the basic functionality of the instrument.

Duration days	Receiver tasks	Detector tasks	Other tasks	Number of people
3	Receiver unpacking	Setup clean-room	Setup SAT mount	10
3	Remove shipping braces	Check detector modules	Continue	10
9	Install FPUs, close receiver	Install detector in FPUs	Prep LAT, DAQ, DM	12
2	Pump out receiver	Setup, test det readout	Continue	10
2	Install SATR on mount	Setup, test det readout	DAQ, DM	10
20	Cooldown, cryo validation	Continue	Continue	10
14	Receiver basic validation	Detector basic validation	Commissioning prep	10
53	Total			

Table 4-20. Task, durations, and staffing levels during SAT integration. This is for a single SAT integration. For multiple SAT integrations, resource estimates should be approximately doubled, but calendar days do not need to be, since things can be done in parallel.

Confirm that the detectors and readout are operational: Next, we will perform current versus voltage (IV) curves to measure the saturation power and verify that it is in the nominal range. We will measure the time constants to verify that they are in the nominal range, using bias steps to get the electrical time constants and a chopped thermal calibration source to measure optical time constants. Using elevation nods, beam filling thermal loads (e.g., with a sparse wire grid), or galactic source observations we will measure the relative responsivity of the detectors and confirm that it is within the nominal range. Finally, using rotating wire grid we will measure the relative polarization angles of the detectors to better than 2° uncertainty. If SATs are deployed Chile, another set of polarization angle measurements with the HWPs rotating will be performed to verify that the polarization angle is stable to 1° .

Validate initial data timestreams: Next we will take initial timestream data to verify that there are not excess glitches, SQUID jumps, or other issues. Less than $< 10\%$ of the timestream should require masking for glitches, etc.

Measure the noise performance of the system: We will take noise spectra curves to characterize the noise performance of the timestreams to verify that they are in the nominal range. Noise measurements will be performed both while the telescopes are stationary and performing various scan patterns. If any excess pickup is detected, we will power cycle various components to determine the source of the excess noise and mitigate it. Additionally, if SATs with polarization modulators are deployed to Chile, we will repeat the noise tests with their polarization modulators on and off to verify that there is less than an order of magnitude of extra noise in the noise curve measurements.

Measure the pointing and beams: Prior to integration with the receivers, the LAT mirrors will be mapped out with photogrammetry to align the mirror panels to $25 \mu\text{m}$ half wavefront error (HWFE). Next, the LAT platform pointing will be validated with a star camera to the pointing specifications laid out in the vendor contract. After full integration is complete, planet scans will be used to verify the LAT pointing to 20 arcsec and beams to $\approx -10 \text{ dB}$. The SAT pointing will be measured to 4 arcmin with a combination of star cameras and planet scans. The SAT beams will be verified to $\approx -10 \text{ dB}$ with planet scans. Additional

near-field measurements of the LATR and SATs with artificial sources could be used to verify that these measurements match those that will be taken in North America prior to deployment.

Survey validation: At the end of commissioning, we will aim to validate instrument performance through a measurement of the instrument noise-equivalent temperature (NET) within a factor of a few via CMB field observations.

4.8.5 Calibration

During science observations, we will perform calibration measurements to provide the calibrations necessary for scientific analyses. There will be two sets of calibration hardware, one for each site. We note that the calibration requirements are dependent on the instrument configurations and require in-depth studies to fully define. The requirements presented here are based on existing instrument requirement studies extrapolated to the scientific goals and estimated noise performance of CMB-S4. In this section, we present a baseline plan for calibration that is designed to meet the most stringent requirements. The final requirements will be determined by an instrument flowdown analysis before the CMB-S4 project baseline design.

Bandpass (same for LAT and SAT): The bandpass requirements are set by the requirements on foreground removal. The bandpasses of the detectors must be known to $< 0.5\%$ uncertainty [541]. Fourier-transform spectrometers (FTSs) will be used to characterize and validate the designed bandpasses of the detectors in situ in the field. Using different optical setups to couple the FTS to the telescope optics will enable us to use the same FTS design for measurements of the SATs and LATs. To achieve the low uncertainty required on the bandpass measurements, the transfer functions of the FTSs must be well understood. We can measure the FTS transfer functions by coupling two FTSs together and measuring their responses. To track the atmospheric-generated variations in the bandpasses, we will use a combination of thermometers, and barometers, and telescope elevation nods calibrated to existing radiometers.

Beams: The polarized and unpolarized beams of the instruments can be decomposed into three types: the main beam, near sidelobes, and far sidelobes. Each of these requires different methods of calibration or measurement that can include both celestial sources and artificial sources. Observations of planets or the Moon can be performed whenever the relevant sources are sufficiently above the horizon; measurements with artificial sources require dedicated setups, typically done once a year if needed. Near-field and mid-field beam measurements (with artificial sources) can be used to validate the beam performance and to inform the beam modeling, while far-field beam measurements feed directly into the analysis. The far-fields of the LATs (10s of km) are much further than those of the SATs (100s of m). Tower-mounted sources can be used in the SAT near/mid-field and the LAT near field, while drone-mounted sources can reach the SAT far-field and the LAT mid-field.

Polarized beams: Polarized sources mounted on a drone can be used for polarized beam measurements, e.g., [542]. The drone can be tethered or coupled with a weather balloon for further stability. These measurements could be further supplemented with polarized beam measurements from tower-mounted polarized sources. These sources and the tower/drone mounting can be also be used for polarization angle calibration.

LAT main beams and near sidelobes: The LAT main beams and near sidelobes will be calibrated with planets. Uranus and Neptune will be used for the main beam down to ≈ -30 dB, and current large aperture experiments have used Venus, Jupiter, Saturn, and Mars to reach the ≈ -50 dB requirement with radial

binning. While we expect that celestial sources will be sufficient for the LAT calibration requirements, tower-mounted thermal sources could be used for deeper measurements of the LAT near-field beams.

SAT main beams and near sidelobes: A combination of planets and artificial sources will be used to calibrate the SAT main beams and near sidelobes. Jupiter and Venus can be used to map main beam response down to ≈ -30 dB. Deeper measurements to the ≈ -50 dB requirement will be achieved using tower-mounted thermal sources coupled to a far-field flat mirror as was done with BICEP. At the South Pole, the planets are low in elevation, so can only be observed if the fixed ground shields are designed to be lowered.

Far sidelobes: Far sidelobes will be characterized using the Sun and Moon in Chile. At the South Pole, tower-mounted sources will be used for response above the elevations of the Sun and Moon. Strong broadband noise sources offer higher signal-to-noise than thermal sources for these measurements. Many far-sidelobe effects can be mitigated with proper Sun and Moon avoidance in the scan strategy.

Polarization angle: To determine the polarization angle of the instruments, we will use a combination of self-calibration, observations of polarized astronomical sources like Tau A and Cen A, and artificial sources. Any polarization angle mis-calibration could give a bias on r . For CMB-S4, this means that the polarization angle requirement is $\lesssim 0.1^\circ$ [4].

Self-calibration: In standard models of cosmology, the EB and TB power spectra should be zero. Any mis-calibration in the polarization angle causes non-zero EB and TB signals. Self-calibration assumes null EB and TB spectra, and uses a minimization of these to recover the polarization angle. However, foregrounds can contaminate EB and TB. The extent to which CMB-S4 will be able to use self-calibration is currently under study.

Astronomical sources: Tau A has been measured to $\approx 0.33^\circ$ and is available to both the SAT and LAT instruments in Chile. The LATs will also be able to observe Cen A, which has been measured to $\approx 1^\circ$, from both Chile and the South Pole.

Artificial sources: Given the uncertainties in astronomical sources and the uncertainty in the level of self-calibration we can achieve with foregrounds, we require the artificial sources used for polarization angle calibration to be able to reach the requirement of $\lesssim 0.1^\circ$. Polarized sources mounted on a tower or a drone can currently achieve $\lesssim 0.5^\circ$ uncertainties in polarization angle calibration and are limited by the local gravity reference to ≈ 20 arcsec. However, further refinements and integrating these technologies with a star camera could open the possibility to achieve the required $\approx 0.01^\circ$ uncertainty, e.g., [542]. The ideal polarization calibrator would measure the polarization angle in the far-field. A drone would enable polarization angle calibration in the SATs far-field and the LAT mid-field. This means that the SAT polarization angle can be more precisely calibrated with artificial sources than the LAT, but additional precision in the LAT polarization angle can be gained by cross-calibrating the LATs and SATs.

Additional relative (SATs and LATs) and absolute (SAT) polarization angle calibration to supplement these measurements could be achieved with a sparse or full wire grids. The balloon-borne Spider experiment used this method to achieve $\approx 0.2^\circ$ precision on absolute polarization angles in pre-flight calibrations.

Pointing: We will use observations of Jupiter, Venus, and the Moon in combination with a star camera to determine the SAT pointing model to < 50 arcsec. At the South Pole, where astronomical sources are only available at low elevations, the full SAT pointing model will be reconstructed from the star camera calibrations. Because they have higher resolution, the pointing model for the LATs will be derived using the point sources detected during nominal observations to < 15 arcsec. The per-detector pointing of both the SATs and the LATs can be further refined by cross-correlating with the CMB.

Time constant: The time constant requirements are set by the beam and pointing requirements. The effects on the beam and pointing decrease for smaller time constants, so faster detectors have less-stringent requirements. For TES detectors, optical time constants are a function of optical loading, which requires that they be measured regularly.

The optical time constants can be calibrated by analyzing beam-smearing from planet observations. We will also consider using a chopped thermal calibrator coupled to a small fraction of the telescope's throughput (as done by SPT), to measure and track changes in optical time constants. Electrical bias steps can also be used to monitor changes in time constant.

Gain: The season-long, focal-plane average detector response can be calibrated by comparison of temperature anisotropy or polarization anisotropy maps with those made by Planck. The detector-to-detector relative gain response and the temporal variation of the gain must be calibrated at the percent level. These can be monitored via regular galactic source and planet calibrations. A variety of other methods can be used to determine relative gain response, including sparse wire grid measurements and elevation nods (modulating the atmospheric signal). We will also track the temporal variations in gain with bias steps. We will also consider using chopped thermal calibrator source signals as an additional option for monitoring temporal gain variations.

Science Analyses

For the purposes of this document, the analysis of CMB-S4 data is divided into two stages: (1) the synthesis of raw detector, pointing, and housekeeping data into single-frequency maps of the sky; and (2) further processing of those maps into data products such as CMB lensing maps, temperature, polarization, and lensing power spectra, and catalogs of galaxy clusters and emissive sources, and using these downstream data products to derive scientific results. In this document, the first stage of analysis is part of the Reference Design under the Data Management subsystem (Sect. 4.7), while the second stage of analysis is referred to as “Science Analysis” and is treated in this chapter (see Fig. 63 for a graphical illustration of this division).

The vast array of CMB-S4 science goals discussed in Chapter 1 will correspondingly necessitate many different types of post-map analyses, and a full accounting of all these types of analyses and the different options for carrying them out would fill a 200-page document by itself. Here we limit ourselves to short discussions of the analyses involved in some of the key CMB-S4 science goals, including the pursuit of a background of primordial gravitational waves, the search for light relics in the early Universe, measuring the mass of the neutrinos, constraining the nature of dark energy, learning about galaxy formation and evolution, and producing legacy catalogs of galaxy clusters and emissive sources. In the following sections, we describe the major analysis steps involved in post-map CMB data analysis and how they relate to these primary science goals, and we discuss major outstanding algorithm choices for certain of the analysis steps. For much more detail on all of these steps, and on the synthesis of raw data to maps, see Chapter 8 of the CMB-S4 Science Book [3].

5.1 Reconstruction of the CMB lensing potential

Crucial to nearly all CMB-S4 science goals will be a maximally accurate and precise estimate of the CMB lensing potential ϕ . The power spectrum of the lensing potential will be the primary CMB-S4 observable used to constrain the nature of dark energy (Sect. 1.3.3) and the sum of the neutrino masses (Sect. 1.3.2). A clean and precise estimate of the lensing potential is also crucial to reaching CMB-S4 goals on the tensor-to-scalar ratio r describing the amplitude of a background of primordial gravitational waves (PGW). Delensing, or using the measured (lensed) CMB temperature and polarization fields and an estimate of ϕ to reconstruct the unlensed fields will be a major part of the PGW analysis pipeline, but it is also a promising avenue to improving constraints from the primary CMB temperature and E -mode polarization power spectra, such as the number of light relic species in the early Universe and the shape of the primordial density perturbation spectrum.

The current state of the art in lensing reconstruction from CMB data is the quadratic estimator method [543], [544]. These quadratic estimators are the first step in an iterative estimation of the true likelihood, and in the weak-lensing limit they are nearly optimal. To harness the full power of CMB-S4 lensing information, however, it will likely be necessary to develop lensing estimators that more closely approximate the maximum likelihood solution. Development of maximum-likelihood algorithms or their equivalent is underway [545, 546], but they have yet to be demonstrated on large CMB data sets.

Another outstanding question, particularly for delensing, is what information to include in the ϕ estimate. The baseline assumption in the CMB-S4 Science Book and in many other publications is that the

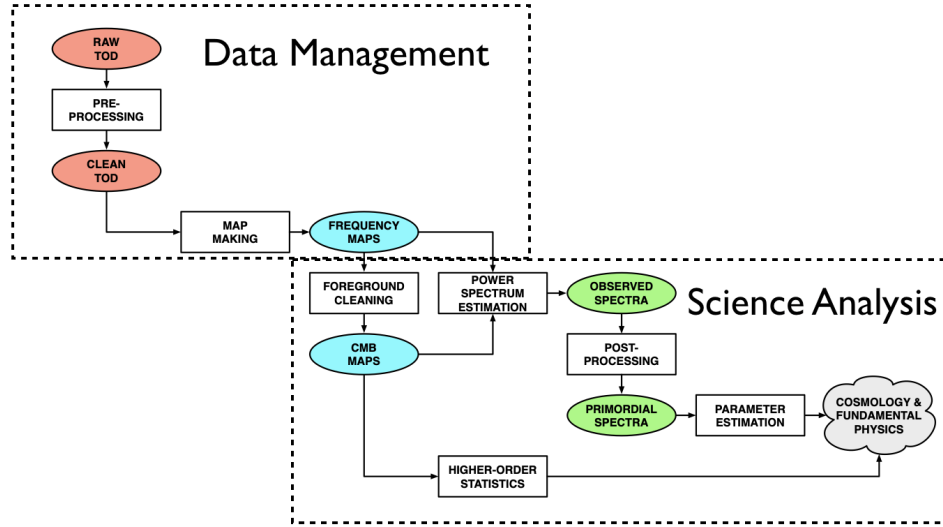


Figure 63. Schematic view of the CMB-S4 data analysis pipeline (from figure 88 of the CMB-S4 Science Book), with boxes illustrating which elements are grouped under Data Management (Sect. 4.7) and which elements are grouped under Science Analyses (this chapter).

ϕ reconstruction for CMB-S4 will come entirely from the combination of CMB E -mode and B -mode polarization, from the so-called EB estimator. This is partially because at CMB-S4 noise levels the EB estimator has the lowest noise, but it is also because any estimators involving temperature are potentially contaminated by foregrounds, particularly at high multipoles. But there is still significant information to be gained from the temperature field, even at CMB-S4 noise levels (cf. figure 43 in Ref. [3]), and recent work has hinted at ways to mitigate foreground contamination in temperature-based lensing estimators [547]. As shown in Ref. [548], information from tracers of ϕ other than the CMB can also improve delensing results, even at CMB-S4 noise levels. One analysis challenge will be to build the optimal combination of all these ϕ estimators that is also robust to foreground contamination.

5.2 Power-spectrum estimation

5.2.1 Power-spectrum estimation methods

Early measurements of CMB temperature anisotropy, with comparatively few map pixels or angular modes measured, often used maximum-likelihood methods to produce maps of the sky (e.g., Ref. [549]) and either a direct evaluation of the full likelihood or a quadratic approximation to that likelihood (e.g., Ref. [550]) to go from maps to angular power spectra. With the advent of the *WMAP* and *Planck* space missions, which would map the entire sky at sub-degree resolution, it became apparent that computing resources could not compete with the $\mathcal{O}(\mathcal{N}^3)$ scaling of the full-likelihood approach (e.g., Ref. [551]). The solution for power spectrum analysis that has been adopted by most current CMB experiments is a Monte-Carlo-based approach advocated in Ref. [552]. In this approach, a biased estimate of the angular power spectrum of the data is obtained by simply binning and averaging the square of the spherical harmonic transform of the sky map. That estimate (known as the “pseudo- \mathcal{C}_ℓ spectrum”) is related to the unbiased estimate that would be

obtained in a maximum-likelihood procedure through the combined effect of noise bias, sky windowing, and any filtering applied to the data before or after mapmaking (including the effects of instrument beam and pixelization). These effects are estimated by “observing” and analyzing simulated data and constructing a matrix describing their net influence on simulated data. This matrix is inverted, and the inverse matrix is applied to the pseudo- C_ℓ s to produce the final data product. Some version of this Monte-Carlo treatment is likely to be adopted for CMB-S4.

5.2.2 Contamination

At CMB-S4 noise levels, as important as raw sensitivity will be control of contamination to power spectrum measurements. This is especially true in the PGW-targeted analysis of degree-scale B -mode polarization, in which the signal of interest is known to be significantly smaller than the two main sources of contamination, Galactic foregrounds and lensed E -mode power (e.g., Ref. [553]). Additional sources of contamination to B -mode searches are intrinsic to observing a finite patch of sky and to pseudo- C_ℓ methods, including ambiguous (not clearly E or B) modes around the edges of the sky patch and spurious B modes introduced in analysis. Mitigation strategies have been developed for all of these contaminants, but it remains to be seen exactly which methods will be most appropriate for CMB-S4 data.

To separate the CMB signal from the contaminating foregrounds, data from multiple bands will be combined, either in a cross-spectrum analysis or by making linear combinations of maps in different bands to produce a “pure-CMB” map for power spectrum estimation. In either case, an underlying model of foreground behavior is assumed—even if that model is simply an assumption regarding the level to which the spectral behavior of foregrounds varies over the sky. There are two challenges related to uncertainties in foreground modeling: one statistical and one systematic. The statistical issue is simply how to propagate the statistical uncertainty on the foreground model to uncertainties on cosmological parameters. In explicitly parameterized foreground models, this happens automatically through the covariance resulting from the fit. For non-parametric models, this covariance can be assessed through Monte-Carlo methods, but making many independent realizations of large-scale Galactic foregrounds is problematic because of the strongly non-Gaussian behavior of these foregrounds.

The primary question for delensing and dealing with spurious or ambiguous B modes is whether to subtract a statistical estimate of their contribution to the B -mode power spectrum or to attempt to construct a phase-ful estimate of the realization of contamination in the actual data. The latter strategy is in general leads to smaller residuals (because there is no sample variance in the estimate of contamination) but is more computationally intensive. For delensing, it has been determined that for a small-area survey which measures a comparatively small number of sky modes to high precision, by-realization cleaning is mandatory, and all the delensing techniques described in the previous chapter are in this category. Spurious/ambiguous B modes can be dealt with either by estimating the statistical bias to the final B -mode spectrum or by constructing a matrix representing the effect of any analysis steps on the true sky [554]. The latter approach involves constructing an $\mathcal{N}_{\text{pixel-by-pixel}}$ matrix, equal in size to the full pixel-pixel covariance, and will not be feasible for high-resolution CMB-S4 data but could be used in analyzing lower-resolution data.

Perhaps the most important aspect of any contamination mitigation is the estimate of the residual contamination after cleaning. If this estimate is inaccurate, it will lead directly to biases in the final determination of the amplitude of the cleaned power spectrum. In the case of PGW searches, this leads directly to a bias in the tensor-to-scalar ratio r . To avoid this bias, all of the contamination-mitigation methods discussed above will need to effectively marginalize over uncertainties in our knowledge of the contamination sources and mechanisms.

5.3 Component maps and cross-correlation with other surveys

It has become increasingly clear, even since the publication of v1 of the CMB-S4 Science Book, that a wealth of exciting science will come from the combination of CMB-S4 data with survey data from other instruments at other wavelengths. These contributions will come both in the area of constraints on cosmological parameters and in increased knowledge of processes on smaller scales and in our local Universe. As shown in Chapter 1, cross-correlations with optical data in particular can lead to improved constraints on the growth of structure, the sum of the neutrino masses, the nature of dark energy, and the primordial power spectrum, as well as informing our understanding of galaxy formation and evolution. Fully exploiting this area of science will require building analysis pipelines jointly with experts in data from other surveys.

Among the challenges and questions in this area of analysis will be whether to perform the analyses in real or harmonic space, how to estimate the covariance of the cross-correlations, and, as in CMB power spectrum analyses, how to deal with contaminating signals and estimate the residual contamination after cleaning. This last point will be key to obtaining precise and accurate determinations of cosmological parameters and constraints on galaxy feedback models through correlations of optical observables with maps of CMB lensing and the tSZ and kSZ effects. In particular, contamination from the tSZ is currently the limiting factor in correlation analyses of optical data and CMB lensing [555, 556], so the foreground mitigation schemes discussed in Sect. 5.1 are also critical to cross-correlation. In general, foreground treatment will be a key aspect of the production of component maps such as the Compton- y map and the CMB blackbody map (see Sect. 1.4.1). As with power spectrum estimation, it will be crucial not only to minimize the contamination in the component maps but also to have accurate estimates of the residual contamination and to propagate these estimates through any cross-correlation analysis involving component maps.

5.4 Parameter estimation

The final step in the analysis of a CMB data set is the estimation of cosmological parameters from the various post-map statistics discussed above. This involves estimating the likelihood of the data given a model parameterized by the standard six Λ CDM parameters, possible extensions of the cosmological model, and any nuisance parameters involving the instrument, foregrounds, and other sources of systematic uncertainty. The current industry standard for this part of the analysis are Markov-chain Monte-Carlo (MCMC) methods, in particular the implementation in CosmoMC [557], and it is expected that CMB-S4 will use similar methods.

5.5 Creating source and cluster catalogs

An additional post-map product of interest for CMB-S4 is the location and properties of compact sources, in particular clusters of galaxies identified through the thermal SZ effect. The standard practice for extracting SZ clusters from multifrequency millimeter-wave maps is through the application of a Fourier-domain spatial-spectral filter [558]. The computational effort involved in this step is small compared to the estimation of power spectra and higher-order correlations, and the algorithms are well-developed and fully implemented for multi-frequency data sets (e.g., Ref. [559, 432])—however, the cluster density could be high enough in CMB-S4 data that approaches more sophisticated than the simple matched filter (e.g., Ref. [560]) could be required to maximize cluster yield.

5.6 Transient and time-domain analyses

One class of analyses that is not represented in Fig. 63 and does not strictly come under the heading of post-single-frequency-map analysis is the area of time-domain and transient astronomy. As discussed in Sects. 1.5.5 and 1.5, CMB-S4 will be a rich data set for detecting mm-wave GRB afterglows, monitoring the light curves of AGN, and discovering new planets. These science goals will require analyses of the time-ordered detector data that are very different from the standard pipeline that will create single-frequency maps of full-depth data. In particular for GRB afterglows, science yield will be maximized by a quasi-real-time alert system linked to the transient alert mechanisms in the wider community. This will require on-site computing and analysis software that runs autonomously.

5.7 Sky simulations

Simulations will be a key aspect of the entire CMB-S4 analysis pipeline, both in the stages handled by the data management subsystem and in post-map analyses. As discussed in Sect. 4.7, we assume here that the software for synthesizing mock skies into simulated time-ordered detector data and maps is part of the data management subsystem, but that the creation of the simulated skies themselves will occur outside of the data management sub-system. Sky simulations that will be useful for CMB-S4 fall into two main categories: (1) simulations of the polarized mm-wave emission from our own galaxy; and (2) simulations of the (mostly) unpolarized emission from extragalactic sources. These present different challenges and will likely come from independent sources.

Comprehensive, accurate, and flexible simulations of Galactic emission will be critical for reaching the CMB-S4 goals in the search for primordial gravitational waves. Key challenges in the generation of simulations of Galactic emission include:

- the level of coherence of diffuse emission across observing frequencies, as any decoherence will limit the efficacy of cleaning a foreground from one observing band using the measurement in a different band;
- the existence or not of a simple parametric emission law for each component emission, such as power laws (for synchrotron) or modified blackbody emission (for dust components);
- the absolute level of foreground emission (in particular for those components that do not scale simply as a function of frequency, such as the superposition of many individual sources with a specific emission law each);
- whether or not emissions for which the level of polarization is unknown or unclear must be modeled and treated for CMB-S4 or can be safely neglected;
- The level at which foregrounds can be treated as Gaussian random fields, which is an assumption of certain foreground cleaning approaches.
- the reliability of models based on observations at angular resolution lower than that of CMB-S4, integrated in broad frequency bands, and with a sensitivity limit at least an order of magnitude worse than what will be achieved with CMB-S4.

Simulations of extragalactic emission will be important for estimating contamination to small-scale power spectra and lensing maps, and they will play a crucial role in validating analysis methods and verifying

results in the area of cross-correlation with surveys at other wavelengths. This requires simulated maps at many wavelengths with the same underlying initial conditions and large-scale structure fluctuations, and with the proper correlations between observables such as CMB lensing potential, tSZ and kSZ, galaxy positions, and galaxy weak-lensing shear. A key challenge in the area of extragalactic simulations will be the competing needs of high accuracy at small angular scales—which can only be achieved with full hydrodynamical simulations—and the ability to create quickly many realizations with different cosmologies and galaxy-formation parameters.

5.8 Implementation

As described in Sect. 4.7, the volume of raw, time-ordered data in CMB-S4 will necessitate both the provisioning of unprecedented computing resources and innovation in CMB data processing algorithms. Once the raw data have been reduced to single-frequency maps, however, the data volume for CMB-S4 will no longer be significantly large than that of *Planck*. Implementation of the standard methods in the field for post-map processing is not expected to overly tax the computing resources of the field. If minimum-variance methods are used for certain post-map analysis steps—for instance the simultaneous estimation of the CMB lensing potential, the unlensed fields, and cosmological parameters as described in [545]—or if full pixel-pixel matrix methods are used for spurious/ambiguous B -mode cleaning, the requirement for computing resources could be higher.

5.9 Validation and verification

The importance of the key CMB-S4 science goals and the challenge of not only achieving the necessary sensitivity but also demonstrating that the final signal is not significantly affected by contamination or bias, compels us to set up a framework in which we can validate any software used for the key science analyses and verify that the algorithms and implementation thereof produce unbiased results with the expected precision on simulated data. This process is already happening in the PGW forecasting working group through a series of data challenges, and we expect to extend the data challenge framework to key results from the large-area, large-aperture surveys in the near future.

Project Overview

6.1 Introduction

The CMB-S4 Collaboration and a pre-Project Development Group of experienced project professionals jointly contributed to the development of a Work Breakdown Structure (WBS), Organization, Cost Book, Resource Loaded Schedule, and Risk Registry. The reference design and project baseline prepared for this document is the basis for subsequent design and project development work to be completed by the Interim Project Office and the Collaboration during 2019–2020. A permanent Integrated Project Office will be established in 2020 to manage the construction phase which starts in 2021.

The CMB-S4 project total estimated cost is currently \$591.6M (fully loaded and escalated to the year of expenditure) including a 35% contingency budget. In-kind contributions delivered by Private and International partners are expected and would reduce the total cost to NSF and DOE. Critical R&D is in progress, funded by DOE. There is a pending proposal to NSF for Design and Development support under the NSF MSRI-R1 program. An eight-year construction project is anticipated (2021–2029) with a transition to operations starting with completion of commissioning of the 1st telescope in 2026. Long lead procurements for the construction project will start in 2021.

6.2 Scope, work breakdown structure, and cost

The CMB-S4 Work Breakdown Structure (WBS) includes twelve major categories as shown in Table 6.1. The distribution of WBS elements by funding agency will be discussed with the Joint Coordination Group established by the funding agencies. As these plans mature there will be a general understanding of the proposed responsibilities of each funding agency and partners providing in-kind contributions. The scope distribution will leverage the capabilities of universities, national laboratories, and industry.

The cost estimate shown in Table 6-1 includes all Materials & Services (M&S) and Salaries, Wages and Fringe Benefit (SWF) costs for the project. The M&S costs and labor resources are estimated at the lowest (task) level in the Project Schedule. The costs in the schedule are given in FY19 dollars. Appropriate overhead and escalation are done external to Primavera, within the Cobra Project cost management and reporting software that will eventually be used to compute earned value. It is foreseen that all Project tracking and reporting will be done using Cobra and Primavera software for the duration of the Project. The cost estimate is the full cost, i.e., does not take credit for contributions from collaborating institutions supported by private and international partners, e.g., Large Aperture Telescopes currently under construction in Chile as part of the Simons Observatory, and Large and Small Aperture Telescopes proposed by international collaborators. The value of in-kind contributions could reduce the total cost of the CMB-S4 project by 20–25%.

Scientific labor resources not in project management roles, i.e., scientific and technical development work, are provided to the Project through research program support outside of the Project. This is the traditional funding model used for DOE and NSF joint projects including LSST, ATLAS, and CMS. The approach

equalizes the treatment of the collaboration’s experimental physicists with respect to their cost to Projects, regardless of their funding source.

The cost contingency estimate was constructed using input from subject matter experts with previous experience in previous CMB experiments and similar NSF MREFC projects and DOE MIE projects. The current estimate of contingency budget or reserve is 35% of the base cost estimate. As the design, cost estimates, and schedules mature the contingency as a percentage of the base cost estimate is expected to decrease to 30% or less. The target range for the start of the CMB-S4 construction project is 25–30%.

The contingency estimate was compared to similar large research infrastructure projects sponsored by the DOE and NSF and is reasonably consistent with other projects at this early stage of project development.

WBS Level 2 Element	Total Cost (\$M)
Total Estimated Cost (TEC)	
1.01 – Project Management	19.6
1.03 – Detectors	39.5
1.04 – Readout	59.9
1.05 – Module Assembly & Testing	31.8
1.06 – Large Aperture Telescopes	86.5
1.07 – Small Aperture Telescopes	52.3
1.08 – Observation Control & Data Acquisition	13.9
1.09 – Data Management	26.9
1.10 – Chile Infrastructure	38.1
1.11 – South Pole Infrastructure	37.0
1.12 – Integration & Commissioning	7.7
Direct TEC	413.2
TEC Contingency (35%)	144.6
Total TEC	557.9
Other Project Cost (OPC)	
1.01 – Project Management	7.0
1.02 – Research & Development	24.2
Direct OPC	31.2
OPC Contingency (35%) – excludes R&D	2.5
Total OPC	33.7
Total Project Cost (TPC)	
TEC + OPC with contingency	591.6

Table 6-1. CMB-S4 WBS structure and cost.

The resource-loaded schedule determines the annual funding profile. The schedule is technically driven starting in 2022, i.e., the technically efficient schedule unconstrained by funding in any year. The resulting funding profile is provided in the following table.

	FY19	FY20	FY21	FY22	FY23	FY24	FY25	FY26	FY27	FY28	Total
OPC	4.3	11.4	17.9	0.0	0.0	0.0	0.0	0.0	0.0	0.0	33.6
TEC	0.0	0.0	7.1	80.8	127.7	184.0	93.5	50.6	13.5	0.5	557.9
TPC	4.3	11.4	25.0	80.8	127.7	184.0	93.5	50.6	13.5	0.5	557.9
Operations	1.0	5.0	10.0	16.0

Table 6-2. *CMB-S4 funding profile.*

6.2.0.1 WBS dictionary

The WBS Dictionary defines scope of each CMB-S4 WBS element as described in Table 6-3.

6.2.0.2 Scope management plan

The scope of the CMB-S4 project is defined to meet the scientific and technical requirements. The reference design and emerging conceptual design meets these requirements as demonstrated by the flow down of science requirements to the proposed project scope. The delivery of the project scope will be time phased and elements of the project scope will be added and deleted as necessary and in accordance with the configuration management plans. Changes to the project scope will be subject to formal change control procedures and approved at the appropriate level of authority. For example, minor changes can be approved by the Level 2 or Level 3 managers, major changes by the Project Director and/or Project Manager, and finally changes affecting the overall scientific performance or key performance parameters will need to be approved by the funding agencies and partners.

Scope contingency options will be identified and managed including the timing of decision points, both for scope reductions and scope restoration or scope additions. A scope management plan will be developed describing this process.

6.2.1 Organization

CMB-S4 is both a scientific collaboration and a nascent DOE/NSF project. While these are certainly tightly-coupled, they do have different roles and responsibilities; the overall organization of CMB-S4 therefore decouples into the organization of the collaboration and the project.

6.2.1.1 Collaboration organization

During the 2013 Snowmass particle physics project planning exercise, the US CMB community came together and recognized:

1. that realizing the extraordinary scientific potential of the CMB would require an increase in the sensitivity of our instruments corresponding to moving from tens to hundreds of thousands of detectors;
2. that the cost of an experiment of this scale would limit us to a single instance, in contrast to the long history of having multiple completing experiments at any epoch;

WBS Element	Description
1.01 – Project Management	Labor, travel, and materials necessary to plan, track, organize, manage, maintain communications, conduct reviews, and perform necessary safety, risk, and QA tasks during all phases of the project. Overall project Systems Engineering is a subsection of this wbs element. However, subsystem-related management and support activities for planning, estimating, tracking, and reporting as well as their specific EH&S and QA tasks are included in each of the subsystems.
1.02 – Research & Development	Labor, travel and materials necessary to support development of a Conceptual Design (pre CD-1 for DOE). Activities include design of detector wafers and the readout systems, data management, optical design, and aspects of the cryostat design for both large and small telescopes.
1.03 – Detectors	Labor, materials, and equipment associated with the design, fabrication and testing of the detector wafers.
1.04 – Readout	Labor, materials, and equipment associated with the design, fabrication and testing of the detector readout system.
1.05 – Module Assembly & Testing	Labor, materials, and equipment associated with the design, parts fabrication, assembly and testing of the detector modules.
1.06 – Large Aperture Telescopes	Labor, materials, and equipment associated with the design, prototyping, materials selection, construction and certification for the Large Aperture Telescope (LAT) System.
1.07 – Small Aperture Telescopes	Labor, materials, and equipment associated with the design, prototyping, materials selection, construction and certification for the Small Aperture Telescope (SAT) System.
1.08 – Observation Control & Data Acquisition	Labor, materials, and equipment associated with the design, construction, certification, and delivery of the control systems for the observatories and data acquisition.
1.09 – Data Management	Labor, materials, and equipment associated with the design, construction, certification, and delivery of the data management system.
1.10 – Chile Infrastructure	Labor, travel, and materials necessary to plan, track, manage, maintain communications, conduct reviews, and perform necessary safety monitoring on site including management of all shipping of CMB-S4 components to Chile and oversight of construction activities on site.
1.11 – South Pole Infrastructure	Labor, travel, and materials necessary to plan, track, manage, maintain communications, conduct reviews, and perform necessary safety monitoring on site including management of all shipping of CMB-S4 components to the South Pole and oversight of construction activities on site.
1.12 – Integration & Commissioning	On Site Integration and Commissioning of the CMB-S4 telescopes and infrastructure in Chile and the South Pole.

Table 6-3. WBS dictionary.

3. that the challenges of fielding such an experiment would primarily be associated with scaling existing technologies (in hardware and data management) to unprecedented levels; and
4. that meeting these scaling challenges would require adding the unique capabilities of DOE laboratories to the long-standing NSF program.

The community therefore proposed CMB-S4 to the 2014 Particle Physics Project Prioritization Process (P5) as single, community-wide, experiment, jointly supported by DOE and NSF. After P5 recommended CMB-S4 under all budget scenarios, in 2015 the CMB community started holding biannual workshops—open to CMB scientists from around the world—to develop and refine the concept. At the request of DOE and NSF, in late 2016 the Astronomy and Astrophysics Advisory Committee (AAAC) convened a Concept Definition Taskforce (CDT) to conduct a CMB-S4 concept study, and the resulting report was unanimously accepted in late 2017.

One recommendation of the CDT report was that the community should organize itself into a formal collaboration, and an Interim Collaboration Coordination Committee was elected to coordinate this process. The resulting draft bylaws were refined at the Spring 2018 community workshop, and overwhelmingly ratified on March 19th 2018, bringing the CMB-S4 collaboration into being, and the first elections for the various officers of the collaboration were completed by the end of April 2018.

Figure 64 shows the organizational structure of the CMB-S4 collaboration. The Governing Board sets policy and provides oversight to an Executive Team led by two equal co-Spokespersons which is responsible for the day-to-day management of the collaboration. A number of Councils, Committees, and Working Groups then carry out the necessary work to enable the overall scientific goals of the collaboration. As of summer 2019 the collaboration has 198 members, 71 of whom hold positions within the organizational structure. These members represent 11 countries on 4 continents, and 76 institutions comprising 16 national laboratories and 60 universities. It should be noted that collaboration members from both national laboratories and universities are engaged in the entire scope of work, and that US institutional affiliation does not automatically map to specifically DOE or NSF scope.

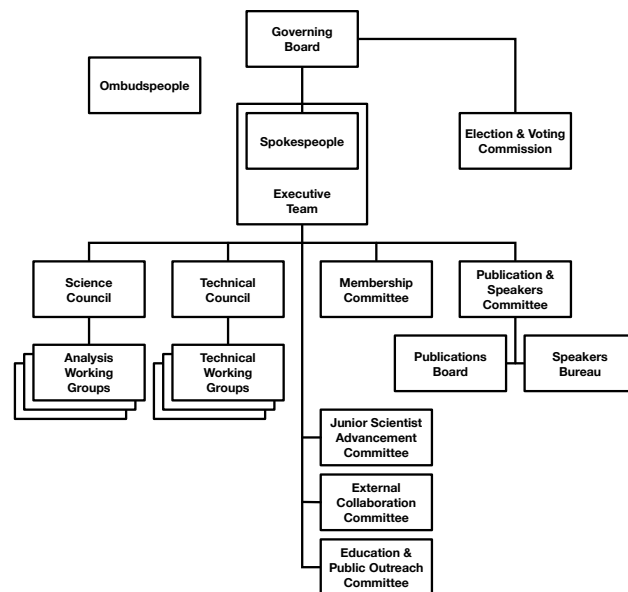


Figure 64. Organizational chart of the CMB-S4 collaboration.

6.2.1.2 Project organization

The CMB-S4 project organization including lines of authority, communication, oversight and advisory committees will be consistent with the expectations for large research infrastructure projects sponsored by the NSF and DOE. This includes a core project office reporting to institutions directly accountable to NSF and DOE for the successful delivery of the CMB-S4 project. The lead institutions will establish oversight committees and will work with the project office to ensure the necessary advisory committees are effective.

The general approach is to develop and deliver the CMB-S4 project with an Integrated Project Office and project organization established by the institutions leading the NSF CMB-S4 MREFC project and the DOE CMB-S4 Major Item of Equipment (MIE) project. The lead NSF institution will be the NSF partner for the MREFC Cooperative Agreement, and the lead DOE institution will be a DOE M&O contractor (National Laboratory) identified by DOE as the lead integrator for the MIE project.

The organization and management approach will adhere to the following principles:

1. one experiment undertaken by a single collaboration and run as one project;
2. joint NSF and DOE oversight and management, currently a Joint Coordination Group (JCG), with a lead agency to be defined prior to approval of the integrated project baseline;
3. lead institutions accountable to NSF and DOE for the MREFC and MIE projects, respectively;
4. lead institutions establish a project governance agreement that engages all major institutional partners, e.g., an Integrated Project Steering Committee (IPSC) comprised of the major institutions delivering project scope including universities, DOE labs, and private and international partners;
5. a single Integrated Project Office established by the lead institutions and their oversight council, e.g., the IPSC, with clear reporting lines to NSF and DOE;
6. clarity in the NSF and DOE scope of work to ensure direct lines of accountability to each agency and clearly defined management interfaces;
7. integrated Project Office prepared agreements for contributions by private and international partners; and,
8. common management systems that meet the needs of both funding agencies, e.g., project cost and schedule tools, systems engineering processes, etc.

Prior to the establishment of the lead institutions, an Interim Project Office (see Fig. 65) will coordinate the development of management plans that adhere to the principles described above and prepare for the project execution phase. The Interim Project Office will report monthly progress on project development efforts and R&D activities to the Joint Coordination Group established by NSF and DOE. The Interim Project Office is led by the Interim Project Director. The transition to the construction phase will include the appointment of a permanent management structure.

The Interim Project Office, with the support of the Collaboration, will continue to mature the experiment design and project execution plans during 2019–2020. The Interim Project Office will transition into a permanent Integrated Project Office in 2020 to prepare for the construction phase which starts in 2021. A key feature of the organization is the role of collaboration members in the Project Office, primarily as leaders of the Level 2 systems. The Level 2 managers are supported by engineering and project management professionals. The NSF/DOE scope distribution will promote the engagement and participation of universities and national laboratories. Graduate students, postdocs, professional technicians and engineers are expected to be involved in all aspects of the project.

6.2.1.3 Private and international partners

The CMB-S4 Project is a collaborative project, with the scientific Collaboration members serving in technical leadership roles in the project. This is similar to many successful projects including IceCube, ATLAS, and




Figure 65. Organizational chart of the interim project office. The figure includes a notional distribution of project scope by funding agency (NSF = blue, DOE = green, Other = yellow). We are actively pursuing partners who could make significant scope contributions in areas aligned with their expertise.

CMS. The lines of accountability for project delivery are clearly defined within the project organization. The project office is responsible for forming partnerships with key stakeholder institutions including DOE National Laboratories, universities, and potential collaborating observatories/projects such as the Simons Observatory, South Pole Observatory, and the CCAT-prime project. Partnerships are also expected to include foreign institutions participating in the CMB-S4 Science Collaboration and contributing to the CMB-S4 Project.

6.2.1.4 In-kind contributions

The CMB-S4 project is expected to include significant contributions from collaborating institutions supported by funding agencies other than NSF and DOE. These “in-kind” contributions will be defined as deliverables to the project. The collaborating institutions agree to deliver items, e.g., instrumentation and effort, required for the success of the CMB-S4 project. The actual cost of each item is the responsibility of the collaborating institution providing the In-kind deliverable. The current best estimate of the value of in-kind contributions is 20–25% or the total cost of the project. This includes both existing infrastructure, telescopes currently under construction, and telescopes and instrumentation proposed by international collaborators.

The CMB-S4 Project Director (PD) is responsible for ensuring the successful delivery of all in-kind contributions required for the CMB-S4 project. The PD is supported in this role by the NSF, DOE, international funding agencies, private foundations, the collaboration and their elected spokesperson, the host institutions, and the Project Office and Level 2 and Level 3 managers. The management approach engages all of these parties in the process of defining and ensuring the delivery of in-kind contributions from partners.

In-kind contributions will be defined in Memoranda of Understanding (MoU) and/or Statements of Work (SoW) executed between the Project Office and the contributing institutions. The MoUs or SoWs define the

in-kind contributions for each collaborating institution including the schedule milestones for the institution's deliverables. The Project Director is responsible for sign off on the in-kind deliverables.

Milestones for the in-kind deliverables will be defined in the CMB-S4 project schedule. Progress against these milestones is evaluated during monthly schedule reviews by the Project Office and Level 2 managers. Large deliverables are the terminal milestone for a sequence of lower level milestones. This procedure includes determining the completion forecast for each milestone and taking corrective action when needed.

The project management approach is the same for the entire project: in-kind scope delivered by partners or scope supported by NSF or DOE. The only difference is in the tracking of actual costs, in-kind actual costs are tracked by the collaborating institution providing the in-kind deliverable and are not tracked by the Project Office.

6.3 Cost, schedule, and risk

The project has developed a task based detailed resource loaded schedule which was reviewed by an external panel of experts in December 2018. The estimate follows the guidance in the NSF Large Facilities Manual, NSF 17-066 and the Project Management for the Acquisition of Capital Assets, DOE 413.3b.

For the task-based estimate, tasks were defined at the lowest level elements of the WBS. The entire ensemble of tasks represents all the required resources, activities, and components of the entire project. Each of the tasks are scheduled and estimated by the teams using accepted techniques. The estimates are documented with a Basis of Estimate (BOE) developed by the subsystems leads and stored in a set of google documents which can be ingested into the PrimaveraTM scheduling program. To facilitate proper integration into the CMB-S4 project control system, standard PrimaveraTM layouts are used to enter the information into the database. The schedule has 1110 activities, 1928 relationships, 5 Level 1, 20 Level 2 and 299 Level 3 Milestones for the CMB-S4 project.

6.3.1 Cost

Each detailed task-based cost estimate corresponds to a task in the project schedule. For that specific task, resources and their quantities are assigned from a standardized list of resources. The list includes multiple resource classes in each of the categories: labor, materials/non-labor, or travel. A task estimate consists of the number of hours of each labor resource class, the base-year dollar cost of each materials/non-labor resource class, the number of trips for each travel resource class, and the basis for each estimate.

6.3.2 Schedule

Table 6-4 shows the proposed the NSF Level 1 Milestones along with the corresponding DOE Critical Decision Milestones and Fig. 66 shows a summary of the schedule and high level milestones. The Interim Project Office, jointly supported by NSF and DOE, will further define the schedule for NSF and DOE reviews and approvals with two guiding objectives, a technically driven schedule and coordinated agency reviews and approvals. This approach is necessary for a single integrated project and the clear delineation of scope and

responsibilities for each funding agency and partner. The critical path is driven by fabrication of detector wafers and delivery of assembled and tested detector modules to the large aperture telescopes.

The schedule development strategy is to define a schedule that is consistent with the funding potentially available during FY2019-FY2021, and subsequently technically driven. The project is working towards an early completion milestone that contains limited schedule float and a year of schedule float following this early project complete milestone is included in the overall project complete milestone (CD-4). The Interim Project Office will continue to optimize the schedule and include explicit float for activities that are not on the critical path. The best opportunity to improve on the schedule is to reduce the time required to deliver the full quantity of the Detectors and Readout (D&R) components. This is a major focus of the R&D program supported by the DOE. The Interim Project Office formed a D&R Task Force in early 2019 to evaluate existing fabrication and testing capabilities and to provide recommendations on production plans. A formal review of the resulting detector fabrication plan will be completed in mid-2019.

NSF Level 1 Milestone (DOE Critical Decision)	Schedule (FY)
Pre-Conceptual Design (CD-0, Mission Need)	Q3 2019
Preliminary Baseline (CD-1/3a, Cost Range/Long-Lead Procurement)	Q3 2021
Preliminary Design Review (CD-2, Performance Baseline)	Q2 2022
Final Design Review (CD-3, Start of Construction)	Q4 2023
Completion of 1st Telescope (CD-4a, Initial Operations)	Q2 2026
Project Completion (CD-4, Operations)	Q1 2029

Table 6-4. *Funding agency milestones.*

6.3.3 Reporting and contingency management

The Project will provide reports on a regular basis to NSF and DOE management. The objective of the reporting is to compile essential technical, cost, schedule and performance data into reports to aid in the monitoring and management of the Project.

All cost account managers (CAMs) will submit monthly written narrative reports to the Project Office detailing specific progress on the pertinent subsystems. These reports summarize the activities of the previous month, describe the activities planned for the upcoming month, and include comments and concerns. In addition, performance reporting, including cost and schedule variance reporting, are submitted by the CAMs. These are collected and summarized in a corresponding monthly report prepared by the CMB-S4 management team. This report outlines progress, problems, and budget and schedule status, including comparisons of projected status versus actual status.

Requests for cost and/or schedule contingency usage will be included in the monthly reports and will generate a change request which must be approved by (depending on the amount requested) the project change control board, project management, laboratory and agency program officers prior to allocation of funds.

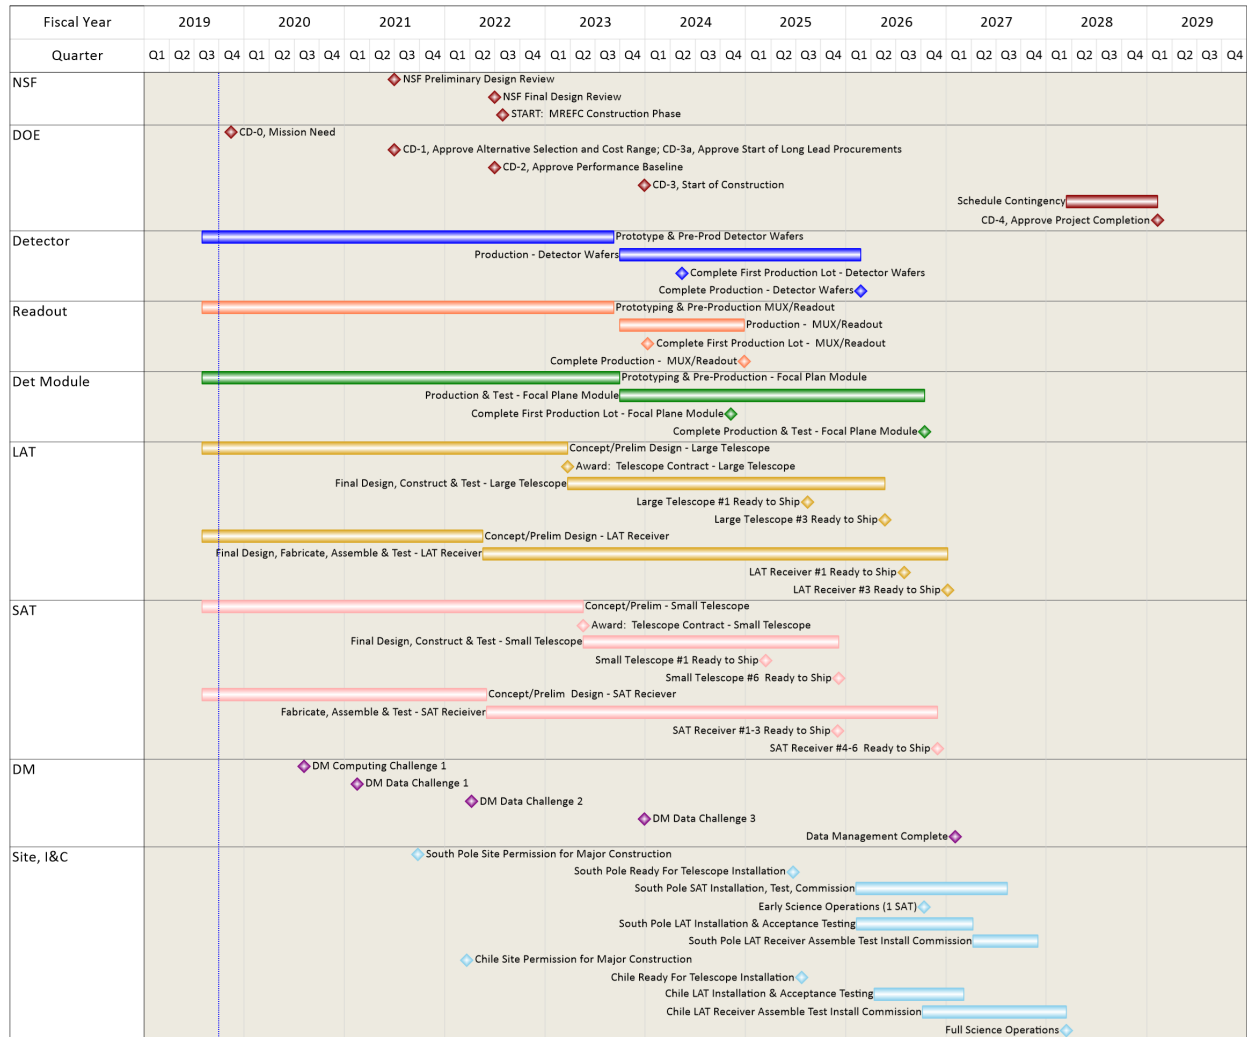


Figure 66. CMB-S4 schedule and milestone summary.

6.4 Risks and opportunities

6.4.1 Risk and opportunity management plan

The CMB-S4 Risk and Opportunity Management Plan, describes the continuous risk and opportunity management (RM) process implemented by the project. RM is a disciplined approach to managing project risks throughout the life cycle of the project. This plan is consistent with DOE O413.3B, “Project Management for the Acquisition of Capital Assets,” and the NSF 17-066, “NSF Large Facilities Manual.” The plan establishes the methods of assessing CMB-S4 project risk and opportunities for all subsystems as well as the system as a whole. Project risk and opportunity are managed throughout the life of the project, from development through construction and commissioning phases.

The primary goal is to manage the risks and opportunities associated with the development and construction of CMB-S4 and focus on understanding, reducing, or eliminating identified risks. Project risks and opportunities are centrally managed, but are the result of project-wide integrated and quantitative assessment which supports management decision-making. The statistical analysis of the residual risk after the planned mitigations informs the project contingency analysis for both cost and schedule.

Current and comprehensive risk updates provide management with additional information in preparing for and reacting to contingent events and adverse outcomes to planned events. The process also provides a uniform language for tracking risk elements and communicating that information. The Risk Registry documents the risk assessment, mitigation strategy, and the residual risk after mitigation. It also includes information about all identified risks within the project. The registry has incorporated lessons learned in several recent projects. Risk Review Board meetings with the project leads will be held on a regular basis to review critical project risks, updates to the registry and status on mitigations. The Risk Register is maintained by the Systems Engineering and the risk management execution is owned by the project management.

The current risk registry for CMB-S4 is being used to define the R&D programs to mitigate risks and to develop the baseline plan for the overall project. A series of risk management meetings were held with each WBS Level 2 system leads and their designated technical experts. The primary purpose of these meetings were to discuss “why” risk management is an essential tool for all levels of project management, as it helps the team communicate and work together to reduce the negative impacts and increase positive impacts. The other goal was the “how” of risk management which included identifying risks, developing informative risk statement, accessing the current probability and impacts of the risks and their possible mitigations. Risk review board meetings will be held regularly to discuss the top risks of the program and update status and changes to the risk registry, as well as action tracking.

6.4.2 Risk/opportunity register

The CMB-S4 risk registry has 140 risks identified and assessed. There are four (4) risks that are currently assessed at critical and 38 risks at high. Table 6-5 provides the current assessment summary for the identified risks and Table 6-6 summarizes the top risks of this program. The project is working on mitigations to ensure that these risks are lowered to reasonable levels that are consistent with our overall project timeline and performance.

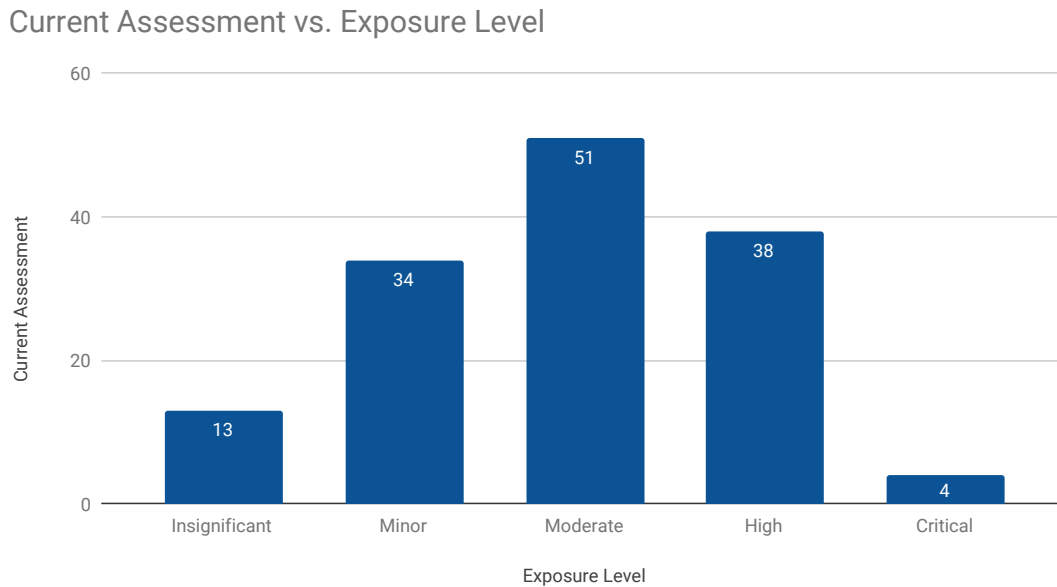


Table 6-5. *CMB-S4 risk registry current assessment exposure level. This table shows the current assessment exposure level of the risks that have been identified.*

In summary, CMB-S4 will require a large increase in the number of detectors and associated readout components and camera optical elements as compared to Stage-3 experiments. Improving reliability and throughput is important. Technology development is being targeted to mitigate cost, schedule, and performance risk, as well as to exploit opportunities to reduce cost.

6.5 Value management—alternative analysis and selection

The work proposed here will be evaluated through a well-established value engineering process. The first of three phases in this process is the Functional Analysis to define the technical scope, including current design decisions in the context of the full project, as well as known requirements (performance, operational, environment, etc.). Each of the functions will be classified as basic or secondary through their connection to the flow down requirements. The second phase is the Alternative Analysis, which establishes which among the space of solutions could best meet project objectives relative to a reference design. For each of these options weighted criteria are defined (i.e. requirements, cost, schedule, production, operations, risks, complexity, flexibility, maintainability, safety, development cost/schedule) that will be used to evaluate the approaches. These considerations inform the weaknesses/risks that will be studied in a prioritized way through the above scenarios, subject to cost and schedule constraints. Through this phase, implications on cost, schedule, risks, and opportunities are evaluated and updated. Finally, the third phase is the Implementation in which the results of the Alternative Analysis with regard to scoring against the weighted criteria will be presented to the Project Office and/or a designated review team to arrive at a decision to either select or reject a design, or to do further work. Through this process, we will manage risk, maximize opportunity, and promote discussion and support for the final decisions.

Risk Description	Current Assessment	Mitigation	Post Mitigated Assessment
IF decrease in fab or performance yield over production period (define criterion) THEN reduces experiment sensitivity and/or increases cost/schedule.	Critical	1) Multiple fab houses with duplicate capabilities hedges risk of any single point failure of affecting sustained yield. 2) Demonstrate sustained yield early, and start fabrication of science-grade wafers after CD-3A. 3) At least 6 m of schedule float to critical path.	Moderate
IF the minimum required foundries can't produce detectors that meet performance at the end of pre-production, THEN impact to cost, schedule and/or performance.	Critical	Investigating two Detector foundries in the DOE complex, other agency foundries and commercial foundries. Holding a review of foundries and supporting prototypes to establish processes.	Moderate
IF detector requirements are not established before the technology downselect, THEN science performance could be impacted.	Critical	Determine flowdown of science to survey to detector/readout requirements. Perform required simulations to support flowdown.	Minor
IF integrated detector module noise performance doesn't meet specification THEN we don't achieve map depth in observing time and/or increases cost/schedule.	Critical	Early development and testing of at least two MUX technologies with orthogonal risks across noise, integration complexity and TRL AND integration testing with sensors End-to-end testing and validation by CD-2/NSF PDR.	Minor
IF 1-2 fully characterized test beds doesn't exist THEN throughput reduction and/or increase cost/schedule.	High	Support construction of testbeds by CD-1/NSF CDR.	Minor
IF pipeline is inefficient on next-generation energy-constrained architectures THEN analysis is resource-limited.	High	1) Demonstrate range of required performance parameters by CD-1/NSF CDR. 2) Fund continued fab effort during pre-CD-0 to characterize processes.	Moderate
IF insufficient quality hires available THEN schedule delay with consequences for experiment design simulations and production pipeline development.	High	Start DM recruiting early.	Moderate
IF wafer performance yield < (eg.80-90%); (yield = Sensor parameter variation {e.g., P _{sat} , R _n , L} meet spec for wafer) THEN reduces experiment sensitivity and/or increases cost/schedule.	High	1) Demonstrate reliability of fab steps and with prototype wafer, 2) Multiple fab houses demonstrate yield statistics on 5-10% of targeted science-grade wafer count, 3) If single step is responsible, explore alternate processes or dedicated machines.	Moderate
If there is residual unmodeled beam shape error THEN polarization systematics too large.	High	1. Science simulations & modeling, 2. New analysis techniques, 3. Better beam measurements so you can correct the problem.	Minor

Table 6-6. CMB-S4 Risk Registry. This table is a summary of the major risks that have been analyzed by the collaboration. A description of the risk and its impact are presented in the first two columns. The 'current assessment' is derived from the product of a probability factor and an impact factor. The planned mitigation action is listed in the third column, followed by a projected post-mitigated risk assessment in the fourth column.

6.6 Operations plan

6.6.1 Transition to operations

Completion of commissioning of the first telescope marks the start of CMB-S4 operations and completion of commissioning of the last telescope is the completion of the construction project. The transition from commissioning to operations will occur over a 2–3 year period and will involve a significant change in staffing at the sites, data centers, and partner institutions. A successful operations phase will require transitioning key personnel from Integration & Commissioning to operations and hiring new staff with skills that are appropriate for operations. A transition plan will be developed to cover this period.

6.6.2 Operations

The basic operations model for CMB-S4 will be observations with multiple telescopes and cameras distributed across two sites, with observing priorities and specifications optimized for the CMB-S4 science goals, and data from all instruments shared throughout the entire CMB-S4 collaboration. Scientists working at laboratories and universities will coordinate the observations, monitor the data, design and implement the data pipeline, and carry out science analyses. Instruments at both sites will collect data nearly continuously. These data will include calibrations as well as CMB observations. Nearly all observations will be automated, so that local operators on the sites will not be needed during routine observations. The two CMB-S4 sites though remote, are both sufficiently well-established that fielding the CMB-S4 instrumentation does not represent a large risk, and costs and schedule estimates can be predicted from past experience. Site-specific planning and management are required because there are many site-specific issues such as power generation, safety monitoring, data storage and transfer. The Chile site is at very high altitude (5200 m, barometric pressure about half sea-level) and must deal with snow removal, but is accessible year round. Staff and visitors are housed in San Pedro de Atacama (2400 m, population 2500) and make the one hour drive up to the site as needed. The South Pole gets very cold (-80 C), is not accessible mid February through late October, and suffers shipping restrictions on the size and weight of parts, but the small winterover staff is housed within 1 km from the telescopes, and is always available. At both sites, major maintenance is generally in summer when the weather is worst for observing.

The operations cost is based on a preliminary bottom-up estimate that includes management, site staff, utilities, instrument maintenance, data transmission, data products, pipeline upgrades, collaboration management, and science analysis. The effort is roughly 30 FTE/year for data products and pipeline upgrades, 15 FTE/year for science analysis, and 25 FTE/year for management and site support. The annual operations cost is \$32M in 2019 dollars, excluding 20 FTE/year of scientist effort supported by DOE research funds. If the DOE-supported scientist effort is included, annual operations are 10% of the construction cost, which is typical for an observatory.

6.7 R&D and pre-conceptual design

The CMB-S4 Collaboration is preparing for CD-1/PDR on the timescale of April 2021. R&D and pre-Conceptual design work is aimed at reducing risk and firming up cost and schedule estimates.

Funding to support these efforts has been provided by DOE and funding has been requested from NSF through a MSRI-R1 proposal.

R&D is focused on Detectors and Readout while the pre-Conceptual Design studies are aimed at 5 distinct areas: Cold Optics, Ground Pickup Sidelobes and Beam systematics, LAT Cryostat Design, SAT Cryocooler test and Data Management.

The interim project office includes an R&D and pre-Conceptual Design manager that works with the collaboration Technical Coordinators (see Fig. 65). Progress on R&D and pre-Conceptual Design work is reported to the manager and to the interim project office on a monthly basis.

A short summary of the R&D and pre-Conceptual Design efforts is given below.

Detectors and Readout: The highest technical risk to the CMB-S4 project was identified by the CDT and the December 2018 review as detector fabrication and testing capacity. ANL and LBNL/HYPRES have already demonstrated the capability to fabricate sinuous antenna detectors that are compatible with fmux readout. The near term R&D will focus on bringing these fabrication facilities up to speed to fabricate another possible combination, OMT horn-coupled detectors that are compatible with umux (and dfMUX) readout, optimized to operate at the planned 100mK focal plane temperature. The long term goal is to have the performance of a CMB-S4 wafer and detector module fully characterized and understood by October 2020.

Cold Optics: CMB-S4 requires more cold optical elements than any previous CMB experiment. These include lenses, filters, half-wave plates (HWP), and baffling materials. The lenses, half-wave plates, and absorbing filters require high quality anti-reflection (AR) coatings. The goal of this effort is to demonstrate implementations of these technologies to enable performance forecasting and technology selection by CD-1.

Ground Pickup & Sidelobes (SATs and LATs): Systematic errors due to ground pickup, and pickup from the Sun, Moon, and galaxy are a critical issue for CMB-S4 and have been identified as one of the most significant project risks. Pickup adds noise and biases in measurements of r and N_{eff} , and must be carefully controlled if CMB-S4 is to meet its ambitious science goals. This effort is focused on modeling pickup for the CMB-S4 large and small telescopes, using a mix of physical optics and ray tracing tools, with inputs and techniques verified by sidelobe measurements from Stage-3 telescopes and laboratory measurements.

LAT Cryostat Design: The key design driver for the large telescope cameras is accommodating enough detectors to achieve the required mapping speed. Existing large camera designs, e.g., the Simons Observatory LAT cryostat, are too small. This effort is aimed at developing a new concept for CMB-S4, including consideration of the optimum pixel spacing, optics tube diameter and spacing, and detector wafer configuration and will develop a preliminary solid model of the new camera concept that will provide a basis for detailed design.

SAT Cryocooler Test: The cooling capacity of 4-K and 40-K stages are the driver for the cryostat design of CMB-S4 SAT. The cryocooling system is the dominant contributor to the power budget at the site. This activity covers a quick test of the cooling capacity for currently available dilution refrigerators and pulse tube systems and will enable an early start to the SAT cryostat design effort.

Data Management: CMB-S4 faces the twin challenges of controlling systematic effects to unprecedented precision and the extraordinary volume of the data to be processed. This effort addresses activities associated with these challenges that are on the critical path to CD-1 either because they are required to inform the baseline design (systematics simulations) or because they have critical milestones on that timescale (data processing at scale and on new computing architectures).

Appendices

A

Science Forecasting

Here we present the forecasts that form the basis of our flowdown from science requirements to measurement requirements. We begin in Sect. A.1 with the forecasts for constraints on the tensor-to-scalar ratio, r . In Sect. A.2 we present the forecasts for constraints on the light relics parameter N_{eff} . Forecasts related to our two design-driving legacy survey goals are presented in Sect. A.3 and A.4.

A.1 Ultra-deep field targeting the degree-scale signature of gravitational waves

Using simulations to optimize the design of a CMB experiment inevitably involves a trade-off between the degree of detail that the simulations are able to capture and the computational (and human) cost of generating and analyzing them. This trade-off includes the choice of domain in which the simulation is generated, ranging from the most detailed but most expensive time domain, through the map domain, to the most simplified, but most flexible, spectral domain. Inclusion of additional detail can help to validate general results, to explore their sensitivity to assumptions about foreground models, sky coverage, and instrumental noise and systematics, and in more mature stages of design, can inform specific instrument and survey strategy choices.

To ensure realism, our general forecast/simulation approach has been an iterative one. We rely on a closed forecasting loop to tie the semi-analytic tools, which allow for fast optimizations, with map-based studies, which can include multiple layers of additional complexity. Our measurement requirements and the baseline experiment configurations which can achieve them are established as a result of multiple passes through this loop. The main steps describing this process are as follows.

1. Develop a (semi-) analytic spectral forecast that makes use of noise performance that is informed by scaling from actual analyses of real experiments from time-streams to power spectra.
2. Use this forecasting tool to optimize the allocation of detector effort across frequencies, determining certain baseline “checkpoints” in survey definition space.
3. Validate these checkpoint configurations with standardized, version-numbered map-based data challenges. If independent analyses show recovery of science parameters from these challenge maps that does not match analytic forecasts (either in terms of variance or bias), we revise the forecasts accordingly.
4. Iterate between steps 1 and 3, injecting increasing realism in the form of: (a) sky model complexity informed by the latest data and modeling efforts; (b) survey coverage based on proven observing strategies; and (c) systematics whose form, parameterization, and likely amplitude is likewise guided by real-world experience.

For the CMB-S4 Science Book [3] an r -forecasting machinery was assembled based on scaling the bandpower covariance matrices and noise spectra of published BICEP/Keck analyses. Given a defined set of bandpasses, and assumptions about foreground power spectra, this semi-analytical approach is capable of optimizing

allocation of detectors across the sky (sky coverage), frequency channel, and delensing vs. degree-scale surveys, for the lowest $\sigma(r)$ at fixed effort, where “effort”= total number of 150-GHz-equivalent detector-years of observation. (It makes sense to define “effort” in these units since it is equivalent to focal plane area, which is in turn the strongest driver of overall project cost.)

For the CDT report we extended this work to map-domain simulations in order to be able to capture additional complexities that cannot be represented in the spectral domain, while remaining computationally tractable. These complexities include:

- non-Gaussianity and statistical anisotropy of the Galactic foregrounds;
- instrumental systematic effects;
- inconsistency between the data and the assumptions (either explicit or implicit) of any given analysis method;
- foreground contamination in the delensing map.

We also used these simulations to validate the spectral domain forecasts for configurations where the approaches are directly comparable.

For this Decadal Survey Report, we have continued the practice of using the semi-analytic approach to explore a wide number of options, while using sets of map-based simulations at specific checkpoints to calibrate the semi-analytic approaches and confirm their validity. Here we review the methods used to explore parameter space for the ultra-deep field, including map level noise simulations, sky models, and observation strategy. We also describe our approach to modeling instrumental systematics, the delensing map, and the analysis methods. We present results of our calculations that guide our flowdown to the measurement requirements.

A strength of CMB-S4 is the access to two different observing sites with complementary capabilities. The sites differ significantly in the range of sky area on which coverage can be concentrated, with the Chilean site allowing for greater sky coverage, and the Pole site allowing for more concentrated sky coverage. We present here forecasts for $\sigma(r)$ as a function of the distribution of SATs across site, forecasts that inform our choice of a baseline configuration, and motivate the flexible approach to deployment that is part of our reference design.

In Sect. A.1.1 we present our flowdown to sky coverage, over-all sensitivity level for both the SATs and the delensing LAT, and for the allocation of detectors across SAT frequencies using our semi-analytic optimization framework. In the next two subsections we build on these results with improved realism in two different ways. In Sect. A.1.2 we present map-based simulations as a cross check on the semi-analytic results, and as an examination of robustness of the forecasted performance given a variety of foreground models as well as systematic errors of instrumental origin. To provide the calculations we need as a basis for our optimization of SAT and delensing LAT siting, we return to our semi-analytic framework. We present these calculations in Sect. A.1.3 for survey coverage maps that include constraints imposed by each site and the finite extent of the instrument field of view.

A.1.1 Flowdown to total number of detector years and allocation of detectors across frequency and delensing effort: semi-analytic calculations

To obtain the optimal allocation of detectors across frequencies, we use the aforementioned performance-based forecasting framework. This semi-analytic tool is grounded in published BICEP/Keck achieved

performances, in the form of bandpower covariance matrices and noise angular power spectra from end-to-end analyses of multiple on-sky receiver years at {95, 150, 220} GHz [553]. For projections we assume that we can scale down the noise based on increased detector count and integration time and that we can apply beam-size and NET rescaling to estimate the achieved performance at other frequencies. This “achieved performance” approach automatically builds in all real world inefficiencies, including (but not limited to) imperfect detector yield, non-uniform detector performance, read-out noise, observing inefficiency, losses due to timestream filtering, beam smoothing, and non-uniform sky coverage. A detailed presentation of the framework and optimization process is given in Ref. [561].

To span the four available atmospheric windows (Fig. 67) and have enough channels to mitigate against complex foregrounds, we assume eight channels at {30, 40, 85, 95, 145, 155, 220, 270} GHz, which are placed on small aperture telescopes. In addition, we also include a 20-GHz channel on a large aperture telescope. This latter inclusion is the result of insight gained from an early iteration through the forecasting loop, which demonstrated that for certain foreground models sizeable biases were present due to synchrotron residuals. The CDT strawperson design was updated accordingly to mitigate against such biases, with the 20-GHz channel being placed on a large aperture telescope due to resolution constraints.

The procedure used to come up with the split in each window was to separate the overlapping bands as far as possible while still keeping the calculated per-detector NET within 10–15% of the NET for a detector that spans the full window. The ideal per-detector NETs were calculated with `NETlib.py`¹ at Pole and Chile, using the 10-year MERRA2 median atmospheric profiles. We use the mean over the two sites, which are {214, 177, 224, 270, 238, 309, 331, 747, 1281} $\mu\text{K}_{\text{CMB}}\sqrt{\text{s}}$ for our nine channels. These NETs are calculated for a 100 mK thermal bath, as opposed to 250 mK for the Science Book, and are therefore lower. Note that this is the only departure from achieved performance. We want to emphasize that these NET numbers are only used to determine the appropriate scalings between different channels to allow current achieved performance numbers to be applied to proposed CMB-S4 instrument configurations, and not to calculate ab-initio sensitivities. The scaling procedure is presented in detail in Ref. [561].

The optimization process specifically includes the need to delens and assigns a fraction of the detectors for that purpose. We assume a separate high-resolution instrument dedicated to measuring the intermediate- and small-scale information necessary to construct a template of lensing B modes, so that their effect can be removed. In this initial semi-analytic optimization process, the delensing instrument is assumed to have 1-arcminute resolution and detector weight at a single frequency with mapping speed equivalent to that of the 145-GHz channel. The translation between detector effort and map noise in the delensing instrument is based on the method used for the low-resolution instrument, but without certain non-idealities specific to low-resolution instruments and low- ℓ analysis (such as mode removal and non-uniform coverage) [561]. Following the formalism in Ref. [562], we convert the map noise in the delensing map to a delensing efficiency, or equivalently a fractional residual in lensed B-mode power, as is shown in Fig. 68.

The trade-off between raw sensitivity, ability to remove foregrounds, and ability to delens results is a complicated optimization problem with respect to sky coverage. Figure 69 (right), shows the r sensitivity forecast for CMB-S4 as a function of the observed sky fraction for the case that we only have an upper limit ($r = 0$). We note that for an initial detection the optimization process requires a deep survey that targets as small an area as possible. This conclusion of course depends on the forecasting assumptions; to that end we would like to draw attention to several key factors. First, holding the desired constraint on r fixed, the level to which we rely on delensing to decrease sample variance increases appropriately at smaller sky fractions. For example, as shown in Fig. 69 (right), achieving the forecasted sensitivity on r for a field targeting 1% of the sky will require an $> 80\%$ reduction in the map rms level of the CMB lensing B modes. While from a sensitivity standpoint it is possible to achieve these levels, the extent to which systematic effects and small-

¹cmb-s4.org/wiki/index.php/New_NET_Calculator_and_Validation

scale foregrounds will need to be constrained may become too stringent. Second, the current optimization assumes identical foreground behaviour across the sky (equivalent to that in the BICEP2/Keck region), while in reality the average amplitude, and possibly the complexity, of foregrounds increase as larger sky areas are targeted. This effect would steepen the optimization curve at high sky fractions and increase our preference for small sky, but it is important to be cautious until we know more about foregrounds at these sensitivity levels. Third, a practical consideration for the robustness of the final r result is its reproducibility across the sky. It is therefore useful to observe multiple roughly 1% patches from which we can derive separate cosmological constraints. Finally, the technical aspect of E/B separation heavily disfavors patches smaller than about 1% of the sky due to cut-sky effects. Balancing the forecasting results with these concerns, we have chosen $\approx 3\%$ as the default sky fraction for CMB-S4 r constraints (assuming a true value of $r = 0$).

For this choice of sky fraction, we find that we need $\approx 1.8 \times 10^6$ 150-GHz-equivalent detector-years (or $\approx 1.2 \times 10^6$ under more optimistic foreground assumptions) to reach our science requirements, as can be seen in Fig. 69. Roughly 30% of this effort is dedicated towards the delensing portion of the observations, yielding a 30% rms lensing residual, with the rest of the effort dedicated towards degree scale component separation. The specific optimal allocation across frequencies and delensing effort and corresponding map depths are shown in Fig. 70. This configuration provides a starting point which leads to the reference design.




Figure 67. Calculated atmospheric brightness spectra (at zenith) for the South Pole at 0.5mm PWV and Atacama at 1.0mm PWV (both are near median values). Atmospheric spectra are generated using Ref. [563]. The tophat bands are plotted on top of these spectra, with the height of each rectangle equal to the band-averaged brightness temperature using the South Pole spectrum.

A.1.2 Cross-checks with map-based simulations

In this section, we present the next step in our iterative forecasting, where we generate and re-analyze map based simulations. As described earlier in this Appendix, this allows us to cross-check the results of the semi-analytic calculations, and also to probe the impacts of foregrounds and instrumental systematic effects.




Figure 68. Forecasted lensing A_L residual (grey scale plus colored contours as labeled) using the EB -only iterative delensing [562], as a function of the beam full width half maximum and noise level in Q and U .




Figure 69. Top panels: forecasted uncertainty on r as a function of effort (left) and sky fraction f_{sky} (right). The left panel is for 3% sky fraction, whereas the right panel is for 1.8×10^6 detector years of effort, as represented by the vertical dashed lines. We included in solid black the case of full delensing, while allowing for decorrelation of the foregrounds, in solid grey the case without delensing, in dotted grey the case where no decorrelation is allowed in the model, and in dashed black the raw sensitivity in the absence of foreground and lensing. Bottom panels: For the delensed case, we show the fraction of effort spent on removing the lensing sample variance and the resulting rms lensing residual.

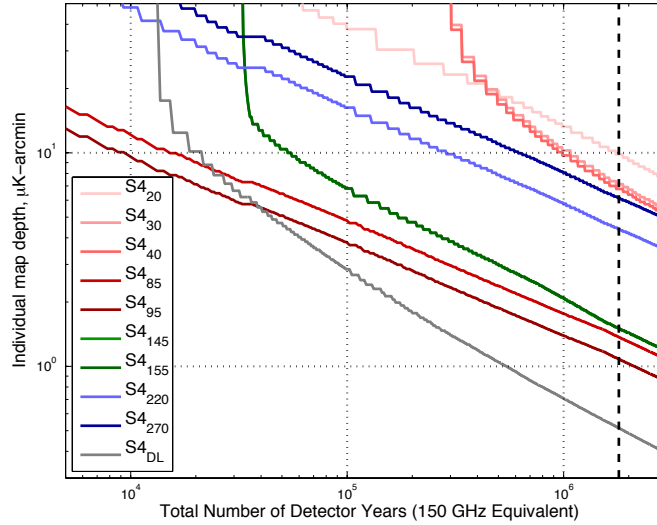


Figure 70. Optimized map-depth in each of the small-aperture channels as well as in the delensing channel, for an $f_{\text{sky}} = 3\%$.

A.1.2.1 Map noise realizations

To produce map-level simulations it is necessary to translate the BICEP/Keck noise bandpowers into a prescription for map noise. We do this by fitting the $N_{\ell S}$ to a white + ℓ^γ model accounting for beam smoothing, etc. For the small-aperture BICEP/Keck data, we find $\ell_{\text{knee}} = 50\text{--}60$ with γ of -2 to -3 . To translate to map noise levels, we must pick a specific sky hit pattern. As we have argued in Sect. A.1.1, we want a patch with 3% area, we use here an idealized circular hit pattern which has this effective area (shown in Fig. 71). We then generate Gaussian noise realizations at each band and divide by the square-root of the assumed coverage pattern such that the noise “blows up around the edge” as it does in real maps.

A.1.2.2 Foreground models

To make simulated sky maps we add realizations of lensed CMB both without and with an r component to models of the Galactic foregrounds. We used the following seven foreground models.

0. Simple Gaussian realizations of synchrotron and dust with power-law angular power spectra at amplitudes set to match the observations in the BICEP/Keck field, and simple uniform SEDs (power law for synchrotron, modified blackbody for dust).
1. The PySM² model **a1d1f1s1**, where the letters refer to anomalous microwave emission, dust, free-free and synchrotron respectively, and the numbers are the base models described in Ref. [564].
2. The PySM model **a2d4f1s3**, where the models have been updated to variants that are also described in Ref. [564]. Note that these include 2% polarized AME, a curvature of the synchrotron SED, and a two-temperature model for dust.

²https://github.com/bthorne93/PySM_public

3. The PySM model `a2d7f1s3`, where the dust model has been updated to a sophisticated physical model of dust grains as described in Ref. [565]. This model is interesting in that it does not necessarily conform to the modified blackbody SED.
4. The dust in model 3 is replaced by a model of polarized dust emission that incorporates HI column density maps as tracers of the dust intensity structures, and a phenomenological description of the Galactic magnetic field as described in Ref. [566]. The model is expanded beyond that described in the paper to produce a modest amount of decorrelation of the dust emission pattern as a function of frequency motivated by the analysis of *Planck* data in Ref. [567].
5. A toy model where the dust decorrelation suggested in figure 3 of Ref. [567] is taken at face value ($\mathcal{R}_{80}^{217 \times 353} = 0.85$). While such a model is not ruled out by current data it appears to be very hard to produce such strong decorrelation in physics-based models. We also note that Ref. [568] have re-analyzed the same *Planck* data and, while they find that the high level of decorrelation in this model is consistent with the data, their best fit to that same data has no decorrelation.
6. A model based on MHD simulations [569] of the Galactic magnetic field, which naturally produces non-Gaussian correlated dust and synchrotron emission.

Models 1 to 4 use the actual large-scale modes of the real sky as measured above the noise in the *Planck* data. This means that these models are intrinsically “single-realization,” and this must be borne in mind when interpreting the results. Models 4 and 6 are not based on *Planck*, but still contain a fixed signal realization. Models 0 and 5 have different seeds for each signal map and include the (Gaussian) sample variance. The PySM models fill in the small-scale structure with power-law Gaussian extrapolations, while models 4 and 6 naturally produce non-Gaussian small-scale structure. However, all of these models are consistent with current data, and we should be careful not to necessarily associate nominal sophistication with greater probability to more closely reflect reality.

A.1.2.3 Delensing

We have started to generate high-resolution simulated maps on which we can run explicit lensing reconstruction and then include that information in the analysis. However, that process is not yet converged, and so for the present we approximate delensing by scaling down the Λ CDM lensing signal. As described in A.1.1, the delensing efficiency at power spectrum level for a coverage of 3% is predicted to reach $\approx 90\%$, so in the map-based results presented here, we use an effective A_L of 0.1.

A.1.2.4 Instrumental systematics

Control of instrumental systematics is a critical design consideration. However, predicting and modeling these effects realistically is a difficult task that is dependent on actual instrument and survey design details, and in any case their impact on an actual result comes not through the modeled effects but through unmodeled residuals. So far we have simulated various generic classes of additive systematic by injecting additional noise-like components into the maps, and then re-analyzing them without knowledge of what was put in. We have experimented with components that are both correlated and uncorrelated across frequency bands, and which have white, $1/\ell$, and white + $1/\ell$ spectra, at varying levels compared to single-frequency map noise or, for correlated cases, combined map noise. Examples of mechanisms that might produce map residuals within this class, after modeling them and either correcting or filtering their leading-order effects, include bandpass mismatches, beam and pointing variations, calibration variations, cross-talk effects, half-wave-plate leakages, ground pickup, and readout irregularities.

A.1.2.5 Analysis methods

To make simulated maps the noise realizations described in Sect. A.1.2.1 are added to the sky models described in Sect. A.1.2.2, and potentially also the systematics realizations described in Sect. A.1.2.4. For each realization one then has a stack of multi-frequency $I/Q/U$ maps containing non-uniform noise, foregrounds and signal, and the challenge is to re-analyze them to recover the parameter of interest (in this case r). This can be done by different teams using different methods, and could be done in a blind manner, although we have not done this yet.

So far we have experimented with two methods. The first is a map-based ILC cleaning method [e.g., 570], which seeks the linear combination of maps that minimizes the remaining CMB signal, followed by a marginalization over residual foregrounds. This method has the advantage that it does not need to know the bandpasses of the frequency channels.

The second method is an evolution of the parametric multi-component fit to the ensemble of auto- and cross-spectra as used for the BICEP/Keck analysis to date [488, 553]. This method fits the observed bandpowers to a model composed of the lensing expectation plus dust and synchrotron contributions and a possible r component. Dust and synchrotron each have an amplitude (A_d and A_s), a spatial spectral parameter (α_d and α_s), and a frequency spectral parameter (β_d and β_s). We also allow dust/synchrotron correlation (ϵ), and decorrelation of the dust patterns over frequency (Δ_d).

Both of these analysis methods are only close to optimal when the foreground behavior is close to uniform across the observing field. For analysis of larger fields, algorithms that fit, for example, the frequency spectral indices individually in (large) pixels, will be required.

A.1.2.6 Results

Table A-1 summarizes the results of re-analysis of simulations of 1.2×10^6 150-GHz-equivalent detector-years and residual lensing power $A_L = 0.1$ as described in Sect. A.1.1. We see that for $r = 0$ the simple Gaussian foreground model 0 gives $\sigma(r) \approx 5 \times 10^{-4}$, as expected from the semi-analytic calculations. As we progress to the more complex foreground models, $\sigma(r)$ is generally in the range $5\text{--}8 \times 10^{-4}$. The bias remains below 1σ in all cases. (These simulations are sets of 500 realizations, so the statistical uncertainty on the bias is $\approx 0.04\sigma$.) The strong decorrelation model 5 does significantly increase $\sigma(r)$, and, while the parametric method is able to resist bias in this case, by construction information is lost. In fact if one believed in such a scenario, re-optimization to concentrate the sensitivity at closer-in frequencies would be called for.

Table A-2 summarizes the results of re-analysis of simulations including additive systematic effects, in different combinations of uncorrelated and correlated contamination with varying spectra, added on top of foreground model 3. The levels of systematic contamination for these simulations were chosen to predict biases on r of $\approx 1 \times 10^{-4}$ in semi-analytic forecasts. We can see that the different combinations explored increase biases on r by amounts that typically vary from $0.5\text{--}1.5 \times 10^{-4}$ for the two different analyses, over the different cases. We find that to restrict bias on r to this level, the sum of additive contamination effects needs to be controlled to 3–7% of the single-frequency survey noise, or (in the case of correlated systematics) 6–11% of the total combined noise levels. Such percentages are consistent with the upper limits currently achieved for residual additive systematic contamination compared to survey noise by small-aperture experiments [e.g., 5]. Assuming CMB-S4 will include a sustained effort to continue to control, understand, and model systematic effects down to levels limited by survey noise, these percentages provide reasonable benchmark requirements.

r value	Sky model	ILC		Parametric	
		$\sigma(r) \times 10^4$	r bias $\times 10^4$	$\sigma(r) \times 10^4$	r bias $\times 10^4$
0	0	4.4	-0.2	5.7	0.3
	1	4.6	0.8	6.4	5.2
	2	4.7	0.7	6.5	1.9
	3	4.6	1.2	6.7	0.7
	4	6.5	4.8	8.3	-7.7
	5 ^a	18	17	15	0.2
0.003	6	4.8	-1.8	6.5	1.8
	0	6.6	-0.7	8.1	0.4
	1	6.9	0.9	8.5	5.4
	2	6.5	-0.1	7.9	1.9
	3	7.0	1.4	8.7	0.9
	4	11	7.1	11	-6.2
	5 ^a	23	17	17	0.4
	6	7.5	-0.2	8.6	2.5

^a An extreme decorrelation model—see Sect. A.1.2.2. The parametric analysis includes a decorrelation parameter. No attempt is made in the ILC analysis to model decorrelation.

Table A-1. Results of two analysis methods applied to map-based simulations using our suite of sky models. All simulations assume an instrument configuration including a (high-resolution) 20-GHz channel, a survey of 3% of the sky with 1.2×10^6 150-GHz-equivalent detector-years, and $A_L = 0.1$, as described in Sect. A.1.1.

Systematic	Uncorrelated		Correlated		ILC		Parametric	
	A [%]	B [%]	A [%]	B [%]	$\sigma(r) \times 10^4$	r bias $\times 10^4$	$\sigma(r) \times 10^4$	r bias $\times 10^4$
None	0	0	0	0	5.3	0.0	7.2	0.0
Uncorrelated white	3.3	0	0	0	6.0	0.84	8.0	0.63
Uncorrelated $1/\ell$	0	6.8	0	0	5.0	0.99	7.0	0.85
Correlated white	0	0	5.8	0	6.3	1.2	7.3	1.4
Correlated $1/\ell$	0	0	0	11	5.2	1.0	6.7	0.97
Uncorrelated white + $1/\ell$...	1.6	3.5	0	0	5.6	0.89	7.5	0.76
Correlated white + $1/\ell$	0	0	2.9	5.3	5.5	0.98	6.9	1.0
Both, white + $1/\ell$	0.8	1.7	1.5	2.6	5.6	1.1	7.9	0.98

Table A-2. Results from re-analysis of map-based simulations containing systematics. We report sky model 3 and $r = 0$, with additive systematic effects in varying combinations, the amplitudes of which are specified as percentages of survey noise, for the white (A) and $1/\ell$ (B) components. (These results are for a slightly different configuration with 1.0×10^6 150-GHz-equivalent detector-years and a low-resolution 20-GHz channel.)

Results of simulating systematic errors in the determination of bandpasses vary by analysis method. The construction of the ILC method makes it largely insensitive to such uncertainties. The parametric analysis, which includes specific models of the frequency spectra of different foregrounds, shows biases on r at the 1×10^{-4} level for uncorrelated random deviations in bandcenter determination of 0.8%, or for correlated deviations of 2%, which we adopt as reasonable benchmark requirements to accommodate a variety of both blind and astrophysical foreground modeling approaches.

A.1.3 Flowdown to distribution of SATs and delensing LAT across sites

The semi-analytic optimization over frequency bands described in Sect. A.1.1 made simplistic assumptions regarding the effect of varying the observed sky area. In this section we describe updates to the framework which attempt to take into account the impact of realistic observing strategies, as well as a slightly more conservative approach to our delensing forecasts.

A.1.3.1 Calculation of noise levels

In reality one is not able to choose the number of detectors in each frequency band in the continuously variable manner shown in Fig. 70 above. For the reference design a realistic mapping of detectors into dichroic optics tubes has been carried out while seeking to maintain the band distribution as determined in the optimization calculations—this results in the configuration described in Sect. 3.3. We then scale the BICEP/Keck noise bandpower covariance matrix in the same way as described in Sect. A.1.1 according to the number of detector-years and ratio of NETs. A further re-scaling is then applied to account for sky coverages, as explained in the next section.

A.1.3.2 Sky coverage effects

The semi-analytic calculations of Sect. A.1.1 assumed a simplified re-scaling for sky area, while the map based simulations of Sect. A.1.2 assumed an idealized circular sky patch which is not actually achievable with a practical instrument from a site at any latitude. Figure 71 compares our prior assumptions to more realistic hit patterns. From Pole it is possible to concentrate the coverage onto a compact region of sky, but from Chile one has to observe different regions as the Earth turns, resulting in a more extended coverage area. The large instantaneous field of view of the SAT telescopes means that there is minimum field size which can be achieved, and also that there is always a strong “edge taper” in the coverage pattern.

We have performed a calculation which attempts to optimize simulated SAT observations from Chile to produce the densest possible coverage on a $\approx 3\%$ patch of low foreground sky resulting in the pattern shown in the figure as “Chile full.” We segment this into its deepest part, which we call “Chile deep,” and the remainder, which we call “Chile shallow.”

From Pole one can scan the same patch 24/7 with the size of the observed patch basically controlled by the length of the scan throw in Right Ascension. A minimal length scan results in the pattern shown in the figure as “Pole deep.” Lengthening the scan while remaining in low foreground sky results in the pattern “Pole wide.” In the results below “Pole deep” and “Pole wide” are therefore “either-or” options.

Because the noise increases in regions with less observing time the effective sky area for noise is larger than the effective sky area for signal—and both of these also depend on the weighting applied when analyzing

the maps. The patterns shown in Fig. 71 have the effective sky fractions reported in Table A-3, assuming inverse noise variance weighting.




Figure 71. *Detector-second hit patterns on the sky for small aperture telescope surveys. top: the actual BICEP3 2017 hit pattern, middle left: idealized circular pattern as used in Sect. A.1.2, middle right: simulated “Chile full” pattern, bottom left: simulated “Pole wide” pattern, and bottom right: simulated “Pole deep” pattern. Each pattern is normalized to the same sum and the color scales are equal. (The “Chile deep” and “Chile shallow” regions referred to in the text are sub regions of the “Chile full” pattern.)*

We can take account of the above effects by re-scaling the BICEP/Keck bandpower covariance matrices (BPCMs) in a more sophisticated manner. First, we need to scale the noise due to distributing the effort on a patch of sky larger than the original BK one. The noise is scaled by the effective noise factor

$$f_{\text{eff}}^{\text{noise}} = \frac{\Omega_{\text{pix}}}{4\pi} \frac{\sum_i w_i^2 h_i^{-1} \sum_i h_i}{\sum_i w_i^2}, \quad (\text{A.1})$$

where Ω_{pix} is the solid angle of a single pixel, w_i are the weights for pixel i , and h_i are the hit counts.

	Pole deep	Pole wide	Chile deep	Chile shallow	Chile full
$f_{\text{sky}}^{\text{noise}}$	2.9	6.5	3.4	20	18
$f_{\text{sky}}^{\text{signal}}$	1.9	4.3	2.4	10	5.9
$f_{\text{sky}}^{\text{non-zero}}$	5.0	12	5.0	47	52

Table A-3. Effective sky fractions for signal and noise in % for the observation patterns shown in Fig. 71, and the case of inverse noise variance weighting (i.e., equations A.2 with $w_i = h_i$). We also report the fraction of sky with non zero coverage, $f_{\text{sky}}^{\text{non-zero}}$.

Second, since we observe a different number of modes, we need to scale the signal, noise and signal-cross-noise contributions of the BPCM by the factors

$$f_{\text{sky}}^{\text{noise}} = \frac{\Omega_{\text{pix}} (\sum_i w_i^2 h_i^{-1})^2}{4\pi \sum_i w_i^4 h_i^{-2}}, \quad f_{\text{sky}}^{\text{signal}} = \frac{\Omega_{\text{pix}} (\sum_i w_i^2)^2}{4\pi \sum_i w_i^4}, \quad f_{\text{sky}}^{\text{cross}} = \frac{\Omega_{\text{pix}} \sum_i w_i^2 \sum_i w_i^2 h_i^{-1}}{4\pi \sum_i w_i^4 h_i^{-2}}. \quad (\text{A.2})$$

We also need to take out the effect of these factors from the original BK BPCM. In the BK analysis, the weights are the inverse noise variance, i.e., $w_i = h_i$. In the case of CMB-S4, we will never be noise dominated, either due to an actual primordial signal, or due to the lensing and foreground residuals, so here we use the inverse variance of the total signal and noise to determine the optimal weighting.

A.1.3.3 Delensing forecasts

The CMB-S4 science goals can only be achieved if the majority of the lensing B modes can be removed. The optimization in Sect. A.1.1 assumed a single frequency channel assigned to the higher resolution delensing observations. The strength of polarized foregrounds at the relevant angular scales is currently poorly constrained by data. The reference design therefore includes some additional coverage at higher and lower frequencies.

To forecast the delensing performance, we proceed in two steps. For a given LAT configuration and sky coverage, we derive the noise levels for an internal linear combination (ILC) that minimizes the variance of components with a frequency dependence that differs from that of a blackbody [571]. In this step, we assume that polarized foreground emission is dominated by Galactic synchrotron and thermal dust emission. Using the ILC noise power spectrum, we then forecast the performance expected for iterative EB delensing [562], as shown in Fig. 68 above.

For the reference design for the Chile LATs described in Sect. 3.3 and a wide area survey covering 70% of the sky, the two-step procedure predicts that 73% of the lensing power can be removed in the ‘‘Chile shallow’’ region after seven years of observation. Similarly, for the single LAT at the South Pole dedicated to delensing of the approximately 3% ‘‘Chile deep’’ and ‘‘Pole deep’’ regions, we expect to be able to remove close to 90% of the lensing power after 7 years of observation.

The numbers given above assume inverse noise variance weighting. For the reference design inverse noise variance weighting is typically suboptimal, and in all the forecasts presented below we employ weights that account for both signal (e.g., for $r = 0$ lensing residual after foreground removal) and noise. For the same survey, this leads to slightly higher noise and lensing residuals, but reduced $\sigma(r)$ due to reduced sample variance.

A.1.3.4 Results

The covariance matrices calculated as described above are used to produce the results given in this section where the number, siting and coverage patterns of the SATs are varied. In all cases a delensing LAT at South Pole is assumed concentrating its coverage on a small patch of sky, while delensing over larger sky coverage is assumed to be available from the Chilean LATs.

As mentioned earlier, we split the Chilean coverage shown as “Chile full” in Fig. 71 into a deep patch “Chile deep,” where it overlaps with the “Pole deep” region, and call the remainder “Chile shallow.” We then make separate predictions for each region (with their very different delensing levels) and add the $\sigma(r)$ results in simple inverse quadrature, thereby making the small approximation of independence of the measured modes. When we combine with Pole observations, we mimic a joint analysis by taking the sum of the “Pole deep” and “Chile deep” coverage maps, and computing the corresponding weights and lensing residuals.

Since some parts of the “Chile shallow” coverage lie closer to the galactic plane, we boost the foreground level by a factor of about 3 with respect to the deep patch. This is based on analysis of maps with data driven foreground spatial variations. We also apply Galactic cuts on top of the coverage maps, based on *Planck* polarization data. We focus here on two cases, one where we mask out areas that are not part of the 58% cleanest part of the full sky, and one where we use the 28% cleanest.

In Figs. 72 and 73, we show the dependence of $\sigma(r)$ on r for the different coverage masks. We calculated these constraints for $r = 0, 0.003, 0.01$ and 0.03 , and show the linear interpolation between these points, for the different galactic cuts. In figure 74, we show how different sitings of a total of 18 tubes change our constraints on r . Note that for $r \lesssim 0.003$ “Pole deep” is always better than “Pole wide” so we use the former in the results below when considering smaller values of r .

In Tables A-4–A-8 we present a set of $\sigma(r)$ results for seven years of observations and $r=0$, while varying the number of optics tubes in Pole and Chile over a wide range. We show results for five different variants: (i) with no decorrelation parameters and observing efficiency in Chile equal to that at Pole for the 28% cleanest polarized sky; (ii) assuming observing efficiency in Chile is half that at Pole; (iii) assuming 1% unmodeled foreground residual uncertainty; (iv) assuming foreground decorrelation parameters are required in the re-analysis; and (v) keeping the 58% cleanest sky. The results are moderately degraded in each of the first three variants.

While the forecast variant with 50% observing efficiency from Chile has no impact on the sensitivity of the reference design, which locates all the small aperture telescopes at the South Pole, it reinforces the preference for Pole-centric configurations. Despite long histories of CMB observations at both sites, it is still quite difficult to make a clean comparison of their observing efficiencies. Results from the BICEP/Keck program are responsible for leading constraints on r for the last decade, but it is not possible to disentangle the role of the observing site from other factors that contribute to the success of that program, such as detector performance, instrument design, observing strategy, and operations management. However, note also that even if we assume equal observing efficiency at both sites (Table A-4), there is still a difference between $\sigma(r)$ obtained for equivalent numbers of optics tubes in Chile vs Pole, due to sky fraction and delensing.

Note that the results where the total number of optics tubes is significantly different to eighteen are subject to a caveat. The delensing effort is assumed to be held fixed in these calculations—one LAT at South Pole and two in Chile. In principle as the total effort it varied away from the reference design one should re-optimize the fraction of delensing effort as per Fig. 69.




Figure 72. Uncertainty on r as a function of the value of r , for the different hit patterns shown in Fig. 71. These results correspond to seven years of observation, assuming no decorrelation parameters and observing efficiency in Chile equal to that at Pole. Each band spans the range from a galactic cut which retains the cleanest 58% of the sky to one which retains the cleanest 28% (based on *Planck* polarization data). Note that we split the “Chile full” pattern into the deepest part (in green) that overlaps with the “Pole deep” map, thus delensed by the Pole delensing survey, and a shallow part that is delensed by the Chilean LATs (in gray). The combined “Chile full” (in blue) is the quadratic sum of these two.




Figure 73. Uncertainty on r as a function of the value of r , for 18 tubes, for “Pole deep,” “Pole wide” and “Chile full.” The upper left is a subset of Fig. 72. In the right panels, we turn on dust decorrelation ($\Delta_d = 0.97$, where Δ is defined in equation (F4) of the BK15 paper [553]). (Here we use a quadratic ℓ -dependence, $g(\ell) = (\ell/80)^2$, as we did in the CDT report and the science book.) In the lower panels, we add a foreground bias in quadrature to $\sigma(r)$. Its value is 1% of the equivalent r of the current foreground minimum of the BK15 data at $\ell = 80$. As in the previous figure, these results correspond to seven years of observation and observing efficiency in Chile equal to that at Pole, and each band shows different galactic cuts, based on *Planck* polarized foregrounds: the upper edge keeps the cleanest 28% of the full sky, whereas the lower edge keeps the 58% cleanest.

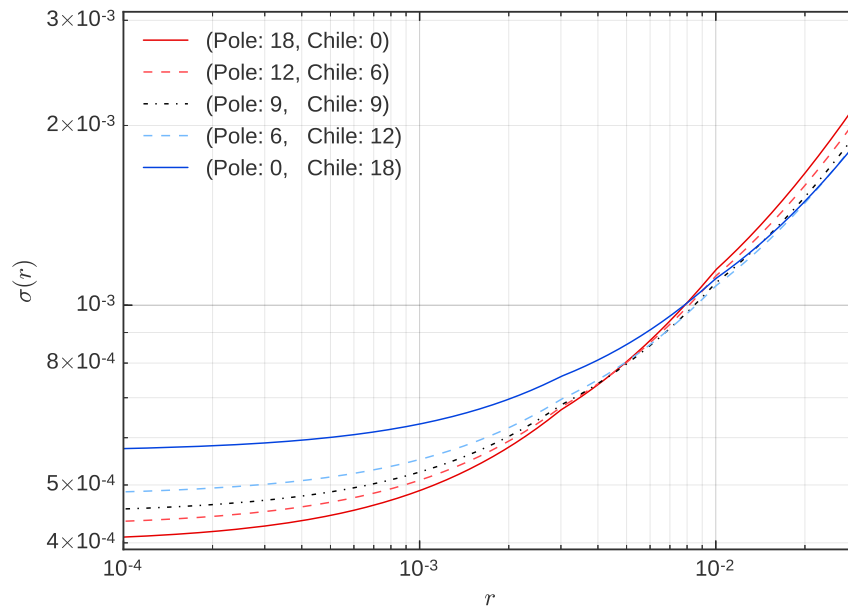


Figure 74. Uncertainty on r as a function of the value of r , for 18 optics tubes, and various splits between Pole and Chile siting. As in the previous figures, these results correspond to seven years of observation and observing efficiency in Chile equal to that at Pole. For clarity, we only show the forecast using the 28% cleanest polarized sky. Note that in the case where all the tubes are at Pole, we use the “Pole wide” pattern for $r \geq 0.01$ since it yields better constraints.

In Tables A-9–A-13 and Tables A-14–A-18 we convert the results into 95% confidence limit and detection significance under the scenario that $r = 0$ and $r = 0.003$ respectively, and show the same four variant cases. For these statistics, we assume the bandpowers have the model expectations values and then calculate the posterior probability distribution for r , marginalized over all foreground parameters. We use the posterior obtained in the map based analysis for different masks to compute a functional form that depends on the noise residuals and the value of r . We scale these posteriors based on the noise sky fraction factors from A.2. This distribution is significantly non-Gaussian, leading to detection significance that is higher than we would naively calculate from $\sigma(r)$. Combined detection significance for Deep and Shallow surveys is calculated by simply multiplying the posteriors. Note that we do not quote these results in the main body of the report as we have not yet taken into account the impact of expected fluctuations in the measured bandpowers. We plan to take these into account with simulations in future work.

A.1.4 Conclusion

Here is a high-level summary of the key points regarding the survey targeting the signature of gravitational waves.

- The semi-analytic calculations described in Sect. A.1.1 indicates that for a 3% sky fraction 1.8×10^6 150 GHz-equivalent detector-years of observation are required to reach the science requirements, with 30% of this assigned to the delensing observations. This calculation provides an optimal distribution of these detectors across frequency bands to achieve the goal in the presence of foregrounds.
- The map based simulations described in Sect. A.1.2 confirm the $\sigma(r)$ results from the semi-analytic calculations. These simulations also indicate that bias in the recovered r value is within 1σ for a suite of different foreground models. However, we note that foregrounds remain a serious issue which must be periodically revisited as the project progresses.
- Additional map based simulations indicate that systematic bias on r can be controlled to $< 1\sigma$ provided fractional contamination levels similar to those already achieved by small aperture telescopes can be maintained.
- Mapping the requirements from the semi-analytic calculations onto realizable instruments results in the reference design described in Sect. 3.3 for a seven-year survey period.
- A larger fraction of the sky can be observed from Chile, but, due to the rotation of the Earth, one can concentrate the available sensitivity more deeply from the South Pole (see Fig. 71). In Sect. A.1.3 we have extended the semi-analytic calculations to account for realistic observation patterns and probed the dependence of $\sigma(r)$ on r , finding that Pole is always favored in the limit of small r , with the cross over point depending on the specific assumptions (see Figs. 73 and 74).

Chile\Pole	0	6	9	12	18	30
0		6.3	5.0	4.5	4.0	3.5
6	12	5.5	4.7	4.3	3.9	3.5
9	8.7	5.1	4.5	4.1	3.8	3.4
12	7.1	4.8	4.3	4.0	3.7	3.4
18	5.7	4.4	4.1	3.8	3.6	3.3
30	4.4	3.9	3.7	3.6	3.4	3.2

Table A-4. Combined $10^4 \times \sigma(r)$, assuming $r=0$ after 7 years of observation, keeping only the 28% cleanest part of the sky, assuming no decorrelation and observing efficiency in Chile same as at Pole.

Chile\Pole	0	6	9	12	18	30
0		6.3	5.0	4.5	4.0	3.5
6	22	5.9	4.9	4.4	3.9	3.5
9	15	5.6	4.8	4.3	3.9	3.5
12	12	5.4	4.6	4.2	3.9	3.5
18	8.8	5.1	4.5	4.1	3.8	3.4
30	6.2	4.6	4.2	3.9	3.7	3.3

Table A-5. Same as Table A-4, but assuming 50% Chilean efficiency.

Chile\Pole	0	6	9	12	18	30
0		7.0	5.9	5.5	5.1	4.7
6	13	6.3	5.6	5.2	4.9	4.6
9	9.2	5.9	5.4	5.1	4.9	4.6
12	7.8	5.7	5.2	5.0	4.8	4.5
18	6.5	5.3	5.0	4.9	4.7	4.5
30	5.4	4.9	4.8	4.7	4.5	4.4

Table A-6. Same as Table A-4, but assuming 1% unmodeled foreground residual uncertainty.

Chile\Pole	0	6	9	12	18	30
0		8.4	6.7	6.0	5.2	4.4
6	16	7.3	6.2	5.6	5.0	4.3
9	12	6.8	5.9	5.4	4.9	4.3
12	9.7	6.4	5.7	5.3	4.8	4.2
18	7.8	5.8	5.3	5.0	4.6	4.1
30	6.0	5.1	4.8	4.6	4.3	4.0

Table A-7. Same as Table A-4, but assuming additional foreground decorrelation parameters.

Chile\Pole	0	6	9	12	18	30
0		6.3	5.0	4.5	4.0	3.5
6	11	5.4	4.6	4.2	3.8	3.5
9	8.0	5.0	4.4	4.1	3.7	3.4
12	6.6	4.6	4.2	3.9	3.6	3.3
18	5.3	4.2	3.9	3.7	3.5	3.2
30	4.1	3.6	3.5	3.4	3.2	3.1

Table A-8. Same as Table A-4, but assuming we keep the 58% cleanest part of the full sky

Chile\Pole	0	6	9	12	18	30
0		14	11	9.7	8.6	7.6
6	26	12	10	9.1	8.3	7.4
9	18	11	9.6	8.8	8.1	7.3
12	15	10	9.2	8.6	7.9	7.2
18	12	9.3	8.6	8.2	7.7	7.1
30	9.3	8.1	7.8	7.5	7.2	6.8

Table A-9. Combined $10^4 \times 95\%$ C.L., for $r=0$ after 7 years of observation, keeping only the 28% cleanest part of the sky, assuming no decorrelation and observing efficiency in Chile same as at Pole.

Chile\Pole	0	6	9	12	18	30
0		14	11	9.7	8.6	7.6
6	46	13	10	9.4	8.5	7.5
9	31	12	10	9.2	8.4	7.5
12	24	12	9.9	9.1	8.3	7.4
18	19	11	9.5	8.8	8.1	7.3
30	13	9.6	8.8	8.3	7.8	7.1

Table A-10. Same as Table A-9, but assuming 50% Chilean efficiency.

Chile\Pole	0	6	9	12	18	30
0		15	13	12	11	10
6	27	13	12	11	11	9.9
9	20	13	11	11	10	9.8
12	16	12	11	11	10	9.7
18	14	11	11	10	9.9	9.5
30	11	10	10	9.9	9.6	9.3

Table A-11. Same as Table A-9, but assuming 1% unmodeled foreground residual uncertainty.

Chile\Pole	0	6	9	12	18	30
0		18	15	13	11	9.5
6	34	16	13	12	11	9.3
9	25	14	13	12	10	9.2
12	21	14	12	11	10	9.1
18	16	12	11	11	9.8	8.8
30	13	11	10	9.8	9.2	8.5

Table A-12. Same as Table A-9, but assuming additional foreground decorrelation parameters.

Chile\Pole	0	6	9	12	18	30
0		14	11	9.7	8.6	7.6
6	24	12	9.9	9.0	8.2	7.4
9	17	11	9.3	8.6	8.0	7.2
12	14	9.8	8.9	8.3	7.7	7.1
18	11	8.8	8.2	7.8	7.4	6.9
30	8.5	7.6	7.3	7.1	6.8	6.5

Table A-13. Same as Table A-9, but assuming we keep the 58% cleanest part of the full sky

Chile\Pole	0	6	9	12	18	30
0		3.7	4.5	4.9	5.4	6.0
6	2.2	4.3	4.8	5.2	5.6	6.1
9	3.0	4.6	5.1	5.4	5.8	6.3
12	3.6	4.9	5.3	5.6	5.9	6.4
18	4.4	5.3	5.7	5.9	6.2	6.6
30	5.4	6.0	6.2	6.4	6.6	6.9

Table A-14. Combined detection significance for $r=0.003$ after 7 years of observation, keeping only the 28% cleanest part of the sky, assuming no decorrelation and observing efficiency in Chile same as at Polee.

Chile\Pole	0	6	9	12	18	30
0		3.7	4.5	4.9	5.4	6.0
6	1.3	4.0	4.6	5.0	5.5	6.0
9	1.9	4.2	4.8	5.1	5.5	6.1
12	2.4	4.3	4.9	5.2	5.6	6.1
18	3.0	4.6	5.1	5.4	5.8	6.3
30	4.0	5.2	5.5	5.8	6.1	6.5

Table A-15. Same as Table A-14, but assuming 50% Chilean efficiency.

Chile\Pole	0	6	9	12	18	30
0		3.5	4.1	4.4	4.8	5.2
6	2.2	4.0	4.4	4.7	5.0	5.3
9	2.9	4.2	4.6	4.8	5.1	5.4
12	3.4	4.5	4.8	5.0	5.2	5.5
18	4.0	4.7	5.0	5.1	5.4	5.6
30	4.7	5.2	5.3	5.4	5.6	5.8

Table A-16. Same as Table A-14, but assuming 1% unmodeled foreground residual uncertainty.

Chile\Pole	0	6	9	12	18	30
0		3.0	3.6	4.0	4.5	5.1
6	1.7	3.5	3.9	4.3	4.7	5.2
9	2.3	3.7	4.1	4.4	4.8	5.3
12	2.8	3.9	4.3	4.6	4.9	5.4
18	3.4	4.3	4.6	4.8	5.1	5.6
30	4.2	4.8	5.0	5.2	5.5	5.8

Table A-17. Same as Table A-14, but assuming additional foreground decorrelation parameters.

Chile\Pole	0	6	9	12	18	30
0		3.8	4.5	5.0	5.3	6.2
6	2.4	4.4	5.0	5.4	5.8	6.3
9	3.3	4.9	5.3	5.6	6.0	6.5
12	4.0	5.2	5.6	5.9	6.2	6.7
18	4.9	5.8	6.1	6.3	6.6	7.0
30	6.1	6.6	6.8	7.0	7.2	7.5

Table A-18. Same as Table A-14, but assuming we keep the 58% cleanest part of the full sky

A.2 Deep and wide field for measuring light relic density

The impact of a change to the light relic density N_{eff} is most prominent in the damping tail region of the CMB power spectra. The sensitivity of a CMB survey to N_{eff} is therefore driven by its ability to map small scale fluctuations. Measuring the relevant modes requires a low-noise, relatively high-resolution, and wide-coverage survey of temperature and polarization. The impact of foregrounds limits the constraining power of the temperature auto spectrum, and so improvements in constraints on N_{eff} tend to be driven by measurements of the TE spectrum. Exponential suppression of the primary CMB temperature and polarization on small angular scales favors observations of as much sky as possible, for fixed total observing time.

The CMB-S4 Science Book [3] explored constraints on N_{eff} over a wide range of experimental configurations. Many of the qualitative lessons learned from that study remain unchanged. However, for the present effort, we have significantly updated and improved the treatment of atmospheric effects, extragalactic foregrounds, and component separation.

In Fig. 75, we show how the forecasted constraints on N_{eff} are affected by changes to the noise level, beam size, and sky fraction. As can be seen from the figure, reduction of the noise level, either through increased detector count or extended observing time, leads to improved constraints on N_{eff} , though significant improvements require fairly large improvements to the map depth of the survey. This is due to the rapid drop in CMB power in the damping tail on small angular scales and to the impact of residual foregrounds, which act as additional noise. Compared to the reference design, increasing the size of the telescope dish leads to a modest improvement in constraints on light relics, while reducing the dish size produces a slightly sharper drop in the constraining power. Observing a larger fraction of the sky allows access to more independent small-scale modes, and results in tighter constraints on N_{eff} . The dependence on sky fraction in Fig. 75 is shown at fixed map depth, but we will explore in much more detail below how the survey design affects light relics constraints.

Given the results shown in Fig. 75, increasing the sky fraction seems to be the most promising route to reduce the constraints on N_{eff} . At fixed observing time, surveying a larger sky fraction leads to an increased noise, though constraints on light relics still favor as wide a survey as possible (due to the rapid fall off of the damping tail). The preference for larger sky fraction is demonstrated in Fig. 76, which for a simple modeling of the effects of atmosphere and foregrounds shows how the error on N_{eff} is affected by changes to the observed sky fraction at fixed total effort. Observing a larger sky fraction from the ground requires telescopes to point to lower elevations, and thus to observe through a greater column of atmosphere for portions of the survey. We have therefore performed a more detailed study of the fraction of sky that can be realistically observed from Chile, taking into account the additional atmospheric loading when observing at lower elevations.

A.2.1 Opportunistic scheduler for sky coverage

In order to determine the map area and depth achievable from Chile, we built a set of observing schedules using an opportunistic scheduler for the observations. First we tiled the sky in celestial coordinates with $10^\circ \times 20^\circ$ (RA \times Dec) tiles that overlap by half a tile in each direction. Then we ran the scheduler with three choices of minimum observing elevation: 30, 40 and 50°. We required a 30° avoidance region around the Sun and the Moon. The three schedules were run for each of two scan strategies designed with and without

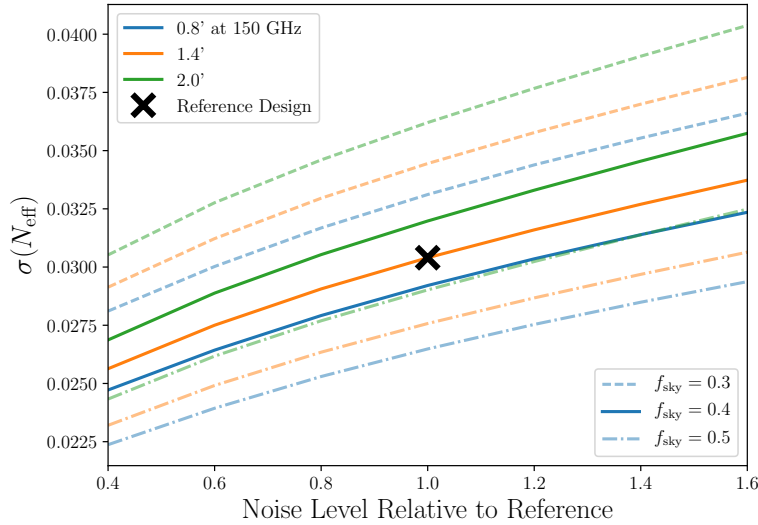


Figure 75. Impact of changes to the noise level, beam size, and sky fraction on forecasted 1σ constraints on N_{eff} with Y_p fixed by BBN consistency. Changes to f_{sky} are taken here at fixed map depth. The forecasts shown in this figure have less detailed modeling of atmospheric effects and foreground cleaning than those shown elsewhere. The results should therefore be taken as a guide to how various experimental design choices impact the constraining power for light relics, but the specific values of the constraints should be taken to be accurate only at the level of about 10%.

an additional “elevation penalty” which tells the scheduler to favor high elevation observations over low elevation observations. We refer to the strategy without the elevation penalty as the “nominal schedule.”

We ran the scheduler on the full sky and subsequently masked pixels inside the galaxy from the resulting hit maps. This strategy was chosen to avoid undesirable boundary effects, since adjusting the tile priority based on galaxy overlap made the final hit distribution around the masked area very uneven. The even-coverage approach led to slightly lower overall observing time but higher effective sky fraction.

Hits were binned into separate maps based on boresight elevation, to facilitate subsequent processing with elevation-dependent noise models.

A.2.2 Atmospheric modeling and depth maps

The elevation-binned hit maps were converted to full survey depth maps in each frequency band as follows. For each elevation bin, an elevation-dependent noise model was used to convert observing time per unit area to map depth in units of $[\mu\text{K-arcmin}]^{-2}$. The per-elevation depth maps were then summed to produce a single depth map, in units of $[\mu\text{K-arcmin}]^{-2}$, for each frequency and each choice of minimum elevation. Intra-tube correlations were captured by binning into cross-frequency depth maps (e.g., 90×150 GHz). This computation is done for the white noise level and at $\ell = 1000, 2000, 3000$ to capture the impact of $1/f$ noise.

The noise model in T includes white noise and atmospheric $1/f$ components. The parameters describing the $1/f$ noise are based on low- ℓ TT spectra measured with ACTPol at 90, 145, and 225 GHz. These power spectra are measured from maps (and thus are expressed in $[\mu\text{K-arcmin}]^2$), but are referenced to ACTPol

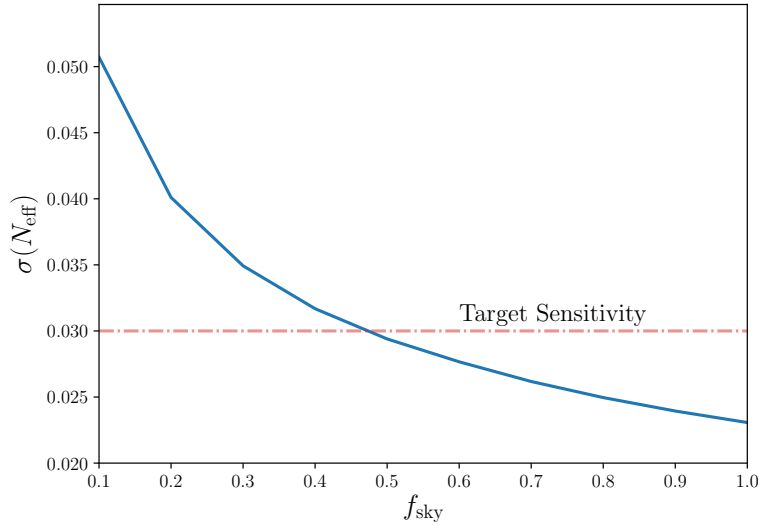


Figure 76. Impact of changes to the sky fraction at fixed effort on forecasted 1σ constraints on N_{eff} with Y_p fixed by BBN consistency. The forecasts shown in this figure have less detailed modeling of atmospheric effects and foreground cleaning than those shown elsewhere and should be taken to be accurate only at the level of about 10%.

array NETs to produce effective atmosphere equivalent noise powers, in $\mu\text{K}^2\text{s}$, that can be combined with the CMB-S4 white noise NETs in each band.

In general, the rescaling of the $1/f$ noise spectrum to a different instrument is complicated because of spatial correlations in the atmospheric contamination. However, our angular scales of greatest interest are smaller than the ACTPol array diameter (0.8°) and in that regime it is reasonable to expect the atmospheric noise power to average down according to the ratio of ACTPol and CMB-S4 focal plane areas. To extrapolate the $1/f$ noise power to other CMB-S4 bands, we use assume that variation of water vapor content is the primary driver of small scale atmospheric noise and use the AM³ atmosphere modeling code v9.2 (median parameters for Atacama, September-November) to obtain a relative calibration between bands. As the ACTPol noise at 90 GHz is 40% higher than predicted by this modeling, we inflate all lower frequency bands by 40% as well.

To add elevation dependence, we assume that the $1/f$ power scales as the square of the airmass. The scaling of the white noise level with elevation is computed separately, using the detector noise model and typical atmospheric loading provided by the AM model.

The polarization $1/f$ noise is not dominated by spatially correlated atmosphere. We assume in this case that $1/f$ noise scales with the detector white noise, with a knee fixed at $\ell = 700$. This is the same assumption made in Simons Observatory LAT forecasting [49].

A.2.3 Foregrounds and point-source removal

For the deep and wide field forecast, the sky model consists of a set of auto- and cross- power spectra for the set of instrument bands (similar to Ref. [572]). It is built from power spectrum templates and source

³<https://www.cfa.harvard.edu/~spaine/am/>

Minimum Elevation	50 degrees				40 degrees				30 degrees			
Galactic Cut (%)	30	20	10	0	30	20	10	0	30	20	10	0
Nominal Schedule												
Sky Fraction	0.39	0.45	0.52	0.57	0.48	0.55	0.63	0.68	0.54	0.62	0.69	0.76
Effective Sky Fraction for Noise	0.35	0.41	0.47	0.52	0.44	0.51	0.58	0.64	0.50	0.57	0.65	0.71
Effective Sky Fraction for Signal	0.32	0.37	0.43	0.47	0.41	0.47	0.54	0.59	0.47	0.53	0.61	0.66
With Elevation Penalty												
Sky Fraction	0.39	0.45	0.52	0.57	0.48	0.55	0.63	0.68	0.54	0.62	0.69	0.76
Effective Sky Fraction for Noise	0.34	0.40	0.47	0.51	0.43	0.49	0.56	0.61	0.46	0.53	0.61	0.66
Effective Sky Fraction for Signal	0.31	0.36	0.42	0.46	0.38	0.44	0.50	0.55	0.40	0.46	0.53	0.58

Table A-19. Sky fractions for various choices of minimum observing elevation and galactic cut for the two scan strategies discussed in the text. The meaning of each type of sky fraction is the same as discussed in Sect. A.1.3.2.

count models. In addition to the CMB, the foreground signals include tSZ, kSZ, Galactic cirrus, radio point sources, and dusty point sources. We scale these with a single frequency dependence per component to translate the power spectra between the frequency bands of CMB-S4, which we treat as delta functions at the central frequencies.

The tSZ and kSZ power spectral templates come from the hydrodynamic simulations of [397]. The tSZ frequency dependence uses the non-relativistic formula, and the kSZ effect has the same frequency dependence at the CMB. The Galactic cirrus template comes from the treatment of [572]. It is a power law $\propto \ell^{-0.7}$, using the amplitude they found for the clean portions of the sky. The frequency dependence is $\propto \nu^{3.8}$ in flux density units.

Radio point-source counts come from [444] at 148 GHz (their model C2Ex). We scale them to other bands as a power law with index -0.5 in flux density units. Dusty point-source counts come from [573] at 217 GHz, but we adjusted them on the faint end to match the SPT 220-GHz source counts [447]. The dusty sources are scaled from 217 GHz with a graybody spectrum with $\beta^{\text{CIB}} = 2.1$ and $T^{\text{CIB}} = 9.7$ K. The power spectral templates for the point sources are flat with an amplitude that depends on the integrated, squared flux density, weighted by the source counts. We model the correlation between dusty sources and the tSZ using a halo model as in Ref. [574, 127].

We estimate the residual point-source power after masking while accounting for the instrument noise and beam. In each band, and based on the total power spectrum of the sky model and the noise model, we compute the variance after applying a filter optimized for point-source detection. The variance is a function of the prospective flux cut, which sets the power of the point-source component. We compare the filtered variance to the amplitude of a filtered point source to determine the signal-to-noise ratio of sources (as a function of the flux cut). We choose the flux cut so that it self-consistently excludes sources with signal-to-noise ratio greater than 5.

For each source population, we identify the best band for excluding sources by scaling the per-band flux cuts back to the reference frequency and comparing them. For radio sources, 90 GHz finds the intrinsically faintest sources due to a combination of the band's noise, beam, and SED. For dusty sources, 270 GHz is the best. Assuming these deepest source cuts determine the source mask, we finally compute the overall level of unmasked residual source power from each population and add it to the total sky model.




Figure 77. Forecasted 1σ constraints on N_{eff} for different choices of the minimum observing elevation as a function of the size of the galactic mask shown for the nominal scan strategy discussed in the text.

A.2.4 Fisher forecasts

For the forecasts, we begin with the multi-frequency model described in subsection A.2.3 and noise spectra based on the depth maps A.2.2. A harmonic space ILC algorithm [571] is used to derive foreground reduced temperature and polarization spectra and noise spectra. The component separated noise curves and extragalactic foreground residuals computed from the procedure above were then used in a Fisher forecast to determine the expected constraints on N_{eff} . The CMB-S4 noise was combined with a model for the noise from the *Planck* satellite, in an inverse variance sum. Doing so reduces the noise on the large scale temperature modes which would be contaminated by the atmosphere for CMB-S4 alone. Contamination from extragalactic residuals was then added, which acts much like an additional source of noise. We imposed a cut at $\ell_{\text{min}} = 30$, and we took $\ell_{\text{max}} = 5000$ for all spectra. A model for *Planck* temperature data was included for $\ell < 30$ on $f_{\text{sky}} = 0.8$.

We assumed a cosmology described by $\Lambda\text{CDM} + N_{\text{eff}} + \sum m_\nu$, assuming BBN consistency to fix the primordial helium density Y_p . The TT , TE , EE , and $\phi\phi$ spectra were included, with the lensing reconstruction noise calculated using the minimum variance combination of quadratic estimators [543, 544], including the improvement from iterative EB reconstruction [575, 562]. The small improvements which come from delensing T and E spectra [576] were also included.

The resulting noise curves were then used in a Fisher forecast to compute the errors on N_{eff} . The sky fractions for CMB-S4 observations were calculated for each choice of minimum elevation, galactic cut, and scan strategy as shown in Table A-19. *Planck* observations were assumed to cover the rest of the sky which is not observed by CMB-S4 and also lies outside the galactic mask. As can be seen in Table A-19, the elevation penalty reduces the usable effective sky fraction which leads to the slightly weaker constraints compared to the nominal scan strategy without elevation penalty. The results for the forecasts for the nominal schedule are shown in Fig. 77.

A.3 Mapping baryons: angular resolution and allocation of detectors across frequency for the LATs

Mapping the pressure, temperature, and distribution of baryons requires high-resolution, multifrequency data in order to separate dust emission, the thermal Sunyaev-Zeldovich (tSZ) effect, and pure blackbody fluctuations, as expected from both the lensed primary CMB and the kinematic Sunyaev-Zeldovich (kSZ) effect.

A.3.1 Frequency allocation

A broad investigation of possible allocations of detectors across frequency was performed. With three extragalactic components known to be present on small angular scales, in addition to possible low-frequency radio emission, and many Galactic components on large and moderate angular scales, at least three well-separated frequencies with high signal-to-noise are required. In addition, extragalactic dust emission—the cosmic infrared background (CIB)—is already known to exhibit frequency decorrelation due to the different redshift kernels of the CIB signal at different observational frequencies [574]. Thus, multiple channels may be needed in order to fully clean this component from the tSZ, kSZ, and lensing signals.

To assess the options for the distribution of detectors, we employed an end-to-end simulation-based optimization framework based on that used in Ref. [49]. We focused on temperature-based observables as metrics for optimization: the tSZ power spectrum, kSZ power spectrum, reconstructed CMB lensing power spectrum via the TT quadratic estimator (as a proxy for CMB “halo lensing,” which is TT -dominated), and the CMB TT power spectrum (this is already well-measured, but included for completeness). Due to the current lack of knowledge regarding small-scale polarized foregrounds, and the expected stronger need for multifrequency coverage for tSZ and kSZ observables, we did not consider the reconstructed CMB lensing power spectrum from polarization data in this optimization.

We used the CMB-S4 LAT noise calculator described earlier to forecast the S/N of these observables for a large number of experimental configurations. We varied the number of optics tubes of each type (LF=27/39 GHz, MF=93/145 GHz, UHF=225/280 GHz), considering here also the possibility of XHF tubes with channels at 281 and 350 GHz, assumed to be located at a high Chilean site with excellent atmospheric properties. We assumed two identical LAT copies. Several thousand configurations were considered. *Planck* data from 30 to 353 GHz were also assumed in all forecasts; these channels are useful on large angular scales where the CMB-S4 atmospheric noise is large.

We modeled the temperature sky at all frequencies from 27 to 353 GHz using the simulated sky maps described in section 2 of [49]. These maps include models for essentially all Galactic and extragalactic foregrounds (and signals). Simple Galactic-emission-thresholded sky masks that leave the cleanest amount of sky that is visible from Chile were employed to self-consistently include the effect of these sky cuts on the Galactic foreground levels. While several cut levels were explored, it was found that the sky fraction did not strongly affect the frequency allocation, with a general trend of higher S/N for larger sky area surveyed.

For configuration option, we used a harmonic-space internal linear combination (ILC) code to obtain post-component-separation noise power spectra for the blackbody CMB temperature and tSZ fields, using the modeled sky power spectra and the per-frequency noise power spectra computed using the CMB-S4 calculator (as well as *Planck* noise, assumed to be white). We considered the option of explicitly “deprojecting” some contaminants using a constrained ILC, which is a robust way to conservatively assess the frequency

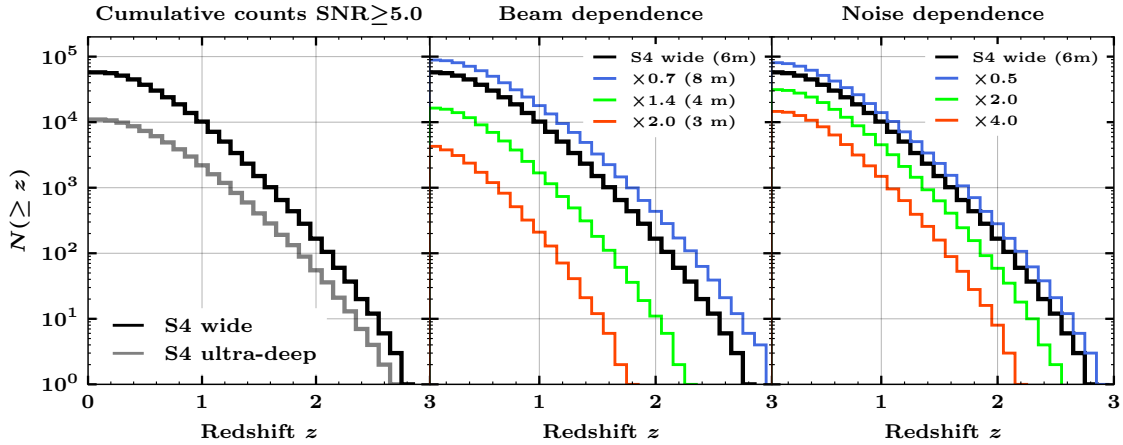


Figure 78. Panel 1 shows the expected number of $SNR \geq 5$ clusters as a function of redshift for the deep and wide (black) and ultra-deep (gray) fields with a 6-m aperture corresponding to 1.15 at 150 GHz. Panels 2, 3 show the dependence on telescope size and noise levels for the deep and wide survey. To get a sample of clusters at $z \gtrsim 2$, given our current understanding of high-redshift clusters, an aperture size ≥ 6 m is required. The sky coverage has been assumed as $f_{\text{sky}} = 0.4$ and $f_{\text{sky}} = 0.03$ for the deep and wide and ultra-deep fields respectively.

coverage that may be needed to sufficiently remove these foregrounds. We considered three deprojection options: no deprojection, deprojection of tSZ (for CMB reconstruction) or CMB (for tSZ reconstruction), and deprojection of a fiducial CIB spectrum (for CMB and tSZ reconstruction). The total number of sky fraction/configuration/deprojection options is 10260. We then flowed down to determine the experimental setup by optimizing the S/N of the observables described above, for each deprojection choice.

For each observable and each deprojection choice, the optimal tube configuration was slightly different, but general trends were clear. The optimization preferred a broad frequency coverage, always including at least one tube of each type. Typical optimal configurations included 1 or 2 LF tubes, 7–14 MF tubes, and 3–11 UHF tubes. Deprojecting tSZ or CIB foregrounds generally put a stronger demand on the need for UHF tubes. We also found that inclusion of XHF tubes at even higher frequencies could lead to 10–20% gains in S/N for some observables, particularly when deprojecting tSZ or CIB foregrounds. This option may be worthy of further study in the future.

The most important conclusion from this optimization is that the CMB-S4 LAT reference configuration (2 LF tubes, 12 MF tubes, 5 UHF tubes) used throughout this work is sufficiently near-optimal to serve as an excellent choice. In nearly all cases, it performed within 5–10% of the maximum S/N found in the optimization; the only exceptions to this were in measurements of the kSZ power spectrum with tSZ or CIB deprojection, where the reference configuration was within 15–20% of the maximal S/N found in the optimization. While future refinements of the detector allocation across frequency may be performed, particularly with updated CIB modeling to inform the high-frequency optimization, the flowdown presented here justifies the reference distribution that has been used throughout this document.




Figure 79. Limiting mass as a function of redshift for CMB-S4 galaxy cluster surveys. Also shown are existing catalogs of clusters selected by either SZ or X-ray, as well as the projected e-Rosita mass limit.

A.3.2 Angular resolution and sensitivity

In terms of angular resolution, higher resolution will always be better for more precisely constraining the astrophysics of galaxy formation; the minimum resolution is set by the typical angular scale of high-redshift galaxy clusters. Section 1.4.2 shows that CMB-S4 provides new views on the thermodynamics of galaxy clusters on the scale of the virial region, while other observations of galaxy feedback are typically on smaller scales. As resolution is improved it will be easier to make direct comparisons with optical/IR observations.

A science goal is to be able to detect the progenitors of massive clusters today at the epoch when they were forming the bulk of their stars ($z \approx 2-3$). Fig. 79 shows that CMB-S4 will be able to detect the typical massive low-redshift massive clusters at $z \approx 2$; clusters will be detected out to $z \approx 3$ with reference design sensitivity, assuming that the thermal pressure is well-described by the model used to simulate the signal [577]. For a typical progenitor cluster ($10^{14} M_{\odot}$ at $z = 3$), the typical cluster diameter ($2R_{200}$) is $2'$. This puts a lower limit on resolution of $2'$ at 90 GHz. Similarly, Fig. 78 shows that detecting clusters at $z = 3$ at CMB-S4 sensitivity requires a beam size no larger than $\text{FWHM} = 1.5'$ at 150 GHz [430].

For detecting the progenitors of massive clusters at $z \approx 2$ and higher, Fig. 80 shows that either a larger beam size or higher noise level leads to an inability to access these clusters.

A.4 Gamma-ray bursts, sensitivity, angular resolution, and cadence

Time-variable sources will be spatially unresolved in CMB-S4. Using difference imaging, source confusion will not be a problem, so the sensitivity to time-variable sources will be determined solely by the instrument noise and aperture size.




Figure 80. Effect on limiting mass as a function of redshift by varying noise levels (left) and beam size (right). Solid lines (red, green, blue) show typical progenitor masses as a function of redshift for three different low-redshift cluster groups.




Figure 81. Projections for 150-GHz transient source counts from Metzger et al. (2015) [468]. Shown are the total number of on-axis long gamma-ray bursts that are expected to be visible in the entire sky at any one time as a function of flux at 150 GHz. The width of the band schematically represents the uncertainty in this estimate. The inset shows some mm-wave follow-up observations of long gamma-ray bursts [494], showing typical variations on time scales of several days.

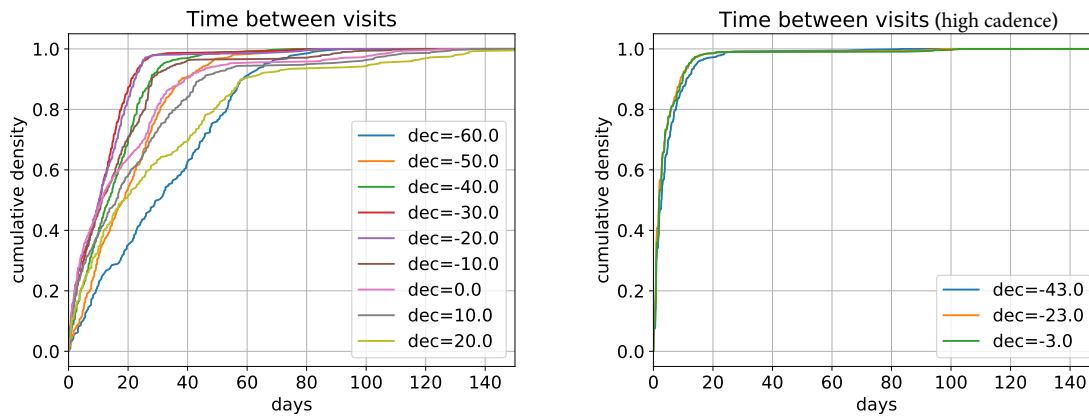


Figure 82. Cumulative fraction of time between visits for the deep and wide field, split by declination. Left panel shows the default scan strategy assumed for calculating noise curves for the reference design; right panel shows results for a scan strategy with a high cadence that results in comparable overall survey performance. The ultra-deep field is expected to revisit every location daily in the reference design.

Some gamma-ray bursts have been measured to be visible at CMB-S4 wavelengths, with measured fluxes typically ranging from 1–10 mJy around 90 GHz in the first week. As shown in Fig. 81, it is not uncommon for the light curves, once they start to decline, to decline by a factor of two within just a few days, motivating a fast cadence.

The detailed number counts are currently unknown, as no systematic surveys have been performed. Theoretical calculations have been performed [468] that are broadly consistent with the targeted follow-up. These theoretical predictions for GRBs are shown in Fig. 81. At fluxes of several mJy the expectation is that there should be a handful of long GRBs visible somewhere on the sky at any moment in time, with each GRB afterglow lasting roughly one week. These calculations are consistent with the dedicated mm-wave follow-up successfully detecting roughly 1/4 of the GRBs that were targeted [494].

The scaling with flux is uncertain at this point, but a reasonable expectation when seeing the bright end of a population, borne out by calculations [468], would be a power-law with $N(> S) \propto S^{-3/2}$, where a factor of two improvement in flux sensitivity gains a factor of 2.8 more sources. A minimum sensitivity of 5 mJy every two days will ensure that there will be a substantial catalog of gamma-ray bursts. The ultra-deep survey, covering only a few percent of the sky but at substantially higher sensitivity, will also be useful for GRBs.

Observing most of the CMB-S4 area every two days (i.e., the entire survey area while avoiding the Sun and Moon) requires careful survey design. The “nominal schedule” scan strategy discussed in the Light Relics section above has not been optimized for cadence, but still leads to good sensitivity to transients. Figure 82 shows that, with the exception of the lowest declinations, 1/3 of the survey area is covered weekly in the reference design. From preliminary exploration of alternative observing schedules we have found we can do significantly better on cadence without degradation to performance on light relics. Specifically, we have found an observing strategy that achieves the cadence shown in the right panel of Fig. 82, with uniform coverage every two days over the same area as that covered by the nominal schedule. In addition, the ultra-deep field (for delensing) will have daily coverage of that entire area (a few percent of the sky).

Some GRBs have shown evidence for a possible “reverse shock,” which will be luminous at mm-wavelengths in the first few hours [495]. Therefore, it is important to have a large area of the sky covered with a cadence

CMB-S4 Survey	Point-source rms noise (mJy)					
	30 GHz	40 GHz	95 GHz	145 GHz	220 GHz	270 GHz
Full depth (wide)	0.8	0.6	0.2	0.2	0.6	1.1
1 week (wide)	15	11	3	4	11	21
Full depth (ultra-deep)	0.1	0.1	0.04	0.07	0.4	0.5
1 week (ultra-deep)	2	3	0.8	1	7	9

Table A-20. *Rms noise levels (in mJy) for point sources at different frequencies for aspects of the CMB-S4 survey, either full depth after coadding the entire survey or the characteristic depth from a single week of data*

close to daily (at least every other day), which would provide a high probability of catching some events in the first 12 hours. The ultra-deep field will do this over a few percent of the sky; with a nearly daily GRB rate there should be several GRBs per year in this patch.

Better angular resolution and sensitivity and more frequency coverage are helpful. As these sources will be spatially unresolved, a key quantity is point-source sensitivity. Sensitivity to unresolved sources is improved by having better angular resolution, with an approximate scaling of signal-to-noise scaling as the inverse of the beam size. It is necessary to have at least two well-measured frequencies to be able to determine an approximate spectral index for follow-up at other wavelengths.

B

Options

B.1 Re-use of existing telescopes and site infrastructure

Instead of building two new 6-m telescopes in Chile, it may be possible to re-use CCAT-prime and/or the SO-LAT, and their associated site infrastructure. Re-use of existing telescopes could reduce the overall cost of CMB-S4. Another option would be to simply add the existing telescopes to CMB-S4 to increase sensitivity. Re-use would likely involve a fee, upgrades to monitoring and control systems, and some additional testing and alignment to meet CMB-S4 requirements. The details are very uncertain, so the reference design assumes all-new construction.

There are two critical decision points for the re-use option: (i) we could agree to re-use telescopes before awarding a design contract, in which case there would be no design or fabrication costs for new telescopes in Chile; or (ii) we could agree to re-use telescopes before awarding a fabrication contract, in which case there would be no fabrication costs for the telescopes, but we would still pay for the design.

B.2 Readout

B.2.1 fmux

Frequency-division multiplexing (fmux) is an alternative TES readout that has been implemented in several Stage-3 experiments. In an fmux readout, each TES is wired in series with an inductor and a capacitor to create a resonant filter, and many of these TES-filter segments are connected in parallel. A comb of AC carriers is applied to the network to AC voltage-bias the TESs, with each TES-filter segment selecting one of the carriers. The currents flowing through the TESs are summed together and amplified by a SQUID. Since each network of TES-filter segments has just one SQUID, the cold electronics are simple, but the SQUIDs must handle a large signal range. Alternative TES readout schemes all have one SQUID per TES, which relaxes the signal range requirements at the expense of more complicated electronics.

The key advantage of the fmux readout is that the amplitude of the AC voltage bias for each TES can be optimized individually to compensate for variations in TES saturation power across a detector wafer. Some variation in saturation power is inevitable because of fabrication tolerances and variations in optical loading. The ability to adjust each TES bias relaxes fabrication tolerances and improves the yield of working detectors.

A second important advantage of the fmux readout is that there are only two wires for each network of many TES-filter segments, so thermal loading on the camera's sub-kelvin stage is small. As an example from an existing fmux implementation, the thermal load for a 5-cm long niobium-titanium stripline connecting an array of TES-filters at 250 mK to a thermal intercept at 350 mK is only ≈ 4 nW. The filter components at 250 mK are all superconducting, so they do not dissipate power onto the 250-mK stage.

Several CMB experiments have deployed fmux readout, including SPT, POLARBEAR, and EBEX. SPT-POL and POLARBEAR used an fmux configuration with 8–12 bolometers multiplexed together, POLARBEAR2 will soon deploy a readout with a MUX factor of 40, and SPT-3G is currently observing with a camera that has a readout MUX factor of 68. SPT-3G provides a concrete demonstration of background-limited performance with an fmux readout at the scale of 16 k detectors.

The main issue for the fmux readout is noise from the readout electronics. To achieve a constant voltage bias, the TES resistance must be larger than the parasitic resistance in the bias filter, but higher TES resistance increases the NEP contribution from SQUID input current noise. The readout NEP in deployed, ground-based fmux systems is a few tens of $\text{aW}\sqrt{\text{Hz}}$, which is comparable with the photon noise. Lower readout noise can be achieved if parasitic resistances are carefully controlled. SRON has demonstrated low readout noise in the laboratory, and ANL and LBNL are working on fmux designs that have the SQUIDs at 250 mK, so there are fewer connections and smaller parasitics. The ANL and LBNL work is supported by LDRD funds.

To be a viable option for CMB-S4, the fmux readout must show integrated performance appropriate for CMB-S4 at the scale of a detector wafer in the laboratory, and preferably on the sky, before the start of the final CMB-S4 design.

B.2.2 μmux

Microwave-multiplexed SQUID readout (μmux) is an evolution of the DC voltage-biased TES readouts that are used in the small-aperture CMB experiments at the South Pole and some of the large-aperture experiments in Atacama. The μmux readout borrows heavily from the DC-SQUID TDM scheme, replacing the TES current sensor (a DC-SQUID) with an equally low-noise extremely high-MUX-factor current sensor that is frequency-division multiplexed. The new current sensor is an RF-SQUID flux-coupled to a GHz-frequency superconducting resonator that can be read out much like an MKID. The signal from each TES modulates the magnetic flux through its SQUID, which in turn modulates the resonant frequency of its resonator.

An important advantage of the μmux readout is that the design of the TES bolometer and current sensor can be optimized independently of each other, as in a TDM readout, enabling large safety factors/margins in both bolometer and readout parameters and thus noise. For instance, μmux can use low-resistance DC-biased TESs ($\approx 10 \text{ m}\Omega$), which significantly reduces the NEP contribution from the readout. Readout NEP levels of $5\text{--}7 \text{ aW}/\sqrt{\text{Hz}}$ (well below the photon noise for a typical ground-based CMB experiment) have been achieved in laboratory tests, using bolometers similar to those that would be deployed in CMB-S4. Additionally, μmux provides much higher multiplexing factors than either TDM or fmux, with MUX factors exceeding 400 having been achieved in laboratory demonstrations and MUX factors of $\gtrsim 2000$ are realistic. A high MUX factor may significantly reduce integration complexity and readout cost per channel, but must be balanced against the risk of a single-point failure.

A drawback of the μmux readout is that large numbers of TESs are connected in series to share the same bias, so it is not possible to optimize the operating point of each bolometer individually. TDM readouts have the same issue, and Stage-3 experiments have developed strategies for optimizing the performance of groups of bolometers that share a common bias. As bolometer uniformity achieved at fabrication has improved, the impact of this drawback has declined.

The μmux readout uses a phase modulation scheme that greatly reduces sensitivity to low-frequency amplitude-modulation noise (the same technique makes FM/XM radio channels sound less noisy than AM

radio channels). In addition, the readout electronics under development at SLAC use a fully digital approach that is largely insensitive to low-frequency noise induced by temperature fluctuations and phase drifts. The readout uses a “tone-tracking” feedback loop in which the microwave carriers that excite the GHz-frequency resonators are continuously maintained at each channel’s resonant frequency. Feedback increases the dynamic range and linearity while reducing sensitivity to a variety of noise sources including gain variations and crosstalk.

The μ mux development effort is supported by LDRD funds. A 64 MUX factor μ mux readout was deployed on the Green Bank Telescope (GBT) to read out the 215-pixel MUSTANG2 array, prototypes of microwave SQUIDs suitable for CMB observations have been produced, and a full prototype readout has been demonstrated with a MUX factor of 528. SLAC and FNAL are working on warm electronics capable of ≥ 2000 MUX factor, nearly an order of magnitude higher than fmux and TDM. In conjunction with this warm readout development, NIST has developed cold μ mux multiplexers achieving a resonator density which will enable MUX factors of 2000. Both BICEP Array and Simons Observatory have baselined μ mux in their instruments.

To be a viable option for CMB-S4, the μ mux readout must show integrated performance appropriate for CMB-S4 at the scale of a detector wafer in the laboratory, and preferably on the sky, before the start of the final CMB-S4 design.

B.3 Kinetic-inductance detectors

KIDs are superconducting thin-film micro resonators that detect radiation through a change in inductance when the superconductor absorbs photons. Many resonators, each at a slightly different frequency, can be coupled to a single readout line to achieve a high multiplexing factor. Multiplexing on the detector wafer is the key advantage of KIDs, because it allows many thousands of detectors on a wafer, with just a few connections. In contrast, TESes read with μ mux require an active squid per TES as well as a pair of wire bonds per TES, requiring high wiring density at the wafer edge. This challenge is particularly acute for bands with small pixels, e.g., CMB-S4 220/270 GHz. KIDs have other advantages: a much larger dynamic range than TESs; no SQUID amplifiers in the readout; simple fabrication; and compatibility with warm μ mux hardware. TKIDs are an alternative that uses the inductor as the thermal sensor of a released bolometer island. This introduces several additional design parameters that have been used to fine tune the detector noise to $20 \text{ aW}/\sqrt{\text{Hz}}$ stable to 100 mHz, measured under loading representative of atmospheric loading at the South Pole. Such performance would be background limited in observing bands at 90 GHz and higher. Their absorptive nature makes them drop-in compatible with antenna-coupling schemes used for multicolor pixels or observing bands below the superconducting band-gap.

Several millimeter-wavelength KID instruments have been deployed on telescopes (e.g., MUSIC, NIKA/NIKA2, MAKO, DESHIMA, A-MKID), but not for CMB observations. To be a viable option for CMB-S4 requires an on-sky demonstration of background limited performance, in a CMB polarization experiment, before CMB-S4 final design starts. Several groups are working towards such a demonstration (e.g., Columbia++ antenna-coupled KIDs, JPL TKIDs, Chicago CMB KIDs). The KID development effort is supported by NSF and NASA funds.

B.4 Coupling to detectors

B.4.1 Lenslets

Several Stage-3 experiments have deployed lenslet-coupled detectors, which have planar antennas on the detector wafer and an array of lenslets to couple the antennas to the telescope. The principal advantage of this approach is wide bandwidth, which allows more detectors to be squeezed into the focal plane, e.g., each pixel in SPT-3G has 90-, 150-, and 220-GHz detectors. Lenslet-coupled detector arrays are compact, there are no issues with differential thermal contraction, because the detector and lenslet wafers are both made of silicon, and it may be possible to fabricate the lenslets commercially; the POLARBEAR group is pursuing commercial fabrication of monolithic silicon lenslet arrays. The main issue with lenslet-coupled detectors is fairly large frequency-dependent polarization errors in wide-band planar antennas. If the technology is to be a viable option for CMB-S4, we must demonstrate correction of the polarization errors leading to residual systematic errors well below the map noise in Stage-3 experiments. Both SPT-3G and POLARBEAR are working towards such a demonstration.

B.4.2 Planar antennas

Antenna-coupled transition-edge sensor bolometers have been developed for a wide range of CMB polarimetry experiments, including BICEP2, BICEP3, Keck Array, and SPIDER. These photo-lithographed planar antenna arrays synthesize symmetric co-aligned beams, with controlled side-lobes and low cross-polarized response (typically 0.5% on boresight, consistent with cross-talk in the multiplexed readout system). End-to-end optical efficiencies in these cameras are routinely 35% or higher, within well-controlled 25% wide spectral bands at 95, 150, 220 and 270 GHz. Thanks to the frequency scalability of this design, more than 90 science wafers have been deployed on the sky to date, achieving a sensitivity of 50 nK-deg at 95 and 150 GHz. Planar antennas show tight control of systematic errors, demonstrated to be sub-dominant to statistical noise at these levels in mapping degree scales with small aperture telescopes[553, 528]. With this demonstrated performance and high fabrication throughput, planar antennas represent a low-risk implementation for CMB-S4.

Planar antenna technology offers a unique detector packing advantage compared to, e.g., single-band hex-packed spline feedhorns, due to the intrinsically steeper profile of a flat-top vs Gaussian illumination. Holding the edge taper on the aperture stop to be the same, square planar antennas pack more densely than spline feedhorns (in the large pixel / low edge taper regime), including gaps for detectors and wiring. Switching to a hexagonal planar antenna may confer some additional efficiencies for filling out focal planes. Planar antennas are currently being developed and tested at 30 and 40 GHz for upcoming observations. Finally, a dual-band planar antenna that offers further improvements in packing density has been recently developed and will undergo extensive testing.

B.5 Commercial detector fabrication

Detectors for Stage-3 CMB experiments have been fabricated by national laboratories and universities, but CMB-S4 will need far more detectors, so fabrication throughput will be a challenge. Moving at least some of the detector fabrication to industry would increase the fabrication rate without an expensive in-house

investment in facilities and staff. The overall cost could be lower because an industrial partner may be able to draw staff and fabrication tools from a larger pool that supports many projects. Shift work is generally also easier to implement in an industrial setting. The main issue with moving detector fabrication to industry is controlling the superconducting properties of the transition edge sensor films, which typically requires dedicated tooling and stringent quality control. Most of the other fabrication steps are standard industry processes.

If commercial detector fabrication is to be a viable option, performance appropriate for CMB-S4 must be demonstrated in the laboratory, and preferably in the field, along with significant fabrication throughput, before CMB-S4 final design starts. LBNL and UCB are working with STARCryo and HYPRES to commercialize the POLARBEAR detector technology, and many process steps have already been demonstrated. The work is supported by SBIR and LDRD awards, and by a DOE Early Career award.

B.6 Other developments

The options described above are supported by LDRD and other funds, with CMB-S4 providing at least some coordination, but there are several other developments that CMB-S4 is watching, and may choose to take advantage of the following.

1. Monolithic versus segmented mirrors for the large telescopes. Eliminating the gaps between mirror panels would reduce ground pickup, which would allow the large telescopes to recover larger angular scales. Ground pickup is a serious concern for all CMB experiments, so a demonstration of a large monolithic mirror would almost certainly result in CMB-S4 adopting the technology.
2. Full boresight rotation for the large telescopes. Small CMB telescopes use boresight rotation to measure and correct polarization errors, but the technique has not yet been attempted on a large telescope, primarily because it requires an additional axis. The Simons Observatory and CCAT-prime large telescope designs offer partial boresight rotation through receiver rotation, which will enable improved correction of polarization errors compared to existing large telescopes.

C

List of Abbreviations

Λ CDM	Λ cold dark matter
μ mux	Microwave multiplexing
AAAC	Astronomy and Astrophysics Advisory Committee
ABS	Atacama B-mode Search
AC	Alternating current
ACT	Atacama Cosmology Telescope
ADC	Analog-to-digital converter
ADM	Arnowitt-Deser-Misner
ADMX	Axion Dark Matter Experiment
AGN	Active galactic nucleus
ALCF	Argonne Leadership Class Facility
ALMA	Atacama Large Millimeter/submillimeter Array
ALP	Axion-like particle
AME	Anomalous microwave emission
AR	Anti-reflection
ASAS-SN	All-Sky Automated Survey for Supernovae
ASKAP	Australian Square Kilometre Array Pathfinder
ATLAS	A Toroidal LHC ApparatuS detector at the LHC
AU	Astronomical unit
BAO	Baryon acoustic oscillations
BBN	Big-bang nucleosynthesis
BDI	Baryon density isocurvature
BOE	Basis of estimate
BOSS	Baryon Oscillation Spectroscopic Survey
BSM	Beyond the Standard Model
CAM	Cost account manager
CASPEr	Cosmic Axion Spin Precession Experiment
CCAT	Cerro Chajnantor Atacama Telescope
CD	Critical decision or Crossed Dragone
CDM	Cold dark matter

CDI	CDM density isocurvature
CDT	Concept Definition Task force
CFRP	Carbon-fiber-reinforced polymer
CGM	Circumgalactic medium
CHIME	Canadian Hydrogen Intensity Mapping Experiment
CIB	Cosmic infrared background
CIP	Compensated isocurvature perturbation
CL	Critical line or confidence limit
CLASS	Cosmology Large Angular Scale Surveyor
CMB	Cosmic microwave background
CMS	Compact Muon Solenoid detector at the Large Hadron Collider
CPU	Central processing unit
CRS	Cold readout stack
CTE	Coefficient of thermal expansion
D&R	Detectors and readout
DAC	Digital-to-analog converter
DAQ	Data acquisition
DE	Dark energy
DESI	Dark Energy Spectroscopic Instrument
DM	Dark matter
DOE	Department of Energy
DR	Dark radiation or Dilution refrigerator
DSFG	Dusty star-forming galaxy
DSL	Dark Sector Laboratory
DSR	Decadal Survey Report
DUNE	Deep Underground Neutrino Experiment
EH&S	Environment, health, and safety
EMI	Electromagnetic interference
EMU	Evolutionary Map of the Universe
ESA	European Space Agency
FACA	Federal Advisory Committee Act
FDM	Frequency-division multiplexing
FET	Field-effect transistor
FLRW	Friedmann-Lemaître-Robertson-Walker
fmux	Frequency-division multiplexing

FNAL	Fermilab National Accelerator Laboratory
FoM	Figure of merit
FOV	Field of view
FP	Focal plane
FPGA	Field-programmable gate array
FRB	Fast radio burst
FTE	Full-time equivalent
FTS	Fourier-transform spectrometer
FWHM	Full width at half maximum
GBT	Green Bank Telescope
GC	Galaxy cluster
GPS	Global positioning system
GRB	Gamma-ray burst
GSR	Generalized slow-roll
GUT	Grand unified theory
HDPE	High-density polyethylene
HEP	High-energy physics
HPC	High-performance computing
HTC	High-throughput computing
HWFE	Half wavefront error
HWP	Half-wave plate
I&C	Integration and commissioning
I&T	Integration and testing
ICM	Intracluster medium
ILC	Internal linear combination
IPO	Interim project office
IPSC	Integrated Project Steering Committee
IR	Infrared
IRIG	Inter-range instrumentation group
ISM	Interstellar medium
JCMT	James Clerk Maxwell Telescope
JPL	Jet Propulsion Laboratory
JWST	James Webb Space Telescope
KATRIN	KARlsruhe TRItium Neutrino (experiment)
KID	Kinetic-inductance detector

kSZ	Kinematic Sunyaev Zeldovich
L1	Level 1
LAMBDA	Legacy Archive for Microwave Background Data
LAT	Large-aperture telescope
LATR	Large-aperture telescope receiver
LBNL	Lawrence Berkeley National Laboratory
LDRD	Laboratory Directed Research and Development
LHC	Large Hadron Collider
LIGO	Laser Interferometer Gravitational-Wave Observatory
LISA	Laser Interferometer Space Antenna
LR	Light relics
LRG	Luminous red galaxy
LSS	Large-scale structure
LSST	Large Synoptic Survey Telescope
M&O	Management and operations
M&S	Materials and services
MAPO	Martin Pomerantz Observatory
MCMC	Markov-chain Monte-Carlo
MHD	Magnetohydrodynamic
MIE	Major Item of Equipment
MoU	Memorandum of understanding
MREFC	Major Research Equipment and Facilities Construction
MKID	Microwave kinetic-inductance detector
MUX	Multiplexing
NASA	National Aeronautics and Space Administration
NDI	Neutrino density isocurvature
NEP	Noise-equivalent power
NERSC	National Energy Research Scientific Computing Center
NET	Noise-equivalent temperature
NIST	National Institute of Standards and Technology
NLDBD	Neutrino-less double-beta decay
NSF	National Science Foundation
NLDBD	Neutrinoless double-beta decay
NLS1	Narrow-line Seyfert 1
NTP	Network time protocol

OCS	Observatory control system
OFHC	Oxygen-free high-conductivity
OMT	Orthomode transducer
OPC	Other project costs
P5	Particle Physics Project Prioritization Panel
PCI-E	Peripheral Component Interconnect Express
PD	Project director
PDR	Preliminary design review
PGW	Primordial gravitational waves
PID	Proportional–integral–derivative
PMNS	Pontecorvo-Maki-Nakagawa-Sakata
pPDG	pre-Project Development Group
pPEP	preliminary Project Execution Plan
PPS	Pulse per second
PTFE	Polytetrafluoroethylene
PTP	Precision Time Protocol
PySM	Python Sky Model
QA	Quality assurance
QCD	Quantum chromodynamics
R&D	Research and development
RF	Radio frequency
RMS	Root mean square
RSD	Redshift-space distortion
RT-MLI	Radio-transparent multi-layer insulation
SAT	Small-aperture telescope
SBIR	Small Business Innovation Research
SCUBA	Submillimeter Common-User Bolometer Array
SDSS	Sloan Digital Sky Survey
SED	Spectral energy distribution
SKA	Square Kilometre Array
SLAC	Stanford Linear Accelerator Center
SM	Standard Model (of particle physics)
SMuRF	SLAC Microresonator Radio Frequency
SO	Simons Observatory
SOFIA	Stratospheric Observatory for Infrared Astronomy,

SoW	Statement of work
SPT	South Pole Telescope
SQUID	Superconducting quantum-interference device
SRON	Netherlands Institute for Space Research
SSA	Series SQUID array
SWF	Salaries, wages, and fringe benefits
TDM	Time-division multiplexing
TEC	Total estimated cost
TES	Transition-edge sensor
TPC	Total project cost
TKID	Thermal kinetic inductance detector
TLS	Two-level system
tSZ	Thermal Sunyaev Zeldovich
UFM	Universal focal-plane module
UHMWPE	Ultra-high-molecular-weight polyethylene
ULA	Ultralight axion
UV	Ultraviolet
VLA	Very Large Array
VLA	VLA Sky Survey
WBS	Work breakdown structure
WFIRST	Wide Field Infrared Survey Telescope
WIMP	Weakly interacting massive particle
WISE	Wide-field Infrared Survey Explorer
WMAP	Wilkinson Microwave Anisotropy Probe
XSEDE	Extreme Science and Engineering Discovery Environment

References

- [1] K. N. Abazajian *et al.*, “Inflation Physics from the Cosmic Microwave Background and Large Scale Structure,” *Astropart. Phys.* **63** (2015) 55–65, [arXiv:1309.5381](#) [[astro-ph.CO](#)].
- [2] K. N. Abazajian *et al.*, “Neutrino Physics from the Cosmic Microwave Background and Large Scale Structure,” *Astropart. Phys.* **63** (2015) 66–80, [arXiv:1309.5383](#) [[astro-ph.CO](#)].
- [3] **CMB-S4** Collaboration, K. N. Abazajian *et al.*, “CMB-S4 Science Book, First Edition,” [arXiv:1610.02743](#) [[astro-ph.CO](#)].
- [4] M. H. Abitbol, Z. Ahmed, D. Barron, R. Basu Thakur, A. N. Bender, B. A. Benson, C. A. Bischoff, S. A. Bryan, J. E. Carlstrom, and C. L. Chang, “CMB-S4 Technology Book, First Edition,” *arXiv e-prints* (Jun, 2017) [arXiv:1706.02464](#), [arXiv:1706.02464](#) [[astro-ph.IM](#)].
- [5] **BICEP2, Keck Array** Collaboration, P. A. R. Ade *et al.*, “Improved Constraints on Cosmology and Foregrounds from BICEP2 and Keck Array Cosmic Microwave Background Data with Inclusion of 95 GHz Band,” *Phys. Rev. Lett.* **116** (2016) 031302, [arXiv:1510.09217](#) [[astro-ph.CO](#)].
- [6] R. Kallosh and A. Linde, “Universality Class in Conformal Inflation,” *JCAP* **1307** (2013) 002, [arXiv:1306.5220](#) [[hep-th](#)].
- [7] A. A. Starobinsky, “A New Type of Isotropic Cosmological Models Without Singularity,” *Phys.Lett.* **B91** (1980) 99–102.
- [8] F. L. Bezrukov and M. Shaposhnikov, “The Standard Model Higgs boson as the inflaton,” *Phys. Lett.* **B659** (2008) 703–706, [arXiv:0710.3755](#) [[hep-th](#)].
- [9] M. Cicoli, C. P. Burgess, and F. Quevedo, “Fibre Inflation: Observable Gravity Waves from IIB String Compactifications,” *JCAP* **0903** (2009) 013, [arXiv:0808.0691](#) [[hep-th](#)].
- [10] K. N. Abazajian *et al.*, “Light Sterile Neutrinos: A White Paper,” [arXiv:1204.5379](#) [[hep-ph](#)].
- [11] R. Essig *et al.*, “Working Group Report: New Light Weakly Coupled Particles,” in *Community Summer Study 2013: Snowmass on the Mississippi (CSS2013)*, Minneapolis, MN, USA, July 29-August 6, 2013. 2013. [arXiv:1311.0029](#) [[hep-ph](#)].
- [12] J. Alexander *et al.*, “Dark Sectors 2016 Workshop: Community Report,” [arXiv:1608.08632](#) [[hep-ph](#)].
- [13] N. Whitehorn, T. Natoli, P. A. R. Ade, J. E. Austermann, J. A. Beall, A. N. Bender, B. A. Benson, L. E. Bleem, J. E. Carlstrom, C. L. Chang, H. C. Chiang, H.-M. Cho, R. Citron, T. M. Crawford, A. T. Crites, T. de Haan, M. A. Dobbs, W. Everett, J. Gallicchio, E. M. George, A. Gilbert, N. W. Halverson, N. Harrington, J. W. Henning, G. C. Hilton, G. P. Holder, W. L. Holzapfel, S. Hoover, Z. Hou, J. D. Hrubes, N. Huang, J. Hubmayr, K. D. Irwin, R. Keisler, L. Knox, A. T. Lee, E. M. Leitch, D. Li, J. J. McMahon, S. S. Meyer, L. Mocuano, J. P. Nibarger, V. Novosad, S. Padin, C. Pryke, C. L. Reichardt, J. E. Ruhl, B. R. Saliwanchik, J. T. Sayre, K. K. Schaffer, G. Smecher, A. A. Stark, K. T. Story, C. Tucker, K. Vanderlinde, J. D. Vieira, G. Wang, and V. Yefremenko, “Millimeter Transient Point Sources in the SPTpol 100 Square Degree Survey,” *Ap. J.* **830** (Oct., 2016) 143, [arXiv:1604.03507](#) [[astro-ph.HE](#)].

- [14] K. D. Alexander, E. Berger, W. Fong, P. K. G. Williams, C. Guidorzi, R. Margutti, B. D. Metzger, J. Annis, P. K. Blanchard, D. Brout, D. A. Brown, H.-Y. Chen, R. Chornock, P. S. Cowperthwaite, M. Drout, T. Eftekhari, J. Frieman, D. E. Holz, M. Nicholl, A. Rest, M. Sako, M. Soares-Santos, and V. A. Villar, “The Electromagnetic Counterpart of the Binary Neutron Star Merger LIGO/Virgo GW170817. VI. Radio Constraints on a Relativistic Jet and Predictions for Late-time Emission from the Kilonova Ejecta,” *Ap. J. Lett.* **848** (Oct., 2017) L21, [arXiv:1710.05457](#) [astro-ph.HE].
- [15] E. Berger, “Short-Duration Gamma-Ray Bursts,” *Ann. Rev. Astron. Astroph.* **52** (Aug., 2014) 43–105, [arXiv:1311.2603](#) [astro-ph.HE].
- [16] **Planck** Collaboration, Planck Collaboration XXIV, “Planck 2013 Results. XXIV. Constraints on primordial non-Gaussianity,” *Astron. Astrophys.* **571** (2014) A24, [arXiv:1303.5084](#) [astro-ph.CO].
- [17] Planck Collaboration XVII, “Planck 2015 results. XVII. Constraints on primordial non-Gaussianity,” [arXiv:1502.01592](#) [astro-ph.CO].
- [18] R. L. Arnowitt, S. Deser, and C. W. Misner, “The Dynamics of general relativity,” *Gen. Rel. Grav.* **40** (2008) 1997–2027, [arXiv:gr-qc/0405109](#) [gr-qc].
- [19] R. Namba, M. Peloso, M. Shiraishi, L. Sorbo, and C. Unal, “Scale-dependent gravitational waves from a rolling axion,” *JCAP* **1601** no. 01, (2016) 041, [arXiv:1509.07521](#) [astro-ph.CO].
- [20] A. Linde, “Gravitational waves and large field inflation,” *JCAP* **1702** no. 02, (2017) 006, [arXiv:1612.00020](#) [astro-ph.CO].
- [21] D. Babich, P. Creminelli, and M. Zaldarriaga, “The Shape of Non-Gaussianities,” *JCAP* **0408** (2004) 009, [arXiv:astro-ph/0405356](#) [astro-ph].
- [22] E. Komatsu *et al.*, “Non-Gaussianity as a Probe of the Physics of the Primordial Universe and the Astrophysics of the Low Redshift Universe,” [arXiv:0902.4759](#) [astro-ph.CO].
- [23] P. D. Meerburg, J. Meyers, A. van Engelen, and Y. Ali-Haïmoud, “CMB B -mode non-Gaussianity,” *Phys. Rev.* **D93** no. 12, (2016) 123511, [arXiv:1603.02243](#) [astro-ph.CO].
- [24] J. M. Maldacena and G. L. Pimentel, “On graviton non-Gaussianities during inflation,” *JHEP* **09** (2011) 045, [arXiv:1104.2846](#) [hep-th].
- [25] L. Bordin, P. Creminelli, M. Mirbabayi, and J. Noreña, “Tensor Squeezed Limits and the Higuchi Bound,” *JCAP* **1609** no. 09, (2016) 041, [arXiv:1605.08424](#) [astro-ph.CO].
- [26] A. Agrawal, T. Fujita, and E. Komatsu, “Large tensor non-Gaussianity from axion-gauge field dynamics,” *PRD* **97** no. 10, (2018) 103526, [arXiv:1707.03023](#) [astro-ph.CO].
- [27] G. Domènech, T. Hiramatsu, C. Lin, M. Sasaki, M. Shiraishi, and Y. Wang, “CMB Scale Dependent Non-Gaussianity from Massive Gravity during Inflation,” *JCAP* **1705** no. 05, (2017) 034, [arXiv:1701.05554](#) [astro-ph.CO].
- [28] H. Lee, D. Baumann, and G. L. Pimentel, “High-Scale Inflation and the Tensor Tilt,” [arXiv:1607.03735](#) [hep-th].
- [29] H. Lee, D. Baumann, and G. L. Pimentel, “Non-Gaussianity as a Particle Detector,” *JHEP* **12** (2016) 040, [arXiv:1607.03735](#) [hep-th].
- [30] D. Baumann, G. Goon, H. Lee, and G. L. Pimentel, “Partially Massless Fields During Inflation,” *JHEP* **04** (2018) 140, [arXiv:1712.06624](#) [hep-th].

- [31] Planck Collaboration, Planck Collaboration X, “Planck 2018 results. X. Constraints on inflation,” arXiv:1807.06211 [astro-ph.CO].
- [32] C. L. Bennett *et al.*, “First Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Preliminary Maps and Basic Results,” *Astrophys. J. Suppl.* **148** (2003) 1, arXiv:astro-ph/0302207.
- [33] D. N. Spergel *et al.*, “First Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Determination of Cosmological Parameters,” *Astrophys. J. Suppl.* **148** (2003) 175, arXiv:astro-ph/0302209.
- [34] Planck Collaboration XX, “Planck 2015 results. XX. Constraints on inflation,” arXiv:1502.02114 [astro-ph.CO].
- [35] A. Slosar *et al.*, “Scratches from the Past: Inflationary Archaeology through Features in the Power Spectrum of Primordial Fluctuations,” arXiv:1903.09883 [astro-ph.CO].
- [36] S. L. Bridle, A. M. Lewis, J. Weller, and G. Efstathiou, “Reconstructing the primordial power spectrum,” *Mon. Not. Roy. Astron. Soc.* **342** (July, 2003) L72–L78, arXiv:astro-ph/0302306 [astro-ph].
- [37] Z.-K. Guo, D. J. Schwarz, and Y.-Z. Zhang, “Reconstruction of the primordial power spectrum from CMB data,” *Journal of Cosmology and Astro-Particle Physics* **2011** (Aug., 2011) 031, arXiv:1105.5916 [astro-ph.CO].
- [38] R. Hlozek, J. Dunkley, G. Addison, J. W. Appel, J. R. Bond, C. Sofia Carvalho, S. Das, M. J. Devlin, R. Dünner, T. Essinger-Hileman, J. W. Fowler, P. Gallardo, A. Hajian, M. Halpern, M. Hasselfield, M. Hilton, A. D. Hincks, J. P. Hughes, K. D. Irwin, J. Klein, A. Kosowsky, T. A. Marriage, D. Marsden, F. Menanteau, K. Moodley, M. D. Niemack, M. R. Nolta, L. A. Page, L. Parker, B. Partridge, F. Rojas, N. Sehgal, B. Sherwin, J. Sievers, D. N. Spergel, S. T. Staggs, D. S. Swetz, E. R. Switzer, R. Thornton, and E. Wollack, “The Atacama Cosmology Telescope: A Measurement of the Primordial Power Spectrum,” *Ap. J.* **749** (Apr., 2012) 90, arXiv:1105.4887 [astro-ph.CO].
- [39] M. Aich, D. K. Hazra, L. Sriramkumar, and T. Souradeep, “Oscillations in the inflaton potential: Complete numerical treatment and comparison with the recent and forthcoming CMB datasets,” *Phys. Rev. D* **87** (Apr., 2013) 083526, arXiv:1106.2798 [astro-ph.CO].
- [40] P. Hunt and S. Sarkar, “Reconstruction of the primordial power spectrum of curvature perturbations using multiple data sets,” *Journal of Cosmology and Astro-Particle Physics* **2014** (Jan., 2014) 025, arXiv:1308.2317 [astro-ph.CO].
- [41] R. de Putter, E. V. Linder, and A. Mishra, “Inflationary freedom and cosmological neutrino constraints,” *Phys. Rev. D* **89** (May, 2014) 103502, arXiv:1401.7022 [astro-ph.CO].
- [42] V. Miranda, W. Hu, C. He, and H. Motohashi, “Nonlinear excitations in inflationary power spectra,” *Phys. Rev. D* **93** (Jan., 2016) 023504.
- [43] D. K. Hazra, A. Shafieloo, G. F. Smoot, and A. A. Starobinsky, “Primordial features and Planck polarization,” *Journal of Cosmology and Astro-Particle Physics* **2016** (Sept., 2016) 009.
- [44] C. Dvorkin and W. Hu, “Generalized Slow Roll for Large Power Spectrum Features,” *Phys. Rev. D* **81** (2010) 023518, arXiv:0910.2237 [astro-ph.CO].
- [45] C. Dvorkin and W. Hu, “CMB Constraints on Principal Components of the Inflaton Potential,” *Phys. Rev. D* **82** (2010) 043513, arXiv:1007.0215 [astro-ph.CO].

- [46] C. Dvorkin and W. Hu, “Complete WMAP Constraints on Bandlimited Inflationary Features,” *Phys. Rev. D* **84** (2011) 063515, [arXiv:1106.4016 \[astro-ph.CO\]](#).
- [47] G. Obied, C. Dvorkin, C. Heinrich, W. Hu, and V. Miranda, “Inflationary Features and Shifts in Cosmological Parameters from Planck 2015 Data,” *Phys. Rev. D* **96** no. 8, (2017) 083526, [arXiv:1706.09412 \[astro-ph.CO\]](#).
- [48] G. Obied, C. Dvorkin, C. Heinrich, W. Hu, and V. Miranda, “Inflationary versus reionization features from Planck 2015 data,” *Phys. Rev. D* **98** no. 4, (2018) 043518, [arXiv:1803.01858 \[astro-ph.CO\]](#).
- [49] The Simons Observatory Collaboration, P. Ade, J. Aguirre, Z. Ahmed, S. Aiola, A. Ali, D. Alonso, M. A. Alvarez, K. Arnold, P. Ashton, and et al., “The Simons Observatory: Science goals and forecasts,” *ArXiv e-prints* (Aug., 2018) , [arXiv:1808.07445](#).
- [50] J. R. Bond and G. Efstathiou, “Cosmic background radiation anisotropies in universes dominated by nonbaryonic dark matter,” *Astrophys. J.* **285** (1984) L45–L48.
- [51] H. Kodama and M. Sasaki, “Evolution of Isocurvature Perturbations. 1. Photon - Baryon Universe,” *Int. J. Mod. Phys. A* **1** (1986) 265.
- [52] H. Kodama and M. Sasaki, “Evolution of Isocurvature Perturbations. 2. Radiation Dust Universe,” *Int. J. Mod. Phys. A* **2** (1987) 491.
- [53] W. Hu and N. Sugiyama, “Toward understanding CMB anisotropies and their implications,” *Phys. Rev. D* **51** (1995) 2599–2630, [astro-ph/9411008](#).
- [54] K. Moodley, M. Bucher, J. Dunkley, P. G. Ferreira, and C. Skordis, “Constraints on isocurvature models from the WMAP first-year data,” *Phys. Rev. D* **70** (2004) 103520, [arXiv:astro-ph/0407304 \[astro-ph\]](#).
- [55] R. Bean, J. Dunkley, and E. Pierpaoli, “Constraining Isocurvature Initial Conditions with WMAP 3-year data,” *Phys. Rev. D* **74** (2006) 063503, [astro-ph/0606685](#).
- [56] K. Enqvist, H. Kurki-Suonio, and J. Valiviita, “Limits on isocurvature fluctuations from boomerang and MAXIMA,” *Phys. Rev. D* **62** (2000) 103003, [arXiv:astro-ph/0006429 \[astro-ph\]](#).
- [57] C. J. MacTavish *et al.*, “Cosmological parameters from the 2003 flight of BOOMERANG,” *Astrophys. J.* **647** (2006) 799–812, [arXiv:astro-ph/0507503 \[astro-ph\]](#).
- [58] **WMAP** Collaboration, J. Dunkley *et al.*, “Five-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Likelihoods and Parameters from the WMAP data,” *Astrophys. J. Suppl.* **180** (2009) 306–329, [arXiv:0803.0586 \[astro-ph\]](#).
- [59] **Planck** Collaboration, Planck Collaboration XXII, “Planck 2013 results. XXII. Constraints on inflation,” *Astron. Astrophys.* **571** (2014) A22, [arXiv:1303.5082 \[astro-ph.CO\]](#).
- [60] S. Mollerach, “Isocurvature Baryon Perturbations and Inflation,” *Phys.Rev.* **D42** (1990) 313–325.
- [61] V. F. Mukhanov, H. Feldman, and R. H. Brandenberger, “Theory of cosmological perturbations. Part 1. Classical perturbations. Part 2. Quantum theory of perturbations. Part 3. Extensions,” *Phys.Rept.* **215** (1992) 203–333.
- [62] T. Moroi and T. Takahashi, “Effects of cosmological moduli fields on cosmic microwave background,” *Phys.Lett.* **B522** (2001) 215–221, [hep-ph/0110096](#).

- [63] D. H. Lyth and D. Wands, “Generating the curvature perturbation without an inflaton,” *Phys. Lett.* **B524** (2002) 5–14, [arXiv:hep-ph/0110002](#) [hep-ph].
- [64] D. H. Lyth, C. Ungarelli, and D. Wands, “The Primordial density perturbation in the curvaton scenario,” *Phys.Rev.* **D67** (2003) 023503, [astro-ph/0208055](#).
- [65] M. Postma, “The Curvaton scenario in supersymmetric theories,” *Phys. Rev.* **D67** (2003) 063518, [hep-ph/0212005](#).
- [66] S. Kasuya, M. Kawasaki, and F. Takahashi, “MSSM curvaton in the gauge mediated SUSY breaking,” *Phys. Lett.* **B578** (2004) 259–268, [hep-ph/0305134](#).
- [67] M. Ikegami and T. Moroi, “Curvaton scenario with Affleck-Dine baryogenesis,” *Phys. Rev.* **D70** (2004) 083515, [hep-ph/0404253](#).
- [68] A. Mazumdar and A. Perez-Lorezana, “Sneutrino condensate as a candidate for the hot big bang cosmology,” *Phys. Rev.* **D70** (2004) 083526, [hep-ph/0406154](#).
- [69] R. Allahverdi, K. Enqvist, A. Jokinen, and A. Mazumdar, “Identifying the curvaton within MSSM,” *JCAP* **0610** (2006) 007, [hep-ph/0603255](#).
- [70] E. Papantonopoulos and V. Zamarias, “Curvaton Dynamics in Brane-worlds,” *JCAP* **0611** (2006) 005, [gr-qc/0608026](#).
- [71] A. Mazumdar and J. Rocher, “Particle physics models of inflation and curvaton scenarios,” *Phys. Rept.* **497** (2011) 85–215, [arXiv:1001.0993](#).
- [72] A. Mazumdar and S. Nadathur, “The curvaton scenario in the MSSM and predictions for non-Gaussianity,” *Phys. Rev. Lett.* **108** (2012) 111302, [arXiv:1107.4078](#).
- [73] C. Gordon and A. Lewis, “Observational constraints on the curvaton model of inflation,” *Phys.Rev.* **D67** (2003) 123513, [astro-ph/0212248](#).
- [74] C. Gordon and K. A. Malik, “WMAP, neutrino degeneracy and non-Gaussianity constraints on isocurvature perturbations in the curvaton model of inflation,” *Phys. Rev.* **D69** (2004) 063508, [arXiv:astro-ph/0311102](#) [astro-ph].
- [75] E. Di Valentino, M. Lattanzi, G. Mangano, A. Melchiorri, and P. Serpico, “Future constraints on neutrino isocurvature perturbations in the curvaton scenario,” *Phys.Rev.* **D85** (2012) 043511, [arXiv:1111.3810](#).
- [76] T. L. Smith and D. Grin, “Probing A Panoply of Curvaton-Decay Scenarios Using CMB Data,” [arXiv:1511.07431](#) [astro-ph.CO].
- [77] D. Grin, O. Doré, and M. Kamionkowski, “Do baryons trace dark matter in the early universe?,” *Phys. Rev. Lett.* **107** (2011) 261301, [arXiv:1107.1716](#) [astro-ph.CO].
- [78] D. Grin, O. Doré, and M. Kamionkowski, “Compensated Isocurvature Perturbations and the Cosmic Microwave Background,” *Phys. Rev.* **D84** (2011) 123003, [arXiv:1107.5047](#) [astro-ph.CO].
- [79] D. Grin, D. Hanson, G. P. Holder, O. Doré, and M. Kamionkowski, “Baryons do trace dark matter 380,000 years after the big bang: Search for compensated isocurvature perturbations with WMAP 9-year data,” *Phys. Rev.* **D89** no. 2, (2014) 023006, [arXiv:1306.4319](#) [astro-ph.CO].
- [80] C. He, D. Grin, and W. Hu, “Compensated isocurvature perturbations in the curvaton model,” *Phys. Rev.* **D92** no. 6, (2015) 063018, [arXiv:1505.00639](#) [astro-ph.CO].

- [81] T. L. Smith, J. B. Muñoz, R. Smith, K. Yee, and D. Grin, “Baryons still trace dark matter: probing CMB lensing maps for hidden isocurvature,” *Phys. Rev.* **D96** no. 8, (2017) 083508, [arXiv:1704.03461](#) [[astro-ph.CO](#)].
- [82] P. D. Meerburg *et al.*, “Primordial Non-Gaussianity,” [arXiv:1903.04409](#) [[astro-ph.CO](#)].
- [83] L. Senatore, K. M. Smith, and M. Zaldarriaga, “Non-Gaussianities in Single Field Inflation and their Optimal Limits from the WMAP 5-year Data,” *JCAP* **1001** (2010) 028, [arXiv:0905.3746](#) [[astro-ph.CO](#)].
- [84] C. Bennett, D. Larson, J. Weiland, N. Jarosik, G. Hinshaw, *et al.*, “Nine-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Final Maps and Results,” [arXiv:1212.5225](#) [[astro-ph.CO](#)].
- [85] N. Arkani-Hamed and J. Maldacena, “Cosmological Collider Physics,” [arXiv:1503.08043](#) [[hep-th](#)].
- [86] X. Chen, M. H. Namjoo, and Y. Wang, “Quantum Primordial Standard Clocks,” *JCAP* **1602** no. 02, (2016) 013, [arXiv:1509.03930](#) [[astro-ph.CO](#)].
- [87] N. Arkani-Hamed, D. Baumann, H. Lee, and G. L. Pimentel, “The Cosmological Bootstrap: Inflationary Correlators from Symmetries and Singularities,” [arXiv:1811.00024](#) [[hep-th](#)].
- [88] N. Barnaby and M. Peloso, “Large Nongaussianity in Axion Inflation,” *Phys. Rev. Lett.* **106** (2011) 181301, [arXiv:1011.1500](#) [[hep-ph](#)].
- [89] N. Barnaby and S. Shandera, “Feeding your Inflaton: Non-Gaussian Signatures of Interaction Structure,” *JCAP* **1201** (2012) 034, [arXiv:1109.2985](#) [[astro-ph.CO](#)].
- [90] N. Barnaby, J. Moxon, R. Namba, M. Peloso, G. Shiu, and P. Zhou, “Gravity waves and non-Gaussian features from particle production in a sector gravitationally coupled to the inflaton,” *Phys. Rev.* **D86** (2012) 103508, [arXiv:1206.6117](#) [[astro-ph.CO](#)].
- [91] R. Namba, E. Dimastrogiovanni, and M. Peloso, “Gauge-flation confronted with Planck,” *JCAP* **1311** (2013) 045, [arXiv:1308.1366](#) [[astro-ph.CO](#)].
- [92] E. Silverstein and D. Tong, “Scalar speed limits and cosmology: Acceleration from D-cceleration,” *Phys. Rev.* **D70** (2004) 103505, [arXiv:hep-th/0310221](#) [[hep-th](#)].
- [93] N. Arkani-Hamed, P. Creminelli, S. Mukohyama, and M. Zaldarriaga, “Ghost inflation,” *JCAP* **0404** (2004) 001, [arXiv:hep-th/0312100](#) [[hep-th](#)].
- [94] M. Alishahiha, E. Silverstein, and D. Tong, “DBI in the sky,” *Phys. Rev.* **D70** (2004) 123505, [arXiv:hep-th/0404084](#) [[hep-th](#)].
- [95] X. Chen, M.-x. Huang, S. Kachru, and G. Shiu, “Observational signatures and non-Gaussianities of general single field inflation,” *JCAP* **0701** (2007) 002, [arXiv:hep-th/0605045](#) [[hep-th](#)].
- [96] C. Cheung, P. Creminelli, A. L. Fitzpatrick, J. Kaplan, and L. Senatore, “The Effective Field Theory of Inflation,” *JHEP* **03** (2008) 014, [arXiv:0709.0293](#) [[hep-th](#)].
- [97] X. Chen and Y. Wang, “Quasi-Single Field Inflation and Non-Gaussianities,” *JCAP* **1004** (2010) 027, [arXiv:0911.3380](#) [[hep-th](#)].
- [98] A. J. Tolley and M. Wyman, “The Gelaton Scenario: Equilateral non-Gaussianity from multi-field dynamics,” *Phys. Rev.* **D81** (2010) 043502, [arXiv:0910.1853](#) [[hep-th](#)].

- [99] S. Cremonini, Z. Lalak, and K. Turzynski, “Strongly Coupled Perturbations in Two-Field Inflationary Models,” *JCAP* **1103** (2011) 016, [arXiv:1010.3021 \[hep-th\]](#).
- [100] A. Achucarro, J.-O. Gong, S. Hardeman, G. A. Palma, and S. P. Patil, “Features of heavy physics in the CMB power spectrum,” *JCAP* **1101** (2011) 030, [arXiv:1010.3693 \[hep-ph\]](#).
- [101] D. Baumann and D. Green, “Signatures of Supersymmetry from the Early Universe,” *Phys. Rev.* **D85** (2012) 103520, [arXiv:1109.0292 \[hep-th\]](#).
- [102] P. Creminelli, “On non-Gaussianities in single-field inflation,” *JCAP* **0310** (2003) 003, [arXiv:astro-ph/0306122 \[astro-ph\]](#).
- [103] D. Baumann, D. Green, and R. A. Porto, “B-modes and the Nature of Inflation,” *JCAP* **1501** no. 01, (2015) 016, [arXiv:1407.2621 \[hep-th\]](#).
- [104] M. Alvarez *et al.*, “Testing Inflation with Large Scale Structure: Connecting Hopes with Reality,” [arXiv:1412.4671 \[astro-ph.CO\]](#).
- [105] D. Baumann, D. Green, H. Lee, and R. A. Porto, “Signs of Analyticity in Single-Field Inflation,” *Phys. Rev.* **D93** no. 2, (2016) 023523, [arXiv:1502.07304 \[hep-th\]](#).
- [106] D. Baumann and D. Green, “Equilateral Non-Gaussianity and New Physics on the Horizon,” *JCAP* **1109** (2011) 014, [arXiv:1102.5343 \[hep-th\]](#).
- [107] D. Green, M. Lewandowski, L. Senatore, E. Silverstein, and M. Zaldarriaga, “Anomalous Dimensions and Non-Gaussianity,” *JHEP* **10** (2013) 171, [arXiv:1301.2630 \[hep-th\]](#).
- [108] V. Assassi, D. Baumann, D. Green, and L. McAllister, “Planck-Suppressed Operators,” *JCAP* **1401** (2014) 033, [arXiv:1304.5226 \[hep-th\]](#).
- [109] P. Creminelli and M. Zaldarriaga, “Single-Field Consistency Relation for the 3-Point Function,” *JCAP* **0410** (2004) 006, [arXiv:astro-ph/0407059 \[astro-ph\]](#).
- [110] E. Nelson and S. Shandera, “Statistical Naturalness and non-Gaussianity in a Finite Universe,” *Phys. Rev. Lett.* **110** no. 13, (2013) 131301, [arXiv:1212.4550 \[astro-ph.CO\]](#).
- [111] M. LoVerde, E. Nelson, and S. Shandera, “Non-Gaussian Mode Coupling and the Statistical Cosmological Principle,” *JCAP* **1306** (2013) 024, [arXiv:1303.3549 \[astro-ph.CO\]](#).
- [112] S. Nurmi, C. T. Byrnes, and G. Tasinato, “A non-Gaussian landscape,” *JCAP* **1306** (2013) 004, [arXiv:1301.3128 \[astro-ph.CO\]](#).
- [113] B. Bonga, S. Brahma, A.-S. Deutsch, and S. Shandera, “Cosmic variance in inflation with two light scalars,” *JCAP* **1605** no. 05, (2016) 018, [arXiv:1512.05365 \[astro-ph.CO\]](#).
- [114] M. Zaldarriaga, “Non-Gaussianities in models with a varying inflaton decay rate,” *Phys. Rev.* **D69** (2004) 043508, [arXiv:astro-ph/0306006 \[astro-ph\]](#).
- [115] J.-L. Lehners and S. Renaux-Petel, “Multifield Cosmological Perturbations at Third Order and the Ekpyrotic Trispectrum,” *Phys. Rev.* **D80** (2009) 063503, [arXiv:0906.0530 \[hep-th\]](#).
- [116] J. C. Hill, “Foreground Biases on Primordial Non-Gaussianity Measurements from the CMB Temperature Bispectrum: Implications for Planck and Beyond,” [arXiv:1807.07324 \[astro-ph.CO\]](#).
- [117] D. Babich and M. Zaldarriaga, “Primordial bispectrum information from cmb polarization,” *Phys. Rev. D* **70** (Oct, 2004) 083005. <https://link.aps.org/doi/10.1103/PhysRevD.70.083005>.

- [118] E. Silverstein and A. Westphal, “Monodromy in the CMB: Gravity Waves and String Inflation,” *Phys.Rev.* **D78** (2008) 106003, [arXiv:0803.3085 \[hep-th\]](#).
- [119] L. McAllister, E. Silverstein, and A. Westphal, “Gravity Waves and Linear Inflation from Axion Monodromy,” *Phys.Rev.* **D82** (2010) 046003, [arXiv:0808.0706 \[hep-th\]](#).
- [120] R. Flauger, L. McAllister, E. Pajer, A. Westphal, and G. Xu, “Oscillations in the CMB from Axion Monodromy Inflation,” *JCAP* **1006** (2010) 009, [arXiv:0907.2916 \[hep-th\]](#).
- [121] S. R. Behbahani, A. Dymarsky, M. Mirbabayi, and L. Senatore, “(Small) Resonant non-Gaussianities: Signatures of a Discrete Shift Symmetry in the Effective Field Theory of Inflation,” *JCAP* **1212** (2012) 036, [arXiv:1111.3373 \[hep-th\]](#).
- [122] S. R. Behbahani and D. Green, “Collective Symmetry Breaking and Resonant Non-Gaussianity,” *JCAP* **1211** (2012) 056, [arXiv:1207.2779 \[hep-th\]](#).
- [123] E. Dimastrogiovanni, M. Fasiello, and M. Kamionkowski, “Imprints of Massive Primordial Fields on Large-Scale Structure,” *JCAP* **1602** (2016) 017, [arXiv:1504.05993 \[astro-ph.CO\]](#).
- [124] S. Endlich, A. Nicolis, and J. Wang, “Solid Inflation,” *JCAP* **1310** (2013) 011, [arXiv:1210.0569 \[hep-th\]](#).
- [125] A. Maleknejad and M. M. Sheikh-Jabbari, “Gauge-flation: Inflation From Non-Abelian Gauge Fields,” *Phys. Lett.* **B723** (2013) 224–228, [arXiv:1102.1513 \[hep-ph\]](#).
- [126] P. Adshead, D. Blas, C. P. Burgess, P. Hayman, and S. P. Patil, “Magnon Inflation: Slow Roll with Steep Potentials,” [arXiv:1604.06048 \[hep-th\]](#).
- [127] **Planck** Collaboration, Planck Collaboration VI, “Planck 2018 results. VI. Cosmological parameters,” [arXiv:1807.06209 \[astro-ph.CO\]](#).
- [128] A. Lewis, A. Challinor, and D. Hanson, “The shape of the CMB lensing bispectrum,” *JCAP* **1103** (2011) 018, [arXiv:1101.2234 \[astro-ph.CO\]](#).
- [129] LSST Science Collaboration, P. A. Abell, J. Allison, S. F. Anderson, J. R. Andrew, J. R. P. Angel, L. Armus, D. Arnett, S. J. Asztalos, T. S. Axelrod, S. Bailey, D. R. Ballantyne, J. R. Bankert, W. A. Barkhouse, J. D. Barr, L. F. Barrientos, A. J. Barth, J. G. Bartlett, A. C. Becker, J. Becla, T. C. Beers, J. P. Bernstein, R. Biswas, M. R. Blanton, J. S. Bloom, J. J. Bochanski, P. Boeshaar, K. D. Borne, M. Bradac, W. N. Brandt, C. R. Bridge, M. E. Brown, R. J. Brunner, J. S. Bullock, A. J. Burgasser, J. H. Burge, D. L. Burke, P. A. Cargile, S. Chandrasekharan, G. Chartas, S. R. Chesley, Y.-H. Chu, D. Cinabro, M. W. Claire, C. F. Claver, D. Clowe, A. J. Connolly, K. H. Cook, J. Cooke, A. Cooray, K. R. Covey, C. S. Culliton, R. de Jong, W. H. de Vries, V. P. Debattista, F. Delgado, I. P. Dell’Antonio, S. Dhital, R. Di Stefano, M. Dickinson, B. Dilday, S. G. Djorgovski, G. Dobler, C. Donalek, G. Dubois-Felsmann, J. Durech, A. Eliasdottir, M. Eracleous, L. Eyer, E. E. Falco, X. Fan, C. D. Fassnacht, H. C. Ferguson, Y. R. Fernandez, B. D. Fields, D. Finkbeiner, E. E. Figueroa, D. B. Fox, H. Francke, J. S. Frank, J. Frieman, S. Fromenteau, M. Furqan, G. Galaz, A. Gal-Yam, P. Garnavich, E. Gawiser, J. Geary, P. Gee, R. R. Gibson, K. Gilmore, E. A. Grace, R. F. Green, W. J. Gressler, C. J. Grillmair, S. Habib, J. S. Haggerty, M. Hamuy, A. W. Harris, S. L. Hawley, A. F. Heavens, L. Hebb, T. J. Henry, E. Hileman, E. J. Hilton, K. Hoadley, J. B. Holberg, M. J. Holman, S. B. Howell, L. Infante, Z. Ivezic, S. H. Jacoby, B. Jain, R. Jedicke, M. J. Jee, J. G. Jernigan, S. W. Jha, K. V. Johnston, R. L. Jones, M. Juric, M. Kaasalainen, Styliani, Kafka, S. M. Kahn, N. A. Kaib, J. Kalirai, J. Kantor, M. M. Kasliwal, C. R. Keeton, R. Kessler, Z. Knezevic, A. Kowalski, V. L. Krabbendam, K. S. Krughoff, S. Kulkarni, S. Kuhlman, M. Lacy,

- S. Lepine, M. Liang, A. Lien, P. Lira, K. S. Long, S. Lorenz, J. M. Lotz, R. H. Lupton, J. Lutz, L. M. Macri, A. A. Mahabal, R. Mandelbaum, P. Marshall, M. May, P. M. McGehee, B. T. Meadows, A. Meert, A. Milani, C. J. Miller, M. Miller, D. Mills, D. Minniti, D. Monet, A. S. Mukadam, E. Nakar, D. R. Neill, J. A. Newman, S. Nikolaev, M. Nordby, P. O'Connor, M. Oguri, J. Oliver, S. S. Olivier, J. K. Olsen, K. Olsen, E. W. Olszewski, H. Oluseyi, N. D. Padilla, A. Parker, J. Pepper, J. R. Peterson, C. Petry, P. A. Pinto, J. L. Pizagno, B. Popescu, A. Prsa, V. Radcka, M. J. Raddick, A. Rasmussen, A. Rau, J. Rho, J. E. Rhoads, G. T. Richards, S. T. Ridgway, B. E. Robertson, R. Roskar, A. Saha, A. Sarajedini, E. Scannapieco, T. Schalk, R. Schindler, S. Schmidt, S. Schmidt, D. P. Schneider, G. Schumacher, R. Scranton, J. Sebag, L. G. Seppala, O. Shemmer, J. D. Simon, M. Sivertz, H. A. Smith, J. A. Smith, N. Smith, A. H. Spitz, A. Stanford, K. G. Stassun, J. Strader, M. A. Strauss, C. W. Stubbs, D. W. Sweeney, A. Szalay, P. Szkody, M. Takada, P. Thorman, D. E. Trilling, V. Trimble, A. Tyson, R. Van Berg, D. V. Berk, J. VanderPlas, L. Verde, B. Vrsnak, L. M. Walkowicz, B. D. Wandelt, S. Wang, Y. Wang, M. Warner, R. H. Wechsler, A. A. West, O. Wiecha, B. F. Williams, B. Willman, D. Wittman, S. C. Wolff, W. M. Wood-Vasey, P. Wozniak, P. Young, A. Zentner, and H. Zhan, "LSST Science Book, Version 2.0," *Science* no. November, (Dec, 2009) 596, [arXiv:0912.0201](#).
- [130] N. Dalal, O. Dore, D. Huterer, and A. Shirokov, "The imprints of primordial non-gaussianities on large-scale structure: scale dependent bias and abundance of virialized objects," *Phys.Rev.* **D77** (2008) 123514, [arXiv:0710.4560](#) [[astro-ph](#)].
- [131] U. Seljak, "Extracting primordial non-gaussianity without cosmic variance," *Phys.Rev.Lett.* **102** (2009) 021302, [arXiv:0807.1770](#) [[astro-ph](#)].
- [132] M. Schmittfull and U. Seljak, "Parameter constraints from cross-correlation of CMB lensing with galaxy clustering," *Phys. Rev. D* **97** no. 12, (June, 2018) 123540, [arXiv:1710.09465](#).
- [133] M. Münchmeyer, M. S. Madhavacheril, S. Ferraro, M. C. Johnson, and K. M. Smith, "Constraining local non-Gaussianities with kSZ tomography," [arXiv:1810.13424](#) [[astro-ph.CO](#)].
- [134] A.-S. Deutsch, E. Dimastrogiovanni, M. C. Johnson, M. Munchmeyer, and A. Terrana, "Reconstruction of the remote dipole and quadrupole fields from the kinetic Sunyaev Zel'dovich and polarized Sunyaev Zel'dovich effects," [arXiv:1707.08129](#) [[astro-ph.CO](#)].
- [135] K. M. Smith, M. S. Madhavacheril, M. Münchmeyer, S. Ferraro, U. Giri, and M. C. Johnson, "KSZ tomography and the bispectrum," [arXiv:1810.13423](#) [[astro-ph.CO](#)].
- [136] J. Aguirre *et al.* (Simons Observatory Collaboration), "The Simons Observatory: Science Goals and Forecasts," *JCAP* **02** (2019) 056, [arXiv:1808.07445](#) [[astro-ph.CO](#)].
- [137] D. Jeong, E. Komatsu, and B. Jain, "Galaxy-CMB and galaxy-galaxy lensing on large scales: Sensitivity to primordial non-Gaussianity," *Phys. Rev. D* **80** no. 12, (Dec., 2009) 123527, [arXiv:0910.1361](#).
- [138] Y. Ono, M. Ouchi, Y. Harikane, J. Toshikawa, M. Rauch, S. Yuma, M. Sawicki, T. Shibuya, K. Shimasaku, M. Oguri, C. Willott, M. Akhlaghi, M. Akiyama, J. Coupon, N. Kashikawa, Y. Komiyama, A. Konno, L. Lin, Y. Matsuoka, S. Miyazaki, T. Nagao, K. Nakajima, J. Silverman, M. Tanaka, Y. Taniguchi, and S.-Y. Wang, "Great Optically Luminous Dropout Research Using Subaru HSC (GOLDRUSH). I. UV luminosity functions at $z \sim 4-7$ derived with the half-million dropouts on the 100 deg² sky," *PASJ* **70** (Jan., 2018) S10, [arXiv:1704.06004](#).
- [139] Y. Harikane, M. Ouchi, Y. Ono, S. Saito, P. Behroozi, S. More, K. Shimasaku, J. Toshikawa, Y.-T. Lin, M. Akiyama, J. Coupon, Y. Komiyama, A. Konno, S.-C. Lin, S. Miyazaki, A. J. Nishizawa,

- T. Shibuya, and J. Silverman, “GOLDRUSH. II. Clustering of galaxies at $z \sim 4-6$ revealed with the half-million dropouts over the 100 deg^2 area corresponding to 1 Gpc^3 ,” *PASJ* **70** (Jan., 2018) S11, [arXiv:1704.06535](#).
- [140] J. S. Dunlop, “Observing the First Galaxies,” *ArXiv e-prints* **396** (2013) 223, [arXiv:1205.1543](#).
- [141] D. J. E. Marsh, “Axion Cosmology,” *Phys. Rept.* **643** (2016) 1–79, [arXiv:1510.07633](#) [[astro-ph.CO](#)].
- [142] K. Abazajian, G. M. Fuller, and M. Patel, “Sterile neutrino hot, warm, and cold dark matter,” *Phys. Rev.* **D64** (2001) 023501, [arXiv:astro-ph/0101524](#) [[astro-ph](#)].
- [143] A. Strumia and F. Vissani, “Neutrino masses and mixings and...,” [arXiv:hep-ph/0606054](#) [[hep-ph](#)].
- [144] L. Ackerman, M. R. Buckley, S. M. Carroll, and M. Kamionkowski, “Dark Matter and Dark Radiation,” *Phys. Rev. D* **79** (2009) 023519, [arXiv:0810.5126](#) [[hep-ph](#)].
- [145] A. Boyarsky, O. Ruchayskiy, and M. Shaposhnikov, “The Role of sterile neutrinos in cosmology and astrophysics,” *Ann. Rev. Nucl. Part. Sci.* **59** (2009) 191–214, [arXiv:0901.0011](#) [[hep-ph](#)].
- [146] A. Arvanitaki, S. Dimopoulos, S. Dubovsky, N. Kaloper, and J. March-Russell, “String Axiverse,” *Phys. Rev.* **D81** (2010) 123530, [arXiv:0905.4720](#) [[hep-th](#)].
- [147] D. Cadamuro, S. Hannestad, G. Raffelt, and J. Redondo, “Cosmological bounds on sub-MeV mass axions,” *JCAP* **1102** (2011) 003, [arXiv:1011.3694](#) [[hep-ph](#)].
- [148] D. E. Kaplan, G. Z. Krnjaic, K. R. Rehermann, and C. M. Wells, “Dark Atoms: Asymmetry and Direct Detection,” *JCAP* **1110** (2011) 011, [arXiv:1105.2073](#) [[hep-ph](#)].
- [149] F.-Y. Cyr-Racine and K. Sigurdson, “Cosmology of atomic dark matter,” *Phys. Rev.* **D87** no. 10, (2013) 103515, [arXiv:1209.5752](#) [[astro-ph.CO](#)].
- [150] C. Brust, D. E. Kaplan, and M. T. Walters, “New Light Species and the CMB,” *JHEP* **12** (2013) 058, [arXiv:1303.5379](#) [[hep-ph](#)].
- [151] S. Weinberg, “Goldstone Bosons as Fractional Cosmic Neutrinos,” *Phys.Rev.Lett.* **110** (2013) 241301, [arXiv:1305.1971](#) [[astro-ph.CO](#)].
- [152] A. Salvio, A. Strumia, and W. Xue, “Thermal axion production,” *JCAP* **1401** (2014) 011, [arXiv:1310.6982](#) [[hep-ph](#)].
- [153] M. Kawasaki, M. Yamada, and T. T. Yanagida, “Observable dark radiation from a cosmologically safe QCD axion,” *Phys. Rev.* **D91** no. 12, (2015) 125018, [arXiv:1504.04126](#) [[hep-ph](#)].
- [154] P. W. Graham, D. E. Kaplan, and S. Rajendran, “Cosmological Relaxation of the Electroweak Scale,” *Phys. Rev. Lett.* **115** no. 22, (2015) 221801, [arXiv:1504.07551](#) [[hep-ph](#)].
- [155] D. Baumann, D. Green, and B. Wallisch, “A New Target for Cosmic Axion Searches,” [arXiv:1604.08614](#) [[astro-ph.CO](#)].
- [156] N. Arkani-Hamed, T. Cohen, R. T. D’Agnolo, A. Hook, H. D. Kim, and D. Pinner, “Nnaturalness,” [arXiv:1607.06821](#) [[hep-ph](#)].
- [157] Z. Chacko, N. Craig, P. Fox, and R. Harnik, “Cosmology in Mirror Twin Higgs and Neutrino Masses,” *JHEP* **07** (2017) 023, [arXiv:1611.07975](#) [[hep-ph](#)].

- [158] N. Craig, S. Koren, and T. Trott, “Cosmological Signals of a Mirror Twin Higgs,” *JHEP* **05** (2017) 038, [arXiv:1611.07977 \[hep-ph\]](#).
- [159] Z. Chacko, D. Curtin, M. Geller, and Y. Tsai, “Cosmological Signatures of a Mirror Twin Higgs,” *JHEP* **09** (2018) 163, [arXiv:1803.03263 \[hep-ph\]](#).
- [160] M. Archidiacono, E. Giusarma, S. Hannestad, and O. Mena, “Cosmic Dark Radiation and Neutrinos,” *Adv. High Energy Phys.* **2013** (2013) 191047, [arXiv:1307.0637 \[astro-ph.CO\]](#).
- [161] J. Bernal, L. Verde, and A. Riess, “The Trouble with H_0 ,” *JCAP* **10** (2016) 019, [arXiv:1607.05617 \[astro-ph.CO\]](#).
- [162] B. Zhang, M. Childress, T. Davis, N. Karpenka, C. Lidman, B. Schmidt, and M. Smith, “A Blinded Determination of H_0 from Low-Redshift Type Ia Supernovae, Calibrated by Cepheid Variables,” *Mon. Not. Roy. Astron. Soc.* **471** (2017) 2254, [arXiv:1706.07573 \[astro-ph.CO\]](#).
- [163] G. Addison, D. Watts, C. Bennett, M. Halpern, G. Hinshaw, and J. Weiland, “Elucidating Λ CDM: Impact of Baryon Acoustic Oscillation Measurements on the Hubble Constant Discrepancy,” *Astrophys. J.* **853** (2018) 119, [arXiv:1707.06547 \[astro-ph.CO\]](#).
- [164] K. Aylor, M. Joy, L. Knox, M. Millea, S. Raghunathan, and K. Wu, “Sounds Discordant: Classical Distance Ladder & Λ CDM-Based Determinations of the Cosmological Sound Horizon,” [arXiv:1811.00537 \[astro-ph.CO\]](#).
- [165] N. MacCrann, J. Zuntz, S. Bridle, B. Jain, and M. Becker, “Cosmic Discordance: Are Planck CMB and CFHTLenS Weak Lensing Measurements Out of Tune?,” *Mon. Not. Roy. Astron. Soc.* **451** (2015) 2877, [arXiv:1408.4742 \[astro-ph.CO\]](#).
- [166] J. Lesgourgues, G. Marques-Tavares, and M. Schmaltz, “Evidence for dark matter interactions in cosmological precision data?,” *JCAP* **1602** no. 02, (2016) 037, [arXiv:1507.04351 \[astro-ph.CO\]](#).
- [167] F. Köhlinger *et al.*, “KiDS-450: The Tomographic Weak Lensing Power Spectrum and Constraints on Cosmological Parameters,” *Mon. Not. Roy. Astron. Soc.* **471** (2017) 4412, [arXiv:1706.02892 \[astro-ph.CO\]](#).
- [168] S. Joudaki *et al.*, “KiDS-450 + 2dFLenS: Cosmological Parameter Constraints from Weak Gravitational Lensing Tomography and Overlapping Redshift-Space Galaxy Clustering,” *Mon. Not. Roy. Astron. Soc.* **474** (2018) 4894, [arXiv:1707.06627 \[astro-ph.CO\]](#).
- [169] D. Weinberg, J. Bullock, F. Governato, R. Kuzio de Naray, and A. Peter, “Cold Dark Matter: Controversies on Small Scales,” *Proc. Nat. Acad. Sci.* **112** (2015) 12249, [arXiv:1306.0913 \[astro-ph.CO\]](#).
- [170] G. Mangano, G. Miele, S. Pastor, T. Pinto, O. Pisanti, *et al.*, “Relic neutrino decoupling including flavor oscillations,” *Nucl.Phys.* **B729** (2005) 221–234, [arXiv:hep-ph/0506164 \[hep-ph\]](#).
- [171] E. Grohs, G. M. Fuller, C. T. Kishimoto, M. W. Paris, and A. Vlasenko, “Neutrino energy transport in weak decoupling and big bang nucleosynthesis,” *Phys. Rev.* **D93** no. 8, (2016) 083522, [arXiv:1512.02205 \[astro-ph.CO\]](#).
- [172] P. F. de Salas and S. Pastor, “Relic neutrino decoupling with flavour oscillations revisited,” *JCAP* **1607** no. 07, (2016) 051, [arXiv:1606.06986 \[hep-ph\]](#).
- [173] P. J. E. Peebles, “Primordial Helium Abundance and the Primordial Fireball. II,” *Astrophys. J.* **146** (1966) 542.

- [174] D. Dicus, E. Kolb, Gleeson, Sudarshan, V. Teplitz, and M. Turner, “Primordial Nucleosynthesis Including Radiative, Coulomb, and Finite-Temperature Corrections to Weak Rates,” *Phys. Rev. D* **26** (1982) 2694.
- [175] Z. Hou, R. Keisler, L. Knox, M. Millea, and C. Reichardt, “How Massless Neutrinos Affect the Cosmic Microwave Background Damping Tail,” *Phys. Rev.* **D87** (2013) 083008, [arXiv:1104.2333](#) [[astro-ph.CO](#)].
- [176] S. Bashinsky and U. Seljak, “Neutrino perturbations in CMB anisotropy and matter clustering,” *Phys. Rev.* **D69** (2004) 083002, [arXiv:astro-ph/0310198](#) [[astro-ph](#)].
- [177] D. Baumann, D. Green, J. Meyers, and B. Wallisch, “Phases of New Physics in the CMB,” *JCAP* **1601** (2016) 007, [arXiv:1508.06342](#) [[astro-ph.CO](#)].
- [178] D. Baumann, D. Green, and M. Zaldarriaga, “Phases of New Physics in the BAO Spectrum,” *JCAP* **11** (2017) 007, [arXiv:1703.00894](#) [[astro-ph.CO](#)].
- [179] N. Bell, E. Pierpaoli, and K. Sigurdson, “Cosmological Signatures of Interacting Neutrinos,” *Phys. Rev. D* **73** (2006) 063523, [arXiv:astro-ph/0511410](#) [[astro-ph](#)].
- [180] A. Friedland, K. Zurek, and S. Bashinsky, “Constraining Models of Neutrino Mass and Neutrino Interactions with the Planck Satellite,” [arXiv:0704.3271](#) [[astro-ph](#)].
- [181] F.-Y. Cyr-Racine and K. Sigurdson, “Limits on Neutrino-Neutrino Scattering in the Early Universe,” *Phys. Rev.* **D90** no. 12, (2014) 123533, [arXiv:1306.1536](#) [[astro-ph.CO](#)].
- [182] I. M. Oldengott, C. Rampf, and Y. Y. Y. Wong, “Boltzmann hierarchy for interacting neutrinos I: formalism,” *JCAP* **1504** no. 04, (2015) 016, [arXiv:1409.1577](#) [[astro-ph.CO](#)].
- [183] L. Lancaster, F.-Y. Cyr-Racine, L. Knox, and Z. Pan, “A Tale of Two Modes: Neutrino Free-Streaming in the Early Universe,” *JCAP* **07** (2017) 033, [arXiv:1704.06657](#) [[astro-ph.CO](#)].
- [184] I. M. Oldengott, T. Tram, C. Rampf, and Y. Y. Y. Wong, “Interacting neutrinos in cosmology: exact description and constraints,” *JCAP* **11** (Nov., 2017) 027, [arXiv:1706.02123](#).
- [185] C. D. Kreisch, F.-Y. Cyr-Racine, and O. Doré, “The Neutrino Puzzle: Anomalies, Interactions, and Cosmological Tensions,” [arXiv:1902.00534](#) [[astro-ph.CO](#)].
- [186] R. J. Wilkinson, C. Boehm, and J. Lesgourgues, “Constraining Dark Matter-Neutrino Interactions using the CMB and Large-Scale Structure,” *JCAP* **1405** (2014) 011, [arXiv:1401.7597](#) [[astro-ph.CO](#)].
- [187] M. Escudero, O. Mena, A. C. Vincent, R. J. Wilkinson, and C. Boehm, “Exploring dark matter microphysics with galaxy surveys,” *jcap* **1509** no. 09, (2015) 034, [arXiv:1505.06735](#) [[astro-ph.CO](#)].
- [188] M. A. Buen-Abad, G. Marques-Tavares, and M. Schmaltz, “Non-Abelian dark matter and dark radiation,” *Phys. Rev.* **D92** no. 2, (2015) 023531, [arXiv:1505.03542](#) [[hep-ph](#)].
- [189] Z. Chacko, Y. Cui, S. Hong, T. Okui, and Y. Tsai, “Partially Acoustic Dark Matter, Interacting Dark Radiation, and Large Scale Structure,” *JHEP* **12** (2016) 108, [arXiv:1609.03569](#) [[astro-ph.CO](#)].
- [190] M. A. Buen-Abad, M. Schmaltz, J. Lesgourgues, and T. Brinckmann, “Interacting Dark Sector and Precision Cosmology,” *JCAP* **1801** no. 01, (2018) 008, [arXiv:1708.09406](#) [[astro-ph.CO](#)].
- [191] C. Brust, Y. Cui, and K. Sigurdson, “Cosmological Constraints on Interacting Light Particles,” *JCAP* **08** (2017) 020, [arXiv:1703.10732](#) [[astro-ph.CO](#)].

- [192] G. Choi, C.-T. Chiang, and M. LoVerde, “Probing Decoupling in Dark Sectors with the Cosmic Microwave Background,” *JCAP* **1806** no. 06, (2018) 044, [arXiv:1804.10180](#) [[astro-ph.CO](#)].
- [193] S. Dodelson, K. Heitmann, C. Hirata, K. Honscheid, A. Roodman, U. Seljak, A. Slosar, and M. Trodden, “Cosmic Visions Dark Energy: Science,” [arXiv:1604.07626](#) [[astro-ph.CO](#)].
- [194] Z. Chacko, Y. Cui, S. Hong, and T. Okui, “Hidden dark matter sector, dark radiation, and the CMB,” *Phys. Rev.* **D92** (2015) 055033, [arXiv:1505.04192](#) [[hep-ph](#)].
- [195] M. Battaglieri *et al.*, “US Cosmic Visions: New Ideas in Dark Matter 2017: Community Report,” in *U.S. Cosmic Visions: New Ideas in Dark Matter College Park, MD, USA, March 23-25, 2017*. 2017. [arXiv:1707.04591](#) [[hep-ph](#)]. <http://lss.fnal.gov/archive/2017/conf/fermilab-conf-17-282-ae-ppd-t.pdf>.
- [196] D. Green and S. Rajendran, “The Cosmology of Sub-MeV Dark Matter,” *JHEP* **10** (2017) 013, [arXiv:1701.08750](#) [[hep-ph](#)].
- [197] S. Knapen, T. Lin, and K. Zurek, “Light Dark Matter: Models and Constraints,” *Phys. Rev. D* **96** (2017) 115021, [arXiv:1709.07882](#) [[hep-ph](#)].
- [198] V. Gluscevic and K. Boddy, “Constraints on Scattering of keV–TeV Dark Matter with Protons in the Early Universe,” *Phys. Rev. Lett.* **121** (2018) 081301, [arXiv:1712.07133](#) [[astro-ph.CO](#)].
- [199] M. Markevitch, A. H. Gonzalez, D. Clowe, A. Vikhlinin, L. David, W. Forman, C. Jones, S. Murray, and W. Tucker, “Direct constraints on the dark matter self-interaction cross-section from the merging galaxy cluster 1E0657-56,” *Astrophys. J.* **606** (2004) 819–824, [arXiv:astro-ph/0309303](#) [[astro-ph](#)].
- [200] R. Essig, R. Harnik, J. Kaplan, and N. Toro, “Discovering New Light States at Neutrino Experiments,” *Phys. Rev. D* **82** (2010) 113008, [arXiv:1008.0636](#) [[hep-ph](#)].
- [201] B. Wallisch, “Cosmological Probes of Light Relics,” [arXiv:1810.02800](#) [[astro-ph.CO](#)].
- [202] P. Adshead, Y. Cui, and J. Shelton, “Chilly Dark Sectors and Asymmetric Reheating,” *JHEP* **06** (2016) 016, [arXiv:1604.02458](#) [[hep-ph](#)].
- [203] L. A. Boyle and A. Buonanno, “Relating gravitational wave constraints from primordial nucleosynthesis, pulsar timing, laser interferometers, and the CMB: Implications for the early Universe,” *Phys. Rev.* **D78** (2008) 043531, [arXiv:0708.2279](#) [[astro-ph](#)].
- [204] A. Stewart and R. Brandenberger, “Observational Constraints on Theories with a Blue Spectrum of Tensor Modes,” *JCAP* **0808** (2008) 012, [arXiv:0711.4602](#) [[astro-ph](#)].
- [205] P. D. Meerburg, R. Hložek, B. Hadzhiyska, and J. Meyers, “Multiwavelength constraints on the inflationary consistency relation,” *Phys. Rev.* **D91** no. 10, (2015) 103505, [arXiv:1502.00302](#) [[astro-ph.CO](#)].
- [206] M. Maggiore, “Gravitational Wave Experiments and Early Universe Cosmology,” *Phys. Rept.* **331** (2000) 283, [arXiv:gr-qc/9909001](#) [[gr-qc](#)].
- [207] R. Easther and E. Lim, “Stochastic Gravitational Wave Production After Inflation,” *JCAP* **04** (2006) 010, [arXiv:astro-ph/0601617](#) [[astro-ph](#)].
- [208] J. Dufaux, A. Bergman, G. Felder, L. Kofman, and J.-P. Uzan, “Theory and Numerics of Gravitational Waves from Preheating after Inflation,” *Phys. Rev. D* **76** (2007) 123517, [arXiv:0707.0875](#) [[astro-ph](#)].

- [209] M. Amin, M. Hertzberg, D. Kaiser, and J. Karouby, “Nonperturbative Dynamics of Reheating After Inflation: A Review,” *Int. J. Mod. Phys. D* **24** (2014) 1530003, [arXiv:1410.3808 \[hep-ph\]](#).
- [210] C. Caprini and D. Figueroa, “Cosmological Backgrounds of Gravitational Waves,” *Class. Quant. Grav.* **35** (2018) 163001, [arXiv:1801.04268 \[astro-ph.CO\]](#).
- [211] P. Adshead, J. Giblin, and Z. Weiner, “Gravitational Waves from Gauge Preheating,” *Phys. Rev. D* **98** (2018) 043525, [arXiv:1805.04550 \[astro-ph.CO\]](#).
- [212] M. Amin, J. Fan, K. Lozanov, and M. Reece, “Cosmological Dynamics of Higgs Potential Fine Tuning,” *Phys. Rev. D* **99** (2019) 035008, [arXiv:1802.00444 \[hep-ph\]](#).
- [213] W. Fischler and J. Meyers, “Dark Radiation Emerging After Big Bang Nucleosynthesis?,” *Phys. Rev. D* **83** (2011) 063520, [arXiv:1011.3501 \[astro-ph.CO\]](#).
- [214] J. Hasenkamp, “Dark Radiation from the Axino Solution of the Gravitino Problem,” *Phys. Lett. B* **707** (2012) 121, [arXiv:1107.4319 \[hep-ph\]](#).
- [215] D. Hooper, F. S. Queiroz, and N. Y. Gnedin, “Non-Thermal Dark Matter Mimicking An Additional Neutrino Species In The Early Universe,” *Phys. Rev. D* **85** (2012) 063513, [arXiv:1111.6599 \[astro-ph.CO\]](#).
- [216] J. Hasenkamp and J. Kersten, “Dark Radiation from Particle Decay: Cosmological Constraints and Opportunities,” *JCAP* **08** (2013) 024, [arXiv:1212.4160 \[hep-ph\]](#).
- [217] M. Zaldarriaga and D. D. Harari, “Analytic approach to the polarization of the cosmic microwave background in flat and open universes,” *Phys. Rev. D* **52** (1995) 3276–3287, [arXiv:astro-ph/9504085 \[astro-ph\]](#).
- [218] B. Follin, L. Knox, M. Millea, and Z. Pan, “First Detection of the Acoustic Oscillation Phase Shift Expected from the Cosmic Neutrino Background,” *Phys. Rev. Lett.* **115** no. 9, (2015) 091301, [arXiv:1503.07863 \[astro-ph.CO\]](#).
- [219] E. Aver, K. A. Olive, and E. D. Skillman, “The effects of He I $\lambda 10830$ on helium abundance determinations,” *JCAP* **1507** no. 07, (2015) 011, [arXiv:1503.08146 \[astro-ph.CO\]](#).
- [220] K. Saikawa and S. Shirai, “Primordial gravitational waves, precisely: The role of thermodynamics in the Standard Model,” *JCAP* **1805** no. 05, (2018) 035, [arXiv:1803.01038 \[hep-ph\]](#).
- [221] **Particle Data Group** Collaboration, K. A. Olive *et al.*, “Review of Particle Physics,” *Chin. Phys. C* **38** (2014) 090001.
- [222] J. R. Bond and A. S. Szalay, “The Collisionless Damping of Density Fluctuations in an Expanding Universe,” *Astrophys. J.* **274** (1983) 443–468.
- [223] J. Lesgourgues and S. Pastor, “Massive neutrinos and cosmology,” *Phys.Rept.* **429** (2006) 307–379, [arXiv:astro-ph/0603494 \[astro-ph\]](#).
- [224] C.-P. Ma, “Linear power spectra in cold + hot dark matter models: Analytical approximations and applications,” *Astrophys. J.* **471** (1996) 13–23, [arXiv:astro-ph/9605198 \[astro-ph\]](#).
- [225] W. Hu and D. J. Eisenstein, “Small scale perturbations in a general MDM cosmology,” *Astrophys.J.* **498** (1998) 497, [arXiv:astro-ph/9710216 \[astro-ph\]](#).
- [226] W. Hu, D. J. Eisenstein, and M. Tegmark, “Weighing neutrinos with galaxy surveys,” *Phys.Rev.Lett.* **80** (1998) 5255–5258, [arXiv:astro-ph/9712057 \[astro-ph\]](#).

- [227] **BOSS** Collaboration, S. Alam *et al.*, “The clustering of galaxies in the completed SDSS-III Baryon Oscillation Spectroscopic Survey: cosmological analysis of the DR12 galaxy sample,” *Mon. Not. Roy. Astron. Soc.* **470** no. 3, (2017) 2617–2652, [arXiv:1607.03155 \[astro-ph.CO\]](#).
- [228] **DES** Collaboration, T. M. C. Abbott *et al.*, “Dark Energy Survey year 1 results: Cosmological constraints from galaxy clustering and weak lensing,” *Phys. Rev.* **D98** no. 4, (2018) 043526, [arXiv:1708.01530 \[astro-ph.CO\]](#).
- [229] **DESI** Collaboration, A. Aghamousa *et al.*, “The DESI Experiment Part I: Science, Targeting, and Survey Design,” [arXiv:1611.00036 \[astro-ph.IM\]](#).
- [230] L. Amendola *et al.*, “Cosmology and fundamental physics with the Euclid satellite,” *Living Rev. Rel.* **21** no. 1, (2018) 2, [arXiv:1606.00180 \[astro-ph.CO\]](#).
- [231] T. Brinckmann, D. C. Hooper, M. Archidiacono, J. Lesgourgues, and T. Sprenger, “The promising future of a robust cosmological neutrino mass measurement,” *JCAP* **1901** (2019) 059, [arXiv:1808.05955 \[astro-ph.CO\]](#).
- [232] C. Kraus *et al.*, “Final results from phase II of the Mainz neutrino mass search in tritium beta decay,” *Eur. Phys. J.* **C40** (2005) 447–468, [arXiv:hep-ex/0412056 \[hep-ex\]](#).
- [233] **Troitsk** Collaboration, V. N. Aseev *et al.*, “An upper limit on electron antineutrino mass from Troitsk experiment,” *Phys. Rev.* **D84** (2011) 112003, [arXiv:1108.5034 \[hep-ex\]](#).
- [234] **KATRIN** Collaboration, J. Angrik *et al.*, “KATRIN design report 2004.”
- [235] **Project 8** Collaboration, A. Ashtari Esfahani *et al.*, “Determining the neutrino mass with cyclotron radiation emission spectroscopy? Project 8,” *J. Phys.* **G44** no. 5, (2017) 054004, [arXiv:1703.02037 \[physics.ins-det\]](#).
- [236] “<http://www.nu-fit.org/>.”
- [237] R. B. Patterson, “Prospects for Measurement of the Neutrino Mass Hierarchy,” *Ann. Rev. Nucl. Part. Sci.* **65** (2015) 177–192, [arXiv:1506.07917 \[hep-ex\]](#).
- [238] P. Langacker, “Neutrino Masses from the Top Down,” *Ann. Rev. Nucl. Part. Sci.* **62** (2012) 215–235, [arXiv:1112.5992 \[hep-ph\]](#).
- [239] K. N. Abazajian, “Telling three from four neutrinos with cosmology,” *Astropart. Phys.* **19** (2003) 303–312, [arXiv:astro-ph/0205238 \[astro-ph\]](#).
- [240] A. Slosar, T. Davis, D. Eisenstein, R. Hložek, M. Ishak-Boushaki, R. Mandelbaum, P. Marshall, J. Sakstein, and M. White, “Dark Energy and Modified Gravity,” *arXiv e-prints* (Mar, 2019) [arXiv:1903.12016](#), [arXiv:1903.12016 \[astro-ph.CO\]](#).
- [241] S. Mishra-Sharma, D. Alonso, and J. Dunkley, “Neutrino masses and beyond- Λ CDM cosmology with LSST and future CMB experiments,” *Phys. Rev.* **D 97** no. 12, (Jun, 2018) 123544, [arXiv:1803.07561 \[astro-ph.CO\]](#).
- [242] **DESI** Collaboration, A. Aghamousa, J. Aguilar, S. Ahlen, S. Alam, L. E. Allen, C. Allende Prieto, J. Annis, S. Bailey, C. Balland, O. Ballester, C. Baltay, L. Beaufore, C. Bebek, T. C. Beers, E. F. Bell, J. L. Bernal, R. Besuner, F. Beutler, C. Blake, H. Bleuler, M. Blomqvist, R. Blum, A. S. Bolton, C. Briceno, D. Brooks, J. R. Brownstein, E. Buckley-Geer, A. Burden, E. Burtin, N. G. Busca, R. N. Cahn, Y.-C. Cai, L. Cardiel-Sas, R. G. Carlberg, P.-H. Carton, R. Casas, F. J. Castander, J. L. Cervantes-Cota, T. M. Claybaugh, M. Close, C. T. Coker, S. Cole, J. Comparat, A. P. Cooper, M. C.

- Cousinou, M. Croce, J.-G. Cuby, D. P. Cunningham, T. M. Davis, K. S. Dawson, A. de la Macorra, J. De Vicente, T. Delubac, M. Derwent, A. Dey, G. Dhungana, Z. Ding, P. Doel, Y. T. Duan, A. Ealet, J. Edelstein, S. Eftekhazadeh, D. J. Eisenstein, A. Elliott, S. Escoffier, M. Evatt, P. Fagrelus, X. Fan, K. Fanning, A. Farahi, J. Farihi, G. Favole, Y. Feng, E. Fernandez, J. R. Findlay, D. P. Finkbeiner, M. J. Fitzpatrick, B. Flaughner, S. Flender, A. Font-Ribera, J. E. Forero-Romero, P. Fosalba, C. S. Frenk, M. Fumagalli, B. T. Gaensicke, G. Gallo, J. Garcia-Bellido, E. Gaztanaga, N. Pietro Gentile Fusillo, T. Gerard, I. Gershkovich, T. Giannantonio, D. Gillet, G. Gonzalez-de-Rivera, V. Gonzalez-Perez, S. Gott, O. Graur, G. Gutierrez, J. Guy, S. Habib, H. Heetderks, I. Heetderks, K. Heitmann, W. A. Hellwing, D. A. Herrera, S. Ho, S. Holland, K. Honscheid, E. Huff, T. A. Hutchinson, D. Huterer, H. S. Hwang, J. M. Illa Laguna, Y. Ishikawa, D. Jacobs, N. Jeffrey, P. Jelinsky, E. Jennings, L. Jiang, J. Jimenez, J. Johnson, R. Joyce, E. Jullo, S. Juneau, S. Kama, A. Karcher, S. Karkar, R. Kehoe, N. Kennamer, S. Kent, M. Kilbinger, A. G. Kim, D. Kirkby, T. Kisner, E. Kitanidis, J.-P. Kneib, S. Kuposov, E. Kovacs, K. Koyama, A. Kremin, R. Kron, L. Kronig, A. Kueter-Young, C. G. Lacey, R. Lafever, O. Lahav, A. Lambert, M. Lampton, M. Land riau, D. Lang, T. R. Lauer, J.-M. Le Goff, L. Le Guillou, A. Le Van Suu, J. H. Lee, S.-J. Lee, D. Leitner, M. Lesser, M. E. Levi, B. L’Huillier, B. Li, M. Liang, H. Lin, E. Linder, S. R. Loebman, Z. Lukić, J. Ma, N. MacCrann, C. Magneville, L. Makarem, M. Manera, C. J. Manser, R. Marshall, P. Martini, R. Massey, T. Matheson, J. McCauley, P. McDonald, I. D. McGreer, A. Meisner, N. Metcalfe, T. N. Miller, R. Miquel, J. Moustakas, A. Myers, M. Naik, J. A. Newman, R. C. Nichol, A. Nicola, L. Nicolati da Costa, J. Nie, G. Niz, P. Norberg, B. Nord, D. Norman, P. Nugent, T. O’Brien, M. Oh, K. A. G. Olsen, C. Padilla, H. Padmanabhan, N. Padmanabhan, N. Palanque-DeLabrouille, A. Palmese, D. Pappalardo, I. Pâris, C. Park, A. Patej, J. A. Peacock, H. V. Peiris, X. Peng, W. J. Percival, S. Perruchot, M. M. Pieri, R. Pogge, J. E. Pollack, C. Poppett, F. Prada, A. Prakash, R. G. Probst, D. Rabinowitz, A. Raichoor, C. H. Ree, A. Refregier, X. Regal, B. Reid, K. Reil, M. Rezaie, C. M. Rockosi, N. Roe, S. Ronayette, A. Roodman, A. J. Ross, N. P. Ross, G. Rossi, E. Roza, V. Ruhlmann-Kleider, E. S. Rykoff, C. Sabiu, L. Samushia, E. Sanchez, J. Sanchez, D. J. Schlegel, M. Schneider, M. Schubnell, A. Secroun, U. Seljak, H.-J. Seo, S. Serrano, A. Shafieloo, H. Shan, R. Sharples, M. J. Sholl, W. V. Shourt, J. H. Silber, D. R. Silva, M. M. Sirk, A. Slosar, A. Smith, G. F. Smoot, D. Som, Y.-S. Song, D. Sprayberry, R. Staten, A. Stefanik, G. Tarle, S. Sien Tie, J. L. Tinker, R. Tojeiro, F. Valdes, O. Valenzuela, M. Valluri, M. Vargas-Magana, L. Verde, A. R. Walker, J. Wang, Y. Wang, B. A. Weaver, C. Weaverdyck, R. H. Wechsler, D. H. Weinberg, M. White, Q. Yang, C. Yeche, T. Zhang, G.-B. Zhao, Y. Zheng, X. Zhou, Z. Zhou, Y. Zhu, H. Zou, and Y. Zu, “The DESI Experiment Part I: Science, Targeting, and Survey Design,” *arXiv e-prints* (Oct, 2016) arXiv:1611.00036, arXiv:1611.00036 [astro-ph.IM].
- [243] LSST Science Collaboration, P. A. Abell, J. Allison, S. F. Anderson, J. R. Andrew, J. R. P. Angel, L. Armus, D. Arnett, S. J. Asztalos, T. S. Axelrod, S. Bailey, D. R. Ballantyne, J. R. Bankert, W. A. Barkhouse, J. D. Barr, L. F. Barrientos, A. J. Barth, J. G. Bartlett, A. C. Becker, J. Becla, T. C. Beers, J. P. Bernstein, R. Biswas, M. R. Blanton, J. S. Bloom, J. J. Bochanski, P. Boeshaar, K. D. Borne, M. Bradac, W. N. Brandt, C. R. Bridge, M. E. Brown, R. J. Brunner, J. S. Bullock, A. J. Burgasser, J. H. Burge, D. L. Burke, P. A. Cargile, S. Chandrasekharan, G. Chartas, S. R. Chesley, Y.-H. Chu, D. Cinabro, M. W. Claire, C. F. Claver, D. Clowe, A. J. Connolly, K. H. Cook, J. Cooke, A. Cooray, K. R. Covey, C. S. Culliton, R. de Jong, W. H. de Vries, V. P. Debattista, F. Delgado, I. P. Dell’Antonio, S. Dhital, R. Di Stefano, M. Dickinson, B. Dilday, S. G. Djorgovski, G. Dobler, C. Donalek, G. Dubois-Felsmann, J. Durech, A. Eliasdottir, M. Eracleous, L. Eyer, E. E. Falco, X. Fan, C. D. Fassnacht, H. C. Ferguson, Y. R. Fernandez, B. D. Fields, D. Finkbeiner, E. E. Figuera, D. B. Fox, H. Francke, J. S. Frank, J. Frieman, S. Fromenteau, M. Furqan, G. Galaz, A. Gal-Yam, P. Garnavich, E. Gawiser, J. Geary, P. Gee, R. R. Gibson, K. Gilmore, E. A. Grace, R. F. Green, W. J. Gressler, C. J. Grillmair, S. Habib, J. S. Haggerty, M. Hamuy, A. W. Harris, S. L. Hawley, A. F. Heavens, L. Hebb, T. J. Henry, E. Hileman, E. J. Hilton, K. Hoadley, J. B. Holberg, M. J. Holman, S. B. Howell, L. Infante, Z. Ivezic, S. H. Jacoby, B. Jain, R. Jedicke, M. J.

- Jee, J. Garrett Jernigan, S. W. Jha, K. V. Johnston, R. L. Jones, M. Juric, M. Kaasalainen, Styliani, Kafka, S. M. Kahn, N. A. Kaib, J. Kalirai, J. Kantor, M. M. Kasliwal, C. R. Keeton, R. Kessler, Z. Knezevic, A. Kowalski, V. L. Krabbendam, K. S. Krughoff, S. Kulkarni, S. Kuhlman, M. Lacy, S. Lepine, M. Liang, A. Lien, P. Lira, K. S. Long, S. Lorenz, J. M. Lotz, R. H. Lupton, J. Lutz, L. M. Macri, A. A. Mahabal, R. Mandelbaum, P. Marshall, M. May, P. M. McGehee, B. T. Meadows, A. Meert, A. Milani, C. J. Miller, M. Miller, D. Mills, D. Minniti, D. Monet, A. S. Mukadam, E. Nakar, D. R. Neill, J. A. Newman, S. Nikolaev, M. Nordby, P. O'Connor, M. Oguri, J. Oliver, S. S. Olivier, J. K. Olsen, K. Olsen, E. W. Olszewski, H. Oluseyi, N. D. Padilla, A. Parker, J. Pepper, J. R. Peterson, C. Petry, P. A. Pinto, J. L. Pizagno, B. Popescu, A. Prsa, V. Radecka, M. J. Raddick, A. Rasmussen, A. Rau, J. Rho, J. E. Rhoads, G. T. Richards, S. T. Ridgway, B. E. Robertson, R. Roskar, A. Saha, A. Sarajedini, E. Scannapieco, T. Schalk, R. Schindler, S. Schmidt, S. Schmidt, D. P. Schneider, G. Schumacher, R. Scranton, J. Sebag, L. G. Seppala, O. Shemmer, J. D. Simon, M. Sivertz, H. A. Smith, J. Allyn Smith, N. Smith, A. H. Spitz, A. Stanford, K. G. Stassun, J. Strader, M. A. Strauss, C. W. Stubbs, D. W. Sweeney, A. Szalay, P. Szkody, M. Takada, P. Thorman, D. E. Trilling, V. Trimble, A. Tyson, R. Van Berg, D. Vand en Berk, J. VanderPlas, L. Verde, B. Vrsnak, L. M. Walkowicz, B. D. Wand elt, S. Wang, Y. Wang, M. Warner, R. H. Wechsler, A. A. West, O. Wiecha, B. F. Williams, B. Willman, D. Wittman, S. C. Wolff, W. M. Wood-Vasey, P. Wozniak, P. Young, A. Zentner, and H. Zhan, "LSST Science Book, Version 2.0," *arXiv e-prints* (Dec, 2009) arXiv:0912.0201, arXiv:0912.0201 [astro-ph.IM].
- [244] The LSST Dark Energy Science Collaboration, R. Mandelbaum, T. Eifler, R. Hložek, T. Collett, E. Gawiser, D. Scolnic, D. Alonso, H. Awan, R. Biswas, J. Blazek, P. Burchat, N. E. Chisari, I. Dell'Antonio, S. Digel, J. Frieman, D. A. Goldstein, I. Hook, Ž. Ivezić, S. M. Kahn, S. Kamath, D. Kirkby, T. Kitching, E. Krause, P.-F. Leget, P. J. Marshall, J. Meyers, H. Miyatake, J. A. Newman, R. Nichol, E. Rykoff, F. J. Sanchez, A. Slosar, M. Sullivan, and M. A. Troxel, "The LSST Dark Energy Science Collaboration (DESC) Science Requirements Document," *arXiv e-prints* (Sep, 2018) arXiv:1809.01669, arXiv:1809.01669 [astro-ph.CO].
- [245] A. Albrecht, G. Bernstein, R. Cahn, W. L. Freedman, J. Hewitt, W. Hu, J. Huth, M. Kamionkowski, E. W. Kolb, L. Knox, J. C. Mather, S. Staggs, and N. B. Suntzeff, "Report of the Dark Energy Task Force," *arXiv e-prints* (Sep, 2006) astro-ph/0609591, arXiv:astro-ph/0609591 [astro-ph].
- [246] S. Hannestad, "Neutrino Masses and the Dark Energy Equation of State:Relaxing the Cosmological Neutrino Mass Bound," *Phys. Rev. L* **95** no. 22, (Nov, 2005) 221301, arXiv:astro-ph/0505551 [astro-ph].
- [247] S. Bhattacharya and A. Kosowsky, "Dark Energy Constraints from Galaxy Cluster Peculiar Velocities," *Phys. Rev. D* **77** (2008) 083004, arXiv:0712.0034 [astro-ph].
- [248] A. Kosowsky and S. Bhattacharya, "A Future Test of Gravitation Using Galaxy Cluster Velocities," *Phys. Rev. D* **80** (2009) 062003, arXiv:0907.4202 [astro-ph.CO].
- [249] E.-M. Mueller, F. de Bernardis, R. Bean, and M. D. Niemack, "Constraints on gravity and dark energy from the pairwise kinematic Sunyaev-Zeldovich effect," *Astrophys. J.* **808** no. 1, (2015) 47, arXiv:1408.6248 [astro-ph.CO].
- [250] E.-M. Mueller, F. de Bernardis, R. Bean, and M. D. Niemack, "Constraints on massive neutrinos from the pairwise kinematic Sunyaev-Zel'dovich effect," *Phys. Rev. D* **92** no. 6, (2015) 063501, arXiv:1412.0592 [astro-ph.CO].
- [251] D. Alonso, J. C. Hill, R. Hložek, and D. N. Spergel, "Measurement of the thermal Sunyaev-Zel'dovich effect around cosmic voids," *Phys. Rev. D* **97** no. 6, (Mar., 2018) 063514, arXiv:1709.01489.

- [252] S. Flender, D. Nagai, and M. McDonald, “Constraints on the optical depth of galaxy groups and clusters,” *Astrophys. J.* **837** no. 2, (2017) 124, [arXiv:1610.08029 \[astro-ph.CO\]](#).
- [253] N. S. Sugiyama, T. Okumura, and D. N. Spergel, “Will kinematic Sunyaev-Zel’dovich measurements enhance the science return from galaxy redshift surveys?,” *JCAP* **1** (Jan., 2017) 057, [arXiv:1606.06367](#).
- [254] M. P. van Daalen, J. Schaye, C. M. Booth, and C. D. Vecchia, “The effects of galaxy formation on the matter power spectrum: A challenge for precision cosmology,” *Mon. Not. Roy. Astron. Soc.* **415** (2011) 3649–3665, [arXiv:1104.1174 \[astro-ph.CO\]](#).
- [255] I. Mohammed, D. Martizzi, R. Teyssier, and A. Amara, “Baryonic effects on weak-lensing two-point statistics and its cosmological implications,” [arXiv:1410.6826 \[astro-ph.CO\]](#).
- [256] **POLARBEAR** Collaboration, P. A. R. Ade *et al.*, “POLARBEAR Constraints on Cosmic Birefringence and Primordial Magnetic Fields,” *Phys. Rev.* **D92** (2015) 123509, [arXiv:1509.02461 \[astro-ph.CO\]](#).
- [257] **Keck Array, BICEP2** Collaboration, P. A. R. Ade *et al.*, “BICEP2 / Keck Array IX: New bounds on anisotropies of CMB polarization rotation and implications for axionlike particles and primordial magnetic fields,” *Phys. Rev.* **D96** no. 10, (2017) 102003, [arXiv:1705.02523 \[astro-ph.CO\]](#).
- [258] S. M. Carroll, “Quintessence and the rest of the world,” *Phys. Rev. Lett.* **81** (1998) 3067–3070, [arXiv:astro-ph/9806099 \[astro-ph\]](#).
- [259] M. Kamionkowski, “How to De-Rotate the Cosmic Microwave Background Polarization,” *Phys. Rev. Lett.* **102** (2009) 111302, [arXiv:0810.1286 \[astro-ph\]](#).
- [260] V. Gluscevic, M. Kamionkowski, and A. Cooray, “De-Rotation of the Cosmic Microwave Background Polarization: Full-Sky Formalism,” *Phys. Rev.* **D80** (2009) 023510, [arXiv:0905.1687 \[astro-ph.CO\]](#).
- [261] **Planck** Collaboration, Planck Collaboration Int. XLIX, “Planck intermediate results. XLIX. Parity-violation constraints from polarization data,” *Astron. Astrophys.* **596** (2016) A110, [arXiv:1605.08633 \[astro-ph.CO\]](#).
- [262] V. Gluscevic, D. Hanson, M. Kamionkowski, and C. M. Hirata, “First CMB Constraints on Direction-Dependent Cosmological Birefringence from WMAP-7,” *Phys. Rev.* **D86** (2012) 103529, [arXiv:1206.5546 \[astro-ph.CO\]](#).
- [263] M. Pospelov, A. Ritz, C. Skordis, A. Ritz, and C. Skordis, “Pseudoscalar perturbations and polarization of the cosmic microwave background,” *Phys. Rev. Lett.* **103** (2009) 051302, [arXiv:0808.0673 \[astro-ph\]](#).
- [264] D. Contreras, P. Boubel, and D. Scott, “Constraints on direction-dependent cosmic birefringence from Planck polarization data,” *JCAP* **12** (Dec., 2017) 046, [arXiv:1705.06387](#).
- [265] T. Emken and C. Kouvaris, “How blind are underground and surface detectors to strongly interacting dark matter?,” *Phys. Rev. D* **97** no. 11, (June, 2018) 115047, [arXiv:1802.04764 \[hep-ph\]](#).
- [266] B. J. Kavanagh, “Earth scattering of superheavy dark matter: Updated constraints from detectors old and new,” *Phys. Rev. D* **97** no. 12, (June, 2018) 123013, [arXiv:1712.04901 \[hep-ph\]](#).

- [267] P. Cushman, C. Galbiati, D. N. McKinsey, H. Robertson, T. M. P. Tait, D. Bauer, A. Borgland, B. Cabrera, F. Calaprice, J. Cooley, T. Empl, R. Essig, E. Figueroa-Feliciano, R. Gaitskell, S. Golwala, J. Hall, R. Hill, A. Hime, E. Hoppe, L. Hsu, E. Hungerford, R. Jacobsen, M. Kelsey, R. F. Lang, W. H. Lippincott, B. Loer, S. Luitz, V. Mandic, J. Mardon, J. Maricic, R. Maruyama, R. Mahapatra, H. Nelson, J. Orrell, K. Palladino, E. Pantic, R. Partridge, A. Ryd, T. Saab, B. Sadoulet, R. Schnee, W. Shepherd, A. Sonnenschein, P. Sorensen, M. Szydagis, T. Volansky, M. Witherell, D. Wright, and K. Zurek, “Snowmass CF1 Summary: WIMP Dark Matter Direct Detection,” *ArXiv e-prints* (Oct., 2013), [arXiv:1310.8327 \[hep-ex\]](#).
- [268] D. Hooper and S. D. McDermott, “Robust Constraints and Novel Gamma-Ray Signatures of Dark Matter That Interacts Strongly With Nucleons,” *Phys. Rev.* **D97** no. 11, (2018) 115006, [arXiv:1802.03025 \[hep-ph\]](#).
- [269] X. Chen, S. Hannestad, and R. J. Scherrer, “Cosmic microwave background and large scale structure limits on the interaction between dark matter and baryons,” *ArXiv Astrophysics e-prints* (Feb., 2002), [astro-ph/0202496](#).
- [270] K. Sigurdson, M. Doran, A. Kurylov, R. R. Caldwell, and M. Kamionkowski, “Dark-matter electric and magnetic dipole moments,” *Phys. Rev.* **D 70** no. 8, (Oct., 2004) 083501, [astro-ph/0406355](#).
- [271] C. Dvorkin, K. Blum, and M. Kamionkowski, “Constraining Dark Matter-Baryon Scattering with Linear Cosmology,” *Phys. Rev. D* **89** no. 2, (2014) 023519, [arXiv:1311.2937 \[astro-ph.CO\]](#).
- [272] V. Gluscevic and K. K. Boddy, “Constraints on Scattering of keV-TeV Dark Matter with Protons in the Early Universe,” *Physical Review Letters* **121** no. 8, (Aug., 2018) 081301, [arXiv:1712.07133](#).
- [273] K. K. Boddy and V. Gluscevic, “First cosmological constraint on the effective theory of dark matter-proton interactions,” *Phys. Rev.* **D 98** no. 8, (Oct, 2018) 083510, [arXiv:1801.08609 \[astro-ph.CO\]](#).
- [274] W. L. Xu, C. Dvorkin, and A. Chael, “Probing sub-GeV dark matter-baryon scattering with cosmological observables,” *Phys. Rev.* **D 97** no. 10, (May, 2018) 103530, [arXiv:1802.06788](#).
- [275] K. K. Boddy and V. Gluscevic, “First Cosmological Constraint on the Effective Theory of Dark Matter-Proton Interactions,” *ArXiv e-prints* (Jan., 2018), [arXiv:1801.08609](#).
- [276] J. Fan, M. Reece, and L.-T. Wang, “Non-relativistic effective theory of dark matter direct detection,” *JCAP* **1011** (2010) 042, [arXiv:1008.1591 \[hep-ph\]](#).
- [277] A. L. Fitzpatrick, W. Haxton, E. Katz, N. Lubbers, and Y. Xu, “The Effective Field Theory of Dark Matter Direct Detection,” *JCAP* **1302** (2013) 004, [arXiv:1203.3542 \[hep-ph\]](#).
- [278] N. Anand, A. L. Fitzpatrick, and W. C. Haxton, “Weakly interacting massive particle-nucleus elastic scattering response,” *Phys. Rev.* **C89** no. 6, (2014) 065501, [arXiv:1308.6288 \[hep-ph\]](#).
- [279] J. B. Dent, L. M. Krauss, J. L. Newstead, and S. Sabharwal, “General analysis of direct dark matter detection: From microphysics to observational signatures,” *Phys. Rev.* **D92** no. 6, (2015) 063515, [arXiv:1505.03117 \[hep-ph\]](#).
- [280] K. K. Boddy, V. Gluscevic, V. Poulin, E. D. Kovetz, M. Kamionkowski, and R. Barkana, “Critical assessment of CMB limits on dark matter-baryon scattering: New treatment of the relative bulk velocity,” *Phys. Rev.* **D 98** no. 12, (Dec, 2018) 123506, [arXiv:1808.00001 \[astro-ph.CO\]](#).
- [281] T. R. Slatyer and C.-L. Wu, “Early-Universe constraints on dark matter-baryon scattering and their implications for a global 21 cm signal,” *Phys. Rev.* **D 98** no. 2, (July, 2018) 023013, [arXiv:1803.09734](#).

- [282] R. de Putter, O. Doré, J. Gleyzes, D. Green, and J. Meyers, “Dark Matter Interactions, Helium, and the CMB,” *ArXiv e-prints* (May, 2018) , [arXiv:1805.11616](#).
- [283] J. D. Bowman, A. E. E. Rogers, R. A. Monsalve, T. J. Mozdzen, and N. Mahesh, “An absorption profile centred at 78 megahertz in the sky-averaged spectrum,” *Nature* **555** (Mar., 2018) 67–70, [arXiv:1810.05912](#).
- [284] R. Barkana, “Possible interaction between baryons and dark-matter particles revealed by the first stars,” *Nature* **555** (Mar., 2018) 71–74, [arXiv:1803.06698](#).
- [285] E. D. Kovetz, V. Poulin, V. Gluscevic, K. K. Boddy, R. Barkana, and M. Kamionkowski, “Tighter limits on dark matter explanations of the anomalous EDGES 21 cm signal,” *Phys. Rev. D* **98** no. 10, (Nov, 2018) 103529, [arXiv:1807.11482](#) [[astro-ph.CO](#)].
- [286] Z. Li, V. Gluscevic, K. K. Boddy, and M. S. Madhavacheril, “Disentangling dark physics with cosmic microwave background experiments,” *Phys. Rev. D* **98** no. 12, (Dec, 2018) 123524, [arXiv:1806.10165](#) [[astro-ph.CO](#)].
- [287] R. J. Thornton, P. A. R. Ade, S. Aiola, F. E. Angilè, M. Amiri, J. A. Beall, D. T. Becker, H. M. Cho, S. K. Choi, P. Corlies, K. P. Coughlin, R. Datta, M. J. Devlin, S. R. Dicker, R. Dünner, J. W. Fowler, A. E. Fox, P. A. Gallardo, J. Gao, E. Grace, M. Halpern, M. Hasselfield, S. W. Henderson, G. C. Hilton, A. D. Hincks, S. P. Ho, J. Hubmayr, K. D. Irwin, J. Klein, B. Koopman, D. Li, T. Louis, M. Lungu, L. Maurin, J. McMahon, C. D. Munson, S. Naess, F. Nati, L. Newburgh, J. Nibarger, M. D. Niemack, P. Niraula, M. R. Nolta, L. A. Page, C. G. Pappas, A. Schillaci, B. L. Schmitt, N. Sehgal, J. L. Sievers, S. M. Simon, S. T. Staggs, C. Tucker, M. Uehara, J. van Lanen, J. T. Ward, and E. J. Wollack, “The Atacama Cosmology Telescope: The Polarization-sensitive ACTPol Instrument,” *The Astrophysical Journal Supplement Series* **227** no. 2, (Dec, 2016) 21, [arXiv:1605.06569](#) [[astro-ph.IM](#)].
- [288] E. O. Nadler, V. Gluscevic, K. K. Boddy, and R. H. Wechsler, “Constraints on Dark Matter Microphysics from the Milky Way Satellite Population,” *arXiv e-prints* (Apr, 2019) [arXiv:1904.10000](#), [arXiv:1904.10000](#) [[astro-ph.CO](#)].
- [289] A. Drlica-Wagner, Y.-Y. Mao, S. Adhikari, R. Armstrong, A. Banerjee, N. Banik, K. Bechtol, S. Bird, K. K. Boddy, A. Bonaca, J. Bovy, M. R. Buckley, E. Bulbul, C. Chang, G. Chapline, J. Cohen-Tanugi, A. Cuoco, F.-Y. Cyr-Racine, W. A. Dawson, A. Díaz Rivero, C. Dvorkin, D. Erkal, C. D. Fassnacht, J. García-Bellido, M. Giannotti, V. Gluscevic, N. Golovich, D. Hendel, Y. D. Hezaveh, S. Horiuchi, M. J. Jee, M. Kaplinghat, C. R. Keeton, S. E. Kogosov, T. S. Li, R. Mandelbaum, S. D. McDermott, M. McNanna, M. Medford, M. Meyer, M. Marc, S. Murgia, E. O. Nadler, L. Necib, E. Nuss, A. B. Pace, A. H. G. Peter, D. A. Polin, C. Prescod-Weinstein, J. I. Read, R. Rosenfeld, N. Shipp, J. D. Simon, T. R. Slatyer, O. Straniero, L. E. Strigari, E. Tollerud, J. A. Tyson, M.-Y. Wang, R. H. Wechsler, D. Wittman, H.-B. Yu, G. Zaharijas, Y. Ali-Haïmoud, J. Annis, S. Birrer, R. Biswas, J. Blazek, A. M. Brooks, E. Buckley-Geer, R. Caputo, E. Charles, S. Digel, S. Dodelson, B. Flaugher, J. Frieman, E. Gawiser, A. P. Hearin, R. Hložek, B. Jain, T. E. Jeltema, S. M. Koushiappas, M. Lisanti, M. LoVerde, S. Mishra-Sharma, J. A. Newman, B. Nord, E. Nourbakhsh, S. Ritz, B. E. Robertson, M. A. Sánchez-Conde, A. Slosar, T. M. P. Tait, A. Verma, R. Vilalta, C. W. Walter, B. Yanny, and A. R. Zentner, “Probing the Fundamental Nature of Dark Matter with the Large Synoptic Survey Telescope,” *arXiv e-prints* (Feb, 2019) [arXiv:1902.01055](#), [arXiv:1902.01055](#) [[astro-ph.CO](#)].
- [290] L. G. van den Aarssen, T. Bringmann, and C. Pfrommer, “Is dark matter with long-range interactions a solution to all small-scale problems of Λ CDM cosmology?,” *Phys. Rev. Lett.* **109** (2012) 231301, [arXiv:1205.5809](#) [[astro-ph.CO](#)].

- [291] F.-Y. Cyr-Racine, R. de Putter, A. Raccanelli, and K. Sigurdson, “Constraints on Large-Scale Dark Acoustic Oscillations from Cosmology,” *Phys. Rev. D* **89** no. 6, (2014) 063517, [arXiv:1310.3278](#) [[astro-ph.CO](#)].
- [292] J. S. Bullock and M. Boylan-Kolchin, “Small-Scale Challenges to the Λ CDM Paradigm,” *Ann. Rev. Astron. Astrophys.* **55** (2017) 343–387, [arXiv:1707.04256](#) [[astro-ph.CO](#)].
- [293] R. Krall, F.-Y. Cyr-Racine, and C. Dvorkin, “Wandering in the Lyman-alpha Forest: A Study of Dark Matter-Dark Radiation Interactions,” *JCAP* **1709** no. 09, (2017) 003, [arXiv:1705.08894](#) [[astro-ph.CO](#)].
- [294] F.-Y. Cyr-Racine, K. Sigurdson, J. Zavala, T. Bringmann, M. Vogelsberger, and C. Pfrommer, “ETHOS – an effective theory of structure formation: From dark particle physics to the matter distribution of the Universe,” *Phys. Rev.* **D93** no. 12, (2016) 123527, [arXiv:1512.05344](#) [[astro-ph.CO](#)].
- [295] S. Tulin, H.-B. Yu, and K. M. Zurek, “Resonant Dark Forces and Small Scale Structure,” *Phys. Rev. Lett.* **110** (2013) 111301, [arXiv:1210.0900](#) [[hep-ph](#)].
- [296] S. Tulin, H.-B. Yu, and K. M. Zurek, “Beyond Collisionless Dark Matter: Particle Physics Dynamics for Dark Matter Halo Structure,” *Phys. Rev. D* **87** (2013) 115007, [arXiv:1302.3898](#) [[hep-ph](#)].
- [297] M. Kaplinghat, S. Tulin, and H.-B. Yu, “Dark Matter Halos as Particle Colliders: A Unified Solution to Small-Scale Structure Puzzles from Dwarfs to Clusters,” [arXiv:1508.03339](#) [[astro-ph.CO](#)].
- [298] M. Vogelsberger, J. Zavala, F.-Y. Cyr-Racine, C. Pfrommer, T. Bringmann, and K. Sigurdson, “ETHOS - An Effective Theory of Structure Formation: Dark matter physics as a possible explanation of the small-scale CDM problems,” *Mon. Not. Roy. Astron. Soc.* **460** (2016) 1399, [arXiv:1512.05349](#) [[astro-ph.CO](#)].
- [299] A. Schneider, “Astrophysical constraints on resonantly produced sterile neutrino dark matter,” *JCAP* **1604** no. 04, (2016) 059, [arXiv:1601.07553](#) [[astro-ph.CO](#)].
- [300] R. Hložek, D. J. E. Marsh, and D. Grin, “Using the Full Power of the Cosmic Microwave Background to Probe Axion Dark Matter,” *Mon. Not. Roy. Astron. Soc.* **476** no. 3, (2018) 3063–3085, [arXiv:1708.05681](#) [[astro-ph.CO](#)].
- [301] R. Hložek, D. Grin, D. J. Marsh, and P. G. Ferreira, “A search for ultralight axions using precision cosmological data,” *Phys. Rev.* **D91** no. 10, (2015) 103512, [arXiv:1410.2896](#) [[astro-ph.CO](#)].
- [302] R. Hložek, D. J. E. Marsh, D. Grin, R. Allison, J. Dunkley, and E. Calabrese, “Future CMB tests of dark matter: ultra-light axions and massive neutrinos,” [arXiv:1607.08208](#) [[astro-ph.CO](#)].
- [303] I. De Martino, T. Broadhurst, S. H. Henry Tye, T. Chiueh, H.-Y. Schive, and R. Lazkoz, “Recognizing Axionic Dark Matter by Compton and de Broglie Scale Modulation of Pulsar Timing,” *Phys. Rev. Lett.* **119** no. 22, (2017) 221103, [arXiv:1705.04367](#) [[astro-ph.CO](#)].
- [304] N. K. Porayko *et al.*, “Parkes Pulsar Timing Array constraints on ultralight scalar-field dark matter,” [arXiv:1810.03227](#) [[astro-ph.CO](#)].
- [305] V. Iršič, M. Viel, M. G. Haehnelt, J. S. Bolton, and G. D. Becker, “First constraints on fuzzy dark matter from Lyman- α forest data and hydrodynamical simulations,” *Phys. Rev. Lett.* **119** no. 3, (2017) 031302, [arXiv:1703.04683](#) [[astro-ph.CO](#)].

- [306] T. Kobayashi, R. Murgia, A. De Simone, V. Iršič, and M. Viel, “Lyman- α constraints on ultralight scalar dark matter: Implications for the early and late universe,” *Phys. Rev.* **D96** no. 12, (2017) 123514, [arXiv:1708.00015](#) [[astro-ph.CO](#)].
- [307] E. Armengaud, N. Palanque-Delabrouille, C. Yèche, D. J. E. Marsh, and J. Baur, “Constraining the mass of light bosonic dark matter using SDSS Lyman- α forest,” *Mon. Not. Roy. Astron. Soc.* **471** no. 4, (2017) 4606–4614, [arXiv:1703.09126](#) [[astro-ph.CO](#)].
- [308] K.-H. Leong, H.-Y. Schive, U.-H. Zhang, and T. Chiueh, “Testing extreme-axion wave dark matter using the BOSS Lyman-Alpha forest data,” [arXiv:1810.05930](#) [[astro-ph.CO](#)].
- [309] M. Nori, R. Murgia, V. Iršič, M. Baldi, and M. Viel, “Lyman- α forest and non-linear structure characterization in Fuzzy Dark Matter cosmologies,” [arXiv:1809.09619](#) [[astro-ph.CO](#)].
- [310] A. Arvanitaki and S. Dubovsky, “Exploring the String Axiverse with Precision Black Hole Physics,” *Phys. Rev.* **D83** (2011) 044026, [arXiv:1004.3558](#) [[hep-th](#)].
- [311] M. J. Stott and D. J. E. Marsh, “Black hole spin constraints on the mass spectrum and number of axionlike fields,” *Phys. Rev.* **D98** no. 8, (2018) 083006, [arXiv:1805.02016](#) [[hep-ph](#)].
- [312] N. Kitajima, J. Soda, and Y. Urakawa, “Gravitational wave forest from string axiverse,” *JCAP* **1810** no. 10, (2018) 008, [arXiv:1807.07037](#) [[astro-ph.CO](#)].
- [313] D. Baumann, H. S. Chia, and R. A. Porto, “Probing Ultralight Bosons with Binary Black Holes,” *Phys. Rev.* **D99** no. 4, (2019) 044001, [arXiv:1804.03208](#) [[gr-qc](#)].
- [314] D. Grin, M. A. Amin, V. Gluscevic, R. Hložek, D. J. E. Marsh, V. Poulin, C. Prescod-Weinstein, and T. L. Smith, “Gravitational probes of ultra-light axions,” [arXiv:1904.09003](#) [[astro-ph.CO](#)].
- [315] D. J. E. Marsh, “WarmAndFuzzy: the halo model beyond CDM,” [arXiv:1605.05973](#) [[astro-ph.CO](#)].
- [316] D. J. E. Marsh and J. Silk, “A Model For Halo Formation With Axion Mixed Dark Matter,” *Mon. Not. Roy. Astron. Soc.* **437** no. 3, (2014) 2652–2663, [arXiv:1307.1705](#) [[astro-ph.CO](#)].
- [317] H.-Y. Schive, M.-H. Liao, T.-P. Woo, S.-K. Wong, T. Chiueh, T. Broadhurst, and W. Y. P. Hwang, “Understanding the Core-Halo Relation of Quantum Wave Dark Matter from 3D Simulations,” *Phys. Rev. Lett.* **113** no. 26, (2014) 261302, [arXiv:1407.7762](#) [[astro-ph.GA](#)].
- [318] H.-Y. Schive, T. Chiueh, and T. Broadhurst, “Cosmic Structure as the Quantum Interference of a Coherent Dark Wave,” *Nature Phys.* **10** (2014) 496–499, [arXiv:1406.6586](#) [[astro-ph.GA](#)].
- [319] H.-Y. Schive, T. Chiueh, T. Broadhurst, and K.-W. Huang, “Contrasting Galaxy Formation from Quantum Wave Dark Matter, ψ DM, with Λ CDM, using Planck and Hubble Data,” *Astrophys. J.* **818** no. 1, (2016) 89, [arXiv:1508.04621](#) [[astro-ph.GA](#)].
- [320] V. H. Robles, V. Lora, T. Matos, and F. J. Sánchez-Salcedo, “Evolution of a dwarf satellite galaxy embedded in a scalar field dark matter halo,” *Astrophys. J.* **810** no. 2, (2015) 99, [arXiv:1404.3424](#) [[astro-ph.GA](#)].
- [321] J. Veltmaat and J. C. Niemeyer, “Cosmological particle-in-cell simulations with ultralight axion dark matter,” *Phys. Rev.* **D94** no. 12, (2016) 123523, [arXiv:1608.00802](#) [[astro-ph.CO](#)].
- [322] X. Du, C. Behrens, J. C. Niemeyer, and B. Schwabe, “Core-halo mass relation of ultralight axion dark matter from merger history,” *Phys. Rev.* **D95** no. 4, (2017) 043519, [arXiv:1609.09414](#) [[astro-ph.GA](#)].

- [323] J. Zhang, Y.-L. S. Tsai, J.-L. Kuo, K. Cheung, and M.-C. Chu, “Ultralight Axion Dark Matter and Its Impact on Dark Halo Structure in N -body Simulations,” *Astrophys. J.* **853** no. 1, (2018) 51, [arXiv:1611.00892 \[astro-ph.CO\]](#).
- [324] S.-R. Chen, H.-Y. Schive, and T. Chiueh, “Jeans Analysis for Dwarf Spheroidal Galaxies in Wave Dark Matter,” *Mon. Not. Roy. Astron. Soc.* **468** no. 2, (2017) 1338–1348, [arXiv:1606.09030 \[astro-ph.GA\]](#).
- [325] H.-Y. Schive and T. Chiueh, “Halo abundance and assembly history with extreme-axion wave dark matter at $z \geq 4$,” *Mon. Not. Roy. Astron. Soc.* **473** no. 1, (2018) L36–L40, [arXiv:1706.03723 \[astro-ph.CO\]](#).
- [326] F. X. Linares Cedeño, A. X. González-Morales, and L. Arturo Ureña López, “Towards a calculation of the halo mass function of a scalar field dark matter,” *J. Phys. Conf. Ser.* **1030** no. 1, (2018) 012006.
- [327] X. Du, B. Schwabe, J. C. Niemeyer, and D. Büurger, “Tidal disruption of fuzzy dark matter subhalo cores,” *Phys. Rev.* **D97** no. 6, (2018) 063507, [arXiv:1801.04864 \[astro-ph.GA\]](#).
- [328] I. De Martino, T. Broadhurst, S. H. H. Tye, T. Chiueh, and H.-Y. Schive, “Dynamical Evidence of a Solitonic Core of $10^9 M_\odot$ in the Milky Way,” [arXiv:1807.08153 \[astro-ph.GA\]](#).
- [329] S.-C. Lin, H.-Y. Schive, S.-K. Wong, and T. Chiueh, “Self-consistent construction of virialized wave dark matter halos,” *Phys. Rev.* **D97** no. 10, (2018) 103523, [arXiv:1801.02320 \[astro-ph.CO\]](#).
- [330] V. H. Robles, J. S. Bullock, and M. Boylan-Kolchin, “Scalar Field Dark Matter: Helping or Hurting Small-Scale Problems in Cosmology?,” [arXiv:1807.06018 \[astro-ph.CO\]](#).
- [331] J. Veltmaat, J. C. Niemeyer, and B. Schwabe, “Formation and structure of ultralight bosonic dark matter halos,” *Phys. Rev.* **D98** no. 4, (2018) 043509, [arXiv:1804.09647 \[astro-ph.CO\]](#).
- [332] M. Nori and M. Baldi, “AX-GADGET: a new code for cosmological simulations of Fuzzy Dark Matter and Axion models,” *Mon. Not. Roy. Astron. Soc.* **478** no. 3, (2018) 3935–3951, [arXiv:1801.08144 \[astro-ph.CO\]](#).
- [333] M. Axenides, R. H. Brandenberger, and M. S. Turner, “Development of Axion Perturbations in an Axion Dominated Universe,” *Phys. Lett.* **B126** (1983) 178.
- [334] P. Fox, A. Pierce, and S. D. Thomas, “Probing a QCD string axion with precision cosmological measurements,” [arXiv:hep-th/0409059 \[hep-th\]](#).
- [335] M. P. Hertzberg, M. Tegmark, and F. Wilczek, “Axion Cosmology and the Energy Scale of Inflation,” *Phys.Rev.* **D78** (2008) 083507, [arXiv:0807.1726 \[astro-ph\]](#).
- [336] P. Svrcek and E. Witten, “Axions In String Theory,” *JHEP* **06** (2006) 051, [arXiv:hep-th/0605206 \[hep-th\]](#).
- [337] J. Preskill, M. B. Wise, and F. Wilczek, “Cosmology of the Invisible Axion,” *Phys. Lett.* **B120** (1983) 127–132.
- [338] L. F. Abbott and P. Sikivie, “A Cosmological Bound on the Invisible Axion,” *Phys. Lett.* **B120** (1983) 133–136.
- [339] M. Dine and W. Fischler, “The Not So Harmless Axion,” *Phys. Lett.* **B120** (1983) 137–141.
- [340] M. S. Turner, “Coherent Scalar Field Oscillations in an Expanding Universe,” *Phys. Rev.* **D28** (1983) 1243.

- [341] P. J. Steinhardt and M. S. Turner, “Saving the Invisible Axion,” *Phys. Lett.* **B129** (1983) 51.
- [342] D. J. E. Marsh and P. G. Ferreira, “Ultra-Light Scalar Fields and the Growth of Structure in the Universe,” *Phys. Rev.* **D82** (2010) 103528, [arXiv:1009.3501 \[hep-ph\]](#).
- [343] **ADMX** Collaboration, S. J. Asztalos *et al.*, “A SQUID-based microwave cavity search for dark-matter axions,” *Phys. Rev. Lett.* **104** (2010) 041301, [arXiv:0910.5914 \[astro-ph.CO\]](#).
- [344] D. Budker, P. W. Graham, M. Ledbetter, S. Rajendran, and A. Sushkov, “Proposal for a Cosmic Axion Spin Precession Experiment (CASPER),” *Phys. Rev.* **X4** no. 2, (2014) 021030, [arXiv:1306.6089 \[hep-ph\]](#).
- [345] D. J. E. Marsh, D. Grin, R. Hlozek, and P. G. Ferreira, “Axiverse cosmology and the energy scale of inflation,” *Phys. Rev.* **D87** (2013) 121701, [arXiv:1303.3008 \[astro-ph.CO\]](#).
- [346] D. J. E. Marsh, D. Grin, R. Hlozek, and P. G. Ferreira, “Tensor Interpretation of BICEP2 Results Severely Constrains Axion Dark Matter,” *Phys. Rev. Lett.* **113** no. 1, (2014) 011801, [arXiv:1403.4216 \[astro-ph.CO\]](#).
- [347] **LSST Collaboration** Collaboration, Z. Ivezić, J. Tyson, R. Allsman, J. Andrew, and R. Angel, “LSST: from Science Drivers to Reference Design and Anticipated Data Products,” [arXiv:0805.2366 \[astro-ph\]](#).
- [348] DESI Collaboration, A. Aghamousa, J. Aguilar, S. Ahlen, S. Alam, L. E. Allen, C. Allende Prieto, J. Annis, S. Bailey, C. Balland, and *et al.*, “The DESI Experiment Part I: Science, Targeting, and Survey Design,” *ArXiv e-prints* (Oct., 2016) , [arXiv:1611.00036 \[astro-ph.IM\]](#).
- [349] A. Merloni, P. Predehl, W. Becker, H. Böhringer, T. Boller, H. Brunner, M. Brusa, K. Dennerl, M. Freyberg, P. Friedrich, A. Georgakakis, F. Haberl, G. Hasinger, N. Meidinger, J. Mohr, K. Nandra, A. Rau, T. H. Reiprich, J. Robrade, M. Salvato, A. Santangelo, M. Sasaki, A. Schwobe, J. Wilms, and *t. German eROSITA Consortium*, “eROSITA Science Book: Mapping the Structure of the Energetic Universe,” *ArXiv e-prints* (Sept., 2012) , [arXiv:1209.3114 \[astro-ph.HE\]](#).
- [350] K. M. Smith, O. Zahn, and O. Dore, “Detection of Gravitational Lensing in the Cosmic Microwave Background,” *Phys. Rev.* **D76** (2007) 043510, [arXiv:0705.3980 \[astro-ph\]](#).
- [351] C. M. Hirata, S. Ho, N. Padmanabhan, U. Seljak, and N. A. Bahcall, “Correlation of CMB with large-scale structure: II. Weak lensing,” *Phys. Rev.* **D78** (2008) 043520, [arXiv:0801.0644 \[astro-ph\]](#).
- [352] **ACT** Collaboration, R. Allison *et al.*, “The Atacama Cosmology Telescope: measuring radio galaxy bias through cross-correlation with lensing,” *Mon. Not. Roy. Astron. Soc.* **451** no. 1, (2015) 849–858, [arXiv:1502.06456 \[astro-ph.CO\]](#).
- [353] J. C. Hill and D. N. Spergel, “Detection of thermal SZ-CMB lensing cross-correlation in Planck nominal mission data,” *JCAP* **1402** (2014) 030, [arXiv:1312.4525 \[astro-ph.CO\]](#).
- [354] G. Holder, M. Viero, O. Zahn, K. Aird, B. Benson, *et al.*, “A Cosmic Microwave Background Lensing Mass Map and Its Correlation with the Cosmic Infrared Background,” *Astrophys. J.* **771** (2013) L16, [arXiv:1303.5048 \[astro-ph.CO\]](#).
- [355] **Planck** Collaboration, Planck Collaboration XVIII, “Planck 2013 results. XVIII. The gravitational lensing-infrared background correlation,” *Astron. Astrophys.* **571** (2014) A18, [arXiv:1303.5078 \[astro-ph.CO\]](#).

- [356] **POLARBEAR** Collaboration, P. A. R. Ade *et al.*, “Evidence for Gravitational Lensing of the Cosmic Microwave Background Polarization from Cross-correlation with the Cosmic Infrared Background,” *Phys. Rev. Lett.* **112** (2014) 131302, [arXiv:1312.6645](#) [astro-ph.CO].
- [357] **ACT** Collaboration, A. van Engelen *et al.*, “The Atacama Cosmology Telescope: Lensing of CMB Temperature and Polarization Derived from Cosmic Infrared Background Cross-Correlation,” *Astrophys. J.* **808** no. 1, (2015) 7, [arXiv:1412.0626](#) [astro-ph.CO].
- [358] F. Bianchini *et al.*, “Cross-correlation between the CMB lensing potential measured by Planck and high- z sub-mm galaxies detected by the Herschel-ATLAS survey,” *Astrophys. J.* **802** no. 1, (2015) 64, [arXiv:1410.4502](#) [astro-ph.CO].
- [359] L. E. Bleem *et al.*, “A Measurement of the Correlation of Galaxy Surveys with CMB Lensing Convergence Maps from the South Pole Telescope,” *Astrophys. J.* **753** (2012) L9, [arXiv:1203.4808](#) [astro-ph.CO].
- [360] **Planck** Collaboration, Planck Collaboration XVII, “Planck 2013 results. XVII. Gravitational lensing by large-scale structure,” *Astron. Astrophys.* **571** (2014) A17, [arXiv:1303.5077](#) [astro-ph.CO].
- [361] J. E. Geach *et al.*, “A Direct Measurement of the Linear Bias of Mid-infrared-selected Quasars at $z \approx 1$ Using Cosmic Microwave Background Lensing,” *Astrophys. J.* **776** (2013) L41, [arXiv:1307.1706](#) [astro-ph.CO].
- [362] A. R. Pullen, S. Alam, S. He, and S. Ho, “Constraining Gravity at the Largest Scales through CMB Lensing and Galaxy Velocities,” *Mon. Not. Roy. Astron. Soc.* **460** no. 4, (2016) 4098–4108, [arXiv:1511.04457](#) [astro-ph.CO].
- [363] S. He, S. Alam, S. Ferraro, Y.-C. Chen, and S. Ho, “The detection of the imprint of filaments on cosmic microwave background lensing,” *Nat. Astron.* **2** no. 5, (2018) 401–406, [arXiv:1709.02543](#) [astro-ph.CO].
- [364] Y.-C. Cai, M. Neyrinck, Q. Mao, J. A. Peacock, I. Szapudi, and A. A. Berlind, “The lensing and temperature imprints of voids on the Cosmic Microwave Background,” *Mon. Not. Roy. Astron. Soc.* **466** no. 3, (2017) 3364–3375, [arXiv:1609.00301](#) [astro-ph.CO].
- [365] B. D. Sherwin *et al.*, “The Atacama Cosmology Telescope: Cross-Correlation of CMB Lensing and Quasars,” *Phys. Rev.* **D86** (2012) 083006, [arXiv:1207.4543](#) [astro-ph.CO].
- [366] N. Hand *et al.*, “First Measurement of the Cross-Correlation of CMB Lensing and Galaxy Lensing,” *Phys. Rev.* **D91** no. 6, (2015) 062001, [arXiv:1311.6200](#) [astro-ph.CO].
- [367] J. Liu and J. C. Hill, “Cross-correlation of Planck CMB Lensing and CFHTLenS Galaxy Weak Lensing Maps,” *Phys. Rev.* **D92** no. 6, (2015) 063517, [arXiv:1504.05598](#) [astro-ph.CO].
- [368] **DES, SPT** Collaboration, E. J. Baxter *et al.*, “Joint Measurement of Lensing-Galaxy Correlations Using SPT and DES SV Data,” *Mon. Not. Roy. Astron. Soc.* **461** (Oct., 2016) 4099–4114, [arXiv:1602.07384](#) [astro-ph.CO].
- [369] **DES** Collaboration, T. Giannantonio *et al.*, “CMB lensing tomography with the DES Science Verification galaxies,” *Mon. Not. Roy. Astron. Soc.* **456** no. 3, (2016) 3213–3244, [arXiv:1507.05551](#) [astro-ph.CO].
- [370] C. Doux, E. Schaan, E. Aubourg, K. Ganga, K.-G. Lee, D. N. Spergel, and J. Tréguer, “First detection of cosmic microwave background lensing and Lyman- α forest bispectrum,” *Phys. Rev.* **D94** no. 10, (2016) 103506, [arXiv:1607.03625](#) [astro-ph.CO].

- [371] G. Hurier, P. Singh, and C. Hernández-Monteagudo, “First measurement of the cross-correlation between CMB weak lensing and X-ray emission,” [arXiv:1711.10774](#) [astro-ph.CO].
- [372] N. Fornengo, L. Perotto, M. Regis, and S. Camera, “Evidence of Cross-correlation between the CMB Lensing and the Γ -ray sky,” *Astrophys. J.* **802** no. 1, (2015) L1, [arXiv:1410.4997](#) [astro-ph.CO].
- [373] E. J. Baxter *et al.*, “A Measurement of Gravitational Lensing of the Cosmic Microwave Background by Galaxy Clusters Using Data from the South Pole Telescope,” *Astrophys. J.* **806** no. 2, (2015) 247, [arXiv:1412.7521](#) [astro-ph.CO].
- [374] **Planck** Collaboration, Planck Collaboration XXIV, “Planck 2015 results. XXIV. Cosmology from Sunyaev-Zeldovich cluster counts,” *Astron. Astrophys.* **594** (2016) A24, [arXiv:1502.01597](#) [astro-ph.CO].
- [375] **ACT** Collaboration, M. Madhavacheril *et al.*, “Evidence of Lensing of the Cosmic Microwave Background by Dark Matter Halos,” *Phys. Rev. Lett.* **114** no. 15, (2015) 151302, [arXiv:1411.7999](#) [astro-ph.CO]. [Addendum: *Phys. Rev. Lett.*114,no.18,189901(2015)].
- [376] **DESI** Collaboration, M. Levi *et al.*, “The DESI Experiment, a whitepaper for Snowmass 2013,” [arXiv:1308.0847](#) [astro-ph.CO].
- [377] Planck Collaboration XXI, “Planck 2013 results. XXI. Power spectrum and high-order statistics of the Planck all-sky Compton parameter map,” *A & A* **571** (Nov., 2014) A21, [arXiv:1303.5081](#).
- [378] Planck Collaboration XXII, “Planck 2015 results. XXII. A map of the thermal Sunyaev-Zeldovich effect,” *A & A* **594** (Sept., 2016) A22, [arXiv:1502.01596](#).
- [379] R. Khatri, “An alternative validation strategy for the Planck cluster catalogue and y -distortion maps,” *A & A* **592** (July, 2016) A48, [arXiv:1505.00778](#).
- [380] M. Fukugita and P. J. E. Peebles, “The Cosmic Energy Inventory,” *Ap. J.* **616** (Dec., 2004) 643–668, [astro-ph/0406095](#).
- [381] A. Gallazzi, J. Brinchmann, S. Charlot, and S. D. M. White, “A census of metals and baryons in stars in the local Universe,” *Mon. Not. Roy. Astron. Soc.* **383** (Feb., 2008) 1439–1458, [arXiv:0708.0533](#).
- [382] K. Schawinski, D. Thomas, M. Sarzi, C. Maraston, S. Kaviraj, S.-J. Joo, S. K. Yi, and J. Silk, “Observational evidence for AGN feedback in early-type galaxies,” *Mon. Not. Roy. Astron. Soc.* **382** (Dec., 2007) 1415–1431, [arXiv:0709.3015](#).
- [383] J. N. Bregman, “The Search for the Missing Baryons at Low Redshift,” *Ann. Rev. Astron. Astroph.* **45** (Sept., 2007) 221–259, [arXiv:0706.1787](#).
- [384] T. Eifler, E. Krause, S. Dodelson, A. R. Zentner, A. P. Hearin, and N. Y. Gnedin, “Accounting for baryonic effects in cosmic shear tomography: determining a minimal set of nuisance parameters using PCA,” *Mon. Not. Roy. Astron. Soc.* **454** (Dec., 2015) 2451–2471, [arXiv:1405.7423](#).
- [385] D. Nagai, A. V. Kravtsov, and A. Vikhlinin, “Effects of Galaxy Formation on Thermodynamics of the Intracluster Medium,” *Astrophys. J.* **668** (2007) 1–14, [arXiv:astro-ph/0703661](#) [astro-ph].
- [386] E. T. Lau, D. Nagai, and K. Nelson, “Weighing Galaxy Clusters with Gas. I. On the Methods of Computing Hydrostatic Mass Bias,” *Ap. J.* **777** no. 2, (Nov, 2013) 151, [arXiv:1306.3993](#) [astro-ph.CO].

- [387] X. Shi, E. Komatsu, D. Nagai, and E. T. Lau, “Analytical model for non-thermal pressure in galaxy clusters - III. Removing the hydrostatic mass bias,” *Mon. Not. Roy. Astron. Soc.* **455** no. 3, (Jan, 2016) 2936–2944, [arXiv:1507.04338](#) [astro-ph.CO].
- [388] E. T. Lau, A. V. Kravtsov, and D. Nagai, “Residual Gas Motions in the Intracluster Medium and Bias in Hydrostatic Measurements of Mass Profiles of Clusters,” *Ap. J.* **705** (Nov., 2009) 1129–1138, [arXiv:0903.4895](#) [astro-ph.CO].
- [389] K. Nelson, E. T. Lau, and D. Nagai, “Hydrodynamic Simulation of Non-thermal Pressure Profiles of Galaxy Clusters,” *Ap. J.* **792** (Sept., 2014) 25, [arXiv:1404.4636](#).
- [390] J. D. Emberson, N. Frontiere, S. Habib, K. Heitmann, P. Larsen, H. Finkel, and A. Pope, “The Borg Cube Simulation: Cosmological Hydrodynamics with CRK-SPH,” *ArXiv e-prints* (Nov., 2018) , [arXiv:1811.03593](#).
- [391] N. Battaglia, S. Ferraro, E. Schaan, and D. N. Spergel, “Future constraints on halo thermodynamics from combined Sunyaev-Zel’dovich measurements,” *JCAP* **11** (Nov., 2017) 040, [arXiv:1705.05881](#).
- [392] K. S. Dawson *et al.*, “The Baryon Oscillation Spectroscopic Survey of SDSS-III,” *Astronomical Journal* **145** (2013) 10, [arXiv:1208.0022](#) [astro-ph.CO].
- [393] S. Ho, S. Dedeo, and D. Spergel, “Finding the Missing Baryons Using CMB as a Backlight,” [arXiv:0903.2845](#) [astro-ph.CO].
- [394] ACTPol Collaboration, E. Schaan *et al.*, “Evidence for the kinematic Sunyaev-Zel’dovich effect with the Atacama Cosmology Telescope and velocity reconstruction from the Baryon Oscillation Spectroscopic Survey,” *Phys. Rev.* **D93** no. 8, (2016) 082002, [arXiv:1510.06442](#) [astro-ph.CO].
- [395] I. G. McCarthy, J. Schaye, S. Bird, and A. M. C. Le Brun, “The BAHAMAS project: calibrated hydrodynamical simulations for large-scale structure cosmology,” *Mon. Not. Roy. Astron. Soc.* **465** (Mar., 2017) 2936–2965, [arXiv:1603.02702](#).
- [396] N. Battaglia, J. R. Bond, C. Pfrommer, J. L. Sievers, and D. Sijacki, “Simulations of the Sunyaev-Zel’dovich Power Spectrum with Active Galactic Nucleus Feedback,” *Ap. J.* **725** (Dec., 2010) 91–99, [arXiv:1003.4256](#).
- [397] N. Battaglia, J. R. Bond, C. Pfrommer, and J. L. Sievers, “On the Cluster Physics of Sunyaev-Zel’dovich and X-Ray Surveys. II. Deconstructing the Thermal SZ Power Spectrum,” *Ap. J.* **758** (Oct., 2012) 75, [arXiv:1109.3711](#).
- [398] J. Schaye, R. A. Crain, R. G. Bower, M. Furlong, M. Schaller, T. Theuns, C. Dalla Vecchia, C. S. Frenk, I. G. McCarthy, J. C. Helly, A. Jenkins, Y. M. Rosas-Guevara, S. D. M. White, M. Baes, C. M. Booth, P. Camps, J. F. Navarro, Y. Qu, A. Rahmati, T. Sawala, P. A. Thomas, and J. Trayford, “The EAGLE project: simulating the evolution and assembly of galaxies and their environments,” *Mon. Not. Roy. Astron. Soc.* **446** (Jan., 2015) 521–554, [arXiv:1407.7040](#).
- [399] R. Weinberger, V. Springel, L. Hernquist, A. Pillepich, F. Marinacci, R. Pakmor, D. Nelson, S. Genel, M. Vogelsberger, J. Naiman, and P. Torrey, “Simulating galaxy formation with black hole driven thermal and kinetic feedback,” *Mon. Not. Roy. Astron. Soc.* **465** (Mar., 2017) 3291–3308, [arXiv:1607.03486](#).
- [400] D. J. Barnes, M. Vogelsberger, R. Kannan, F. Marinacci, R. Weinberger, V. Springel, P. Torrey, A. Pillepich, D. Nelson, R. Pakmor, J. Naiman, L. Hernquist, and M. McDonald, “A census of cool-core galaxy clusters in IllustrisTNG,” *Mon. Not. Roy. Astron. Soc.* **481** (Dec., 2018) 1809–1831, [arXiv:1710.08420](#).

- [401] V. Springel, R. Pakmor, A. Pillepich, R. Weinberger, D. Nelson, L. Hernquist, M. Vogelsberger, S. Genel, P. Torrey, F. Marinacci, and J. Naiman, “First results from the IllustrisTNG simulations: matter and galaxy clustering,” *Mon. Not. Roy. Astron. Soc.* **475** (Mar., 2018) 676–698, [arXiv:1707.03397](#).
- [402] Planck Collaboration Int. XI, “Planck intermediate results. XI. The gas content of dark matter halos: the Sunyaev-Zeldovich-stellar mass relation for locally brightest galaxies,” *A & A* **557** (Sept., 2013) A52, [arXiv:1212.4131](#).
- [403] A. M. C. Le Brun, I. G. McCarthy, and J.-B. Melin, “Testing Sunyaev-Zel’dovich measurements of the hot gas content of dark matter haloes using synthetic skies,” *Mon. Not. Roy. Astron. Soc.* **451** (Aug., 2015) 3868–3881, [arXiv:1501.05666](#).
- [404] J. C. Hill, E. J. Baxter, A. Lidz, J. P. Greco, and B. Jain, “Two-halo term in stacked thermal Sunyaev-Zel’dovich measurements: Implications for self-similarity,” *Phys. Rev. D* **97** no. 8, (Apr., 2018) 083501, [arXiv:1706.03753](#).
- [405] J. C. Hill, S. Ferraro, N. Battaglia, J. Liu, and D. N. Spergel, “Kinematic Sunyaev-Zeldovich Effect with Projected Fields: A Novel Probe of the Baryon Distribution with Planck, WMAP, and WISE Data,” *Phys. Rev. Lett.* **117** no. 5, (2016) 051301, [arXiv:1603.01608](#) [[astro-ph.CO](#)].
- [406] S. Ferraro, J. C. Hill, N. Battaglia, J. Liu, and D. N. Spergel, “The Kinematic Sunyaev-Zel’dovich Effect with Projected Fields II: prospects, challenges, and comparison with simulations,” [arXiv:1605.02722](#) [[astro-ph.CO](#)].
- [407] V. Vikram, A. Lidz, and B. Jain, “A Measurement of the Galaxy Group-Thermal Sunyaev-Zel’dovich Effect Cross-Correlation Function,” *Mon. Not. Roy. Astron. Soc.* **467** (May, 2017) 2315–2330, [arXiv:1608.04160](#).
- [408] A. Hall and A. Challinor, “Detecting the polarization induced by scattering of the microwave background quadrupole in galaxy clusters,” *Phys. Rev. D* **90** no. 6, (Sept., 2014) 063518, [arXiv:1407.5135](#).
- [409] A.-S. Deutsch, M. C. Johnson, M. Münchmeyer, and A. Terrana, “Polarized Sunyaev Zel’dovich tomography,” *JCAP* **4** (Apr., 2018) 034, [arXiv:1705.08907](#).
- [410] C. L. Reichardt, L. Shaw, O. Zahn, K. A. Aird, B. A. Benson, L. E. Bleem, J. E. Carlstrom, C. L. Chang, H. M. Cho, T. M. Crawford, A. T. Crites, T. de Haan, M. A. Dobbs, J. Dudley, E. M. George, N. W. Halverson, G. P. Holder, W. L. Holzappel, S. Hoover, Z. Hou, J. D. Hrubes, M. Joy, R. Keisler, L. Knox, A. T. Lee, E. M. Leitch, M. Lueker, D. Luong-Van, J. J. McMahon, J. Mehl, S. S. Meyer, M. Millea, J. J. Mohr, T. E. Montroy, T. Natoli, S. Padin, T. Plagge, C. Pryke, J. E. Ruhl, K. K. Schaffer, E. Shirokoff, H. G. Spieler, Z. Staniszewski, A. A. Stark, K. Story, A. van Engelen, K. Vanderlinde, J. D. Vieira, and R. Williamson, “A Measurement of Secondary Cosmic Microwave Background Anisotropies with Two Years of South Pole Telescope Observations,” *Ap. J.* **755** (Aug., 2012) 70, [arXiv:1111.0932](#) [[astro-ph.CO](#)].
- [411] O. Zahn, C. L. Reichardt, L. Shaw, A. Lidz, K. A. Aird, B. A. Benson, L. E. Bleem, J. E. Carlstrom, C. L. Chang, H. M. Cho, T. M. Crawford, A. T. Crites, T. de Haan, M. A. Dobbs, O. Doré, J. Dudley, E. M. George, N. W. Halverson, G. P. Holder, W. L. Holzappel, S. Hoover, Z. Hou, J. D. Hrubes, M. Joy, R. Keisler, L. Knox, A. T. Lee, E. M. Leitch, M. Lueker, D. Luong-Van, J. J. McMahon, J. Mehl, S. S. Meyer, M. Millea, J. J. Mohr, T. E. Montroy, T. Natoli, S. Padin, T. Plagge, C. Pryke, J. E. Ruhl, K. K. Schaffer, E. Shirokoff, H. G. Spieler, Z. Staniszewski, A. A. Stark, K. Story, A. van Engelen, K. Vanderlinde, J. D. Vieira, and R. Williamson, “Cosmic Microwave Background

- Constraints on the Duration and Timing of Reionization from the South Pole Telescope,” *Ap. J.* **756** (Sept., 2012) 65, [arXiv:1111.6386](#).
- [412] E. M. George, C. L. Reichardt, K. A. Aird, B. A. Benson, L. E. Bleem, J. E. Carlstrom, C. L. Chang, H.-M. Cho, T. M. Crawford, A. T. Crites, T. de Haan, M. A. Dobbs, J. Dudley, N. W. Halverson, N. L. Harrington, G. P. Holder, W. L. Holzapfel, Z. Hou, J. D. Hrubes, R. Keisler, L. Knox, A. T. Lee, E. M. Leitch, M. Lueker, D. Luong-Van, J. J. McMahon, J. Mehl, S. S. Meyer, M. Millea, L. M. Mocuano, J. J. Mohr, T. E. Montroy, S. Padin, T. Plagge, C. Pryke, J. E. Ruhl, K. K. Schaffer, L. Shaw, E. Shirokoff, H. G. Spieler, Z. Staniszewski, A. A. Stark, K. T. Story, A. van Engelen, K. Vanderlinde, J. D. Vieira, R. Williamson, and O. Zahn, “A Measurement of Secondary Cosmic Microwave Background Anisotropies from the 2500 Square-degree SPT-SZ Survey,” *Ap. J.* **799** (Feb., 2015) 177, [arXiv:1408.3161](#).
- [413] **Planck** Collaboration, Planck Collaboration Int. XLVII, “Planck intermediate results. XLVII. Planck constraints on reionization history,” [arXiv:1605.03507](#) [[astro-ph.CO](#)].
- [414] E. Calabrese *et al.*, “Precision Epoch of Reionization studies with next-generation CMB experiments,” *JCAP* **1408** (2014) 010, [arXiv:1406.4794](#) [[astro-ph.CO](#)].
- [415] C. Dvorkin and K. M. Smith, “Reconstructing Patchy Reionization from the Cosmic Microwave Background,” *Phys. Rev. D* **79** (2009) 043003, [arXiv:0812.1566](#) [[astro-ph](#)].
- [416] C. Dvorkin, W. Hu, and K. M. Smith, “B-mode CMB polarization from patchy screening during reionization,” *Phys. Rev. D* **79** no. 10, (May, 2009) 107302, [arXiv:0902.4413](#) [[astro-ph.CO](#)].
- [417] A. Roy, A. Lapi, D. Spergel, and C. Baccigalupi, “Observing patchy reionization with future CMB polarization experiments,” *JCAP* **5** (May, 2018) 014, [arXiv:1801.02393](#).
- [418] C. Feng and G. Holder, “Detecting electron density fluctuations from cosmic microwave background polarization using a bispectrum approach,” *Phys. Rev. D* **97** no. 12, (June, 2018) 123523, [arXiv:1801.05396](#).
- [419] C. Feng and G. Holder, “Searching for patchy reionization from cosmic microwave background with hybrid quadratic estimators,” *ArXiv e-prints* (Aug., 2018) , [arXiv:1808.01592](#).
- [420] M. A. Alvarez and T. Abel, “The Effect of Absorption Systems on Cosmic Reionization,” *Ap. J.* **747** (Mar., 2012) 126, [arXiv:1003.6132](#) [[astro-ph.CO](#)].
- [421] M. A. Alvarez, “The Kinetic Sunyaev-Zel’dovich Effect from Reionization: Simulated Full-sky Maps at Arcminute Resolution,” *Ap. J.* **824** (June, 2016) 118, [arXiv:1511.02846](#).
- [422] L. D. Shaw, D. H. Rudd, and D. Nagai, “Deconstructing the Kinetic SZ Power Spectrum,” *Ap. J.* **756** (Sept., 2012) 15, [arXiv:1109.0553](#) [[astro-ph.CO](#)].
- [423] H. Park, M. A. Alvarez, and J. R. Bond, “The Impact of Baryonic Physics on the Kinetic Sunyaev-Zel’dovich Effect,” *Ap. J.* **853** (Feb., 2018) 121, [arXiv:1710.02792](#).
- [424] K. M. Smith and S. Ferraro, “Detecting Patchy Reionization in the Cosmic Microwave Background,” *Physical Review Letters* **119** no. 2, (July, 2017) 021301.
- [425] R. A. Sunyaev and Ya. B. Zeldovich, “The Observations of relic radiation as a test of the nature of X-Ray radiation from the clusters of galaxies,” *Comments Astrophys. Space Phys.* **4** (1972) 173–178.
- [426] J. E. Carlstrom, G. P. Holder, and E. D. Reese, “Cosmology with the Sunyaev-Zel’dovich Effect,” *Ann. Rev. Astron. Astroph.* **40** (2002) 643–680.

- [427] D. H. Weinberg, M. J. Mortonson, D. J. Eisenstein, C. Hirata, A. G. Riess, and E. Rozo, “Observational Probes of Cosmic Acceleration,” *Phys. Rept.* **530** (2013) 87–255, [arXiv:1201.2434 \[astro-ph.CO\]](#).
- [428] M. Niemack, P. Ade, J. Aguirre, F. Barrientos, J. Beall, *et al.*, “ACTPol: A polarization-sensitive receiver for the Atacama Cosmology Telescope,” *Proc.SPIE Int.Soc.Opt.Eng.* **7741** (2010) 77411S, [arXiv:1006.5049 \[astro-ph.IM\]](#).
- [429] **SPT-3G** Collaboration, B. A. Benson *et al.*, “SPT-3G: A Next-Generation Cosmic Microwave Background Polarization Experiment on the South Pole Telescope,” *Proc. SPIE Int. Soc. Opt. Eng.* **9153** (2014) 91531P, [arXiv:1407.2973 \[astro-ph.IM\]](#).
- [430] M. S. Madhavacheril, N. Battaglia, and H. Miyatake, “Fundamental physics from future weak-lensing calibrated Sunyaev-Zel’dovich galaxy cluster counts,” *Phys. Rev. D* **96** no. 10, (Nov., 2017) 103525, [arXiv:1708.07502](#).
- [431] Planck Collaboration, P. A. R. Ade, N. Aghanim, M. Arnaud, M. Ashdown, J. Aumont, C. Baccigalupi, A. J. Banday, R. B. Barreiro, R. Barrena, and *et al.*, “Planck 2015 results. XXVII. The second Planck catalogue of Sunyaev-Zeldovich sources,” *A & A* **594** (Sept., 2016) A27, [arXiv:1502.01598](#).
- [432] **SPT** Collaboration, L. E. Bleem *et al.*, “Galaxy Clusters Discovered via the Sunyaev-Zel’dovich Effect in the 2500-square-degree SPT-SZ survey,” *Astrophys. J. Suppl.* **216** no. 2, (2015) 27, [arXiv:1409.0850 \[astro-ph.CO\]](#).
- [433] M. Hilton, M. Hasselfield, C. Sifón, N. Battaglia, S. Aiola, V. Bharadwaj, J. R. Bond, S. K. Choi, D. Crichton, R. Datta, M. J. Devlin, J. Dunkley, R. Dünner, P. A. Gallardo, M. Gralla, A. D. Hincks, S.-P. P. Ho, J. Hubmayr, K. M. Huffenberger, J. P. Hughes, B. J. Koopman, A. Kosowsky, T. Louis, M. S. Madhavacheril, T. A. Marriage, L. Maurin, J. McMahon, H. Miyatake, K. Moodley, S. Naess, F. Nati, L. Newburgh, M. D. Niemack, M. Oguri, L. A. Page, B. Partridge, B. L. Schmitt, J. Sievers, D. N. Spergel, S. T. Staggs, H. Trac, A. van Engelen, E. M. Vavagiakis, and E. J. Wollack, “The Atacama Cosmology Telescope: The Two-season ACTPol Sunyaev-Zeldovich Effect Selected Cluster Catalog,” *Ap. J. Suppl.* **235** (Mar., 2018) 20, [arXiv:1709.05600](#).
- [434] M. Negrello, F. Perrotta, J. González-Nuevo, L. Silva, G. de Zotti, G. L. Granato, C. Baccigalupi, and L. Danese, “Astrophysical and cosmological information from large-scale submillimetre surveys of extragalactic sources,” *Mon. Not. Roy. Astron. Soc.* **377** (June, 2007) 1557–1568, [astro-ph/0703210](#).
- [435] J. Vieira, T. Crawford, E. Switzer, P. Ade, K. Aird, *et al.*, “Extragalactic millimeter-wave sources in South Pole Telescope survey data: source counts, catalog, and statistics for an 87 square-degree field,” *Astrophys.J.* **719** (2010) 763–783, [arXiv:0912.2338 \[astro-ph.CO\]](#).
- [436] M. L. Strandet, A. Weiss, C. De Breuck, D. P. Marrone, J. D. Vieira, M. Aravena, M. L. N. Ashby, M. Béthermin, M. S. Bothwell, C. M. Bradford, J. E. Carlstrom, S. C. Chapman, D. J. M. Cunningham, C.-C. Chen, C. D. Fassnacht, A. H. Gonzalez, T. R. Greve, B. Gullberg, C. C. Hayward, Y. Hezaveh, K. Litke, J. Ma, M. Malkan, K. M. Menten, T. Miller, E. J. Murphy, D. Narayanan, K. A. Phadke, K. M. Rotermund, J. S. Spilker, and J. Sreevani, “ISM Properties of a Massive Dusty Star-forming Galaxy Discovered at $z \sim 7$,” *Ap. J. Lett.* **842** (June, 2017) L15, [arXiv:1705.07912](#).
- [437] D. P. Marrone, J. S. Spilker, C. C. Hayward, J. D. Vieira, M. Aravena, M. L. N. Ashby, M. B. Bayliss, M. Béthermin, M. Brodwin, M. S. Bothwell, J. E. Carlstrom, S. C. Chapman, C.-C. Chen, T. M. Crawford, D. J. M. Cunningham, C. De Breuck, C. D. Fassnacht, A. H. Gonzalez, T. R. Greve, Y. D. Hezaveh, K. Lacaille, K. C. Litke, S. Lower, J. Ma, M. Malkan, T. B. Miller, W. R. Morningstar, E. J. Murphy, D. Narayanan, K. A. Phadke, K. M. Rotermund, J. Sreevani, B. Stalder, A. A. Stark, M. L. Strandet, M. Tang, and A. Weiß, “Galaxy growth in a massive halo in the first billion years of cosmic history,” *Nature* **553** (Jan., 2018) 51–54, [arXiv:1712.03020](#).

- [438] T. B. Miller, S. C. Chapman, M. Aravena, M. L. N. Ashby, C. C. Hayward, J. D. Vieira, A. Weiß, A. Babul, M. Béthermin, C. M. Bradford, M. Brodwin, J. E. Carlstrom, C.-C. Chen, D. J. M. Cunningham, C. De Breuck, A. H. Gonzalez, T. R. Greve, J. Harnett, Y. Hezaveh, K. Lacaille, K. C. Litke, J. Ma, M. Malkan, D. P. Marrone, W. Morningstar, E. J. Murphy, D. Narayanan, E. Pass, R. Perry, K. A. Phadke, D. Rennehan, K. M. Rotermond, J. Simpson, J. S. Spilker, J. Sreevani, A. A. Stark, M. L. Strandet, and A. L. Strom, “A massive core for a cluster of galaxies at a redshift of 4.3,” *Nature* **556** (Apr., 2018) 469–472, [arXiv:1804.09231](#).
- [439] A. Weiß, C. De Breuck, D. P. Marrone, J. D. Vieira, J. E. Aguirre, K. A. Aird, M. Aravena, M. L. N. Ashby, M. Bayliss, B. A. Benson, M. Béthermin, A. D. Biggs, L. E. Bleem, J. J. Bock, M. Bothwell, C. M. Bradford, M. Brodwin, J. E. Carlstrom, C. L. Chang, S. C. Chapman, T. M. Crawford, A. T. Crites, T. de Haan, M. A. Dobbs, T. P. Downes, C. D. Fassnacht, E. M. George, M. D. Gladders, A. H. Gonzalez, T. R. Greve, N. W. Halverson, Y. D. Hezaveh, F. W. High, G. P. Holder, W. L. Holzapfel, S. Hoover, J. D. Hrubes, K. Husband, R. Keisler, A. T. Lee, E. M. Leitch, M. Lueker, D. Luong-Van, M. Malkan, V. McIntyre, J. J. McMahon, J. Mehl, K. M. Menten, S. S. Meyer, E. J. Murphy, S. Padin, T. Plagge, C. L. Reichardt, A. Rest, M. Rosenman, J. Ruel, J. E. Ruhl, K. K. Schaffer, E. Shirokoff, J. S. Spilker, B. Stalder, Z. Staniszewski, A. A. Stark, K. Story, K. Vanderlinde, N. Welikala, and R. Williamson, “ALMA Redshifts of Millimeter-selected Galaxies from the SPT Survey: The Redshift Distribution of Dusty Star-forming Galaxies,” *Ap. J.* **767** (Apr., 2013) 88, [arXiv:1303.2726 \[astro-ph.CO\]](#).
- [440] M. L. Strandet, A. Weiss, J. D. Vieira, C. de Breuck, J. E. Aguirre, M. Aravena, M. L. N. Ashby, M. Béthermin, C. M. Bradford, J. E. Carlstrom, S. C. Chapman, T. M. Crawford, W. Everett, C. D. Fassnacht, R. M. Furstenuau, A. H. Gonzalez, T. R. Greve, B. Gullberg, Y. Hezaveh, J. R. Kamenetzky, K. Litke, J. Ma, M. Malkan, D. P. Marrone, K. M. Menten, E. J. Murphy, A. Nadolski, K. M. Rotermond, J. S. Spilker, A. A. Stark, and N. Welikala, “The Redshift Distribution of Dusty Star-forming Galaxies from the SPT Survey,” *Ap. J.* **822** (May, 2016) 80, [arXiv:1603.05094](#).
- [441] Y. D. Hezaveh, N. Dalal, D. P. Marrone, Y.-Y. Mao, W. Morningstar, D. Wen, R. D. Blandford, J. E. Carlstrom, C. D. Fassnacht, G. P. Holder, A. Kembal, P. J. Marshall, N. Murray, L. Perreault Levasseur, J. D. Vieira, and R. H. Wechsler, “Detection of Lensing Substructure Using ALMA Observations of the Dusty Galaxy SDP.81,” *Ap. J.* **823** (May, 2016) 37, [arXiv:1601.01388](#).
- [442] M. Béthermin, C. De Breuck, M. Sargent, and E. Daddi, “The influence of wavelength, flux, and lensing selection effects on the redshift distribution of dusty, star-forming galaxies,” *A & A* **576** (Apr., 2015) L9, [arXiv:1502.04710](#).
- [443] Planck Collaboration Int. XXXIX, “Planck intermediate results. XXXIX. The Planck list of high-redshift source candidates,” *A & A* **596** (Dec., 2016) A100, [arXiv:1508.04171](#).
- [444] M. Tucci, L. Toffolatti, G. de Zotti, and E. Martínez-González, “High-frequency predictions for number counts and spectral properties of extragalactic radio sources. New evidence of a break at mm wavelengths in spectra of bright blazar sources,” *A & A* **533** (Sept., 2011) A57, [arXiv:1103.5707](#).
- [445] G. Puglisi, V. Galluzzi, L. Bonavera, J. Gonzalez-Nuevo, A. Lapi, M. Massardi, F. Perrotta, C. Baccigalupi, A. Celotti, and L. Danese, “Forecasting Polarized Radio Sources for CMB observations,” *ArXiv e-prints* (Dec., 2017) , [arXiv:1712.09639](#).
- [446] T. Murphy, E. M. Sadler, R. D. Ekers, M. Massardi, P. J. Hancock, E. Mahony, R. Ricci, S. Burke-Spolaor, M. Calabretta, R. Chhetri, G. de Zotti, P. G. Edwards, J. A. Ekers, C. A. Jackson, M. J. Kesteven, E. Lindley, K. Newton-McGee, C. Phillips, P. Roberts, R. J. Sault, L. Staveley-Smith, R. Subrahmanyam, M. A. Walker, and W. E. Wilson, “The Australia Telescope 20 GHz Survey: the source catalogue,” *Mon. Not. Roy. Astron. Soc.* **402** (Mar., 2010) 2403–2423, [arXiv:0911.0002](#).

- [447] L. M. Mocanu, T. M. Crawford, J. D. Vieira, K. A. Aird, M. Aravena, J. E. Austermann, B. A. Benson, M. Béthermin, L. E. Bleem, M. Bothwell, J. E. Carlstrom, C. L. Chang, S. Chapman, H.-M. Cho, A. T. Crites, T. de Haan, M. A. Dobbs, W. B. Everett, E. M. George, N. W. Halverson, N. Harrington, Y. Hezaveh, G. P. Holder, W. L. Holzapfel, S. Hoover, J. D. Hrubes, R. Keisler, L. Knox, A. T. Lee, E. M. Leitch, M. Lueker, D. Luong-Van, D. P. Marrone, J. J. McMahon, J. Mehl, S. S. Meyer, J. J. Mohr, T. E. Montroy, T. Natoli, S. Padin, T. Plagge, C. Pryke, A. Rest, C. L. Reichardt, J. E. Ruhl, J. T. Sayre, K. K. Schaffer, E. Shirokoff, H. G. Spieler, J. S. Spilker, B. Stalder, Z. Staniszewski, A. A. Stark, K. T. Story, E. R. Switzer, K. Vanderlinde, and R. Williamson, “Extragalactic Millimeter-wave Point-source Catalog, Number Counts and Statistics from 771 deg² of the SPT-SZ Survey,” *Ap. J.* **779** (Dec., 2013) 61, [arXiv:1306.3470](#).
- [448] D. Marsden, M. Gralla, T. A. Marriage, E. R. Switzer, B. Partridge, M. Massardi, G. Morales, G. Addison, J. R. Bond, D. Crichton, S. Das, M. Devlin, R. Dünner, A. Hajian, M. Hilton, A. Hincks, J. P. Hughes, K. Irwin, A. Kosowsky, F. Menanteau, K. Moodley, M. Niemack, L. Page, E. D. Reese, B. Schmitt, N. Sehgal, J. Sievers, S. Staggs, D. Swetz, R. Thornton, and E. Wollack, “The Atacama Cosmology Telescope: dusty star-forming galaxies and active galactic nuclei in the Southern survey,” *Mon. Not. Roy. Astron. Soc.* **439** (Apr., 2014) 1556–1574, [arXiv:1306.2288](#) [astro-ph.CO].
- [449] Planck Collaboration XXVI, “Planck 2015 results. XXVI. The Second Planck Catalogue of Compact Sources,” *A & A* **594** (Sept., 2016) A26, [arXiv:1507.02058](#).
- [450] B. Partridge, L. Bonavera, M. López-Caniego, R. Datta, J. Gonzalez-Nuevo, M. Gralla, D. Herranz, A. Lähteenmäki, L. Mocanu, H. Prince, J. Vieira, N. Whitehorn, and L. Zhang, “Can CMB Surveys Help the AGN Community?,” *Galaxies* **5** (Aug., 2017) 47.
- [451] A. Kolodzig, M. Gilfanov, R. Sunyaev, S. Sazonov, and M. Brusa, “AGN and QSOs in the eROSITA All-Sky Survey. I. Statistical properties,” *A & A* **558** (Oct., 2013) A89, [arXiv:1212.2151](#) [astro-ph.CO].
- [452] Planck Collab. Int. XLV, “Planck intermediate results. XLV. Radio spectra of northern extragalactic radio sources,” *A & A* **596** (Dec., 2016) A106, [arXiv:1606.05120](#).
- [453] A. Lähteenmäki, E. Järvelä, T. Hovatta, M. Tornikoski, D. L. Harrison, M. López-Caniego, W. Max-Moerbeck, M. Mingaliev, T. J. Pearson, V. Ramakrishnan, A. C. S. Readhead, R. A. Reeves, J. L. Richards, Y. Sotnikova, and J. Tammi, “37 GHz observations of narrow-line Seyfert 1 galaxies,” *A & A* **603** (July, 2017) A100, [arXiv:1703.10365](#).
- [454] A. Lähteenmäki, E. Järvelä, V. Ramakrishnan, M. Tornikoski, J. Tammi, R. J. C. Vera, and W. Chamani, “Radio jets and gamma-ray emission in radio-silent narrow-line Seyfert 1 galaxies,” *A & A* **614** (June, 2018) L1, [arXiv:1806.02058](#).
- [455] S. Chen, M. Berton, G. La Mura, E. Congiu, V. Cracco, L. Foschini, J. H. Fan, S. Ciroi, P. Rafanelli, and D. Bastieri, “Probing narrow-line Seyfert 1 galaxies in the southern hemisphere,” *A & A* **615** (Aug., 2018) A167, [arXiv:1801.07234](#).
- [456] B.-G. Andersson, A. Lazarian, and J. E. Vaillancourt, “Interstellar Dust Grain Alignment,” *Ann. Rev. Astron. Astroph.* **53** (Aug., 2015) 501–539.
- [457] L. Davis, Jr. and J. L. Greenstein, “The Polarization of Starlight by Aligned Dust Grains,” *Ap. J.* **114** (Sept., 1951) 206.
- [458] S. Chandrasekhar and E. Fermi, “Magnetic Fields in Spiral Arms,” *Ap. J.* **118** (July, 1953) 113.

- [459] R. H. Hildebrand, L. Kirby, J. L. Dotson, M. Houde, and J. E. Vaillancourt, “Dispersion of Magnetic Fields in Molecular Clouds. I,” *Ap. J.* **696** (May, 2009) 567–573.
- [460] M. Houde, J. E. Vaillancourt, R. H. Hildebrand, S. Chitsazzadeh, and L. Kirby, “Dispersion of Magnetic Fields in Molecular Clouds. II,” *Ap. J.* **706** (Dec., 2009) 1504–1516.
- [461] L. M. Fissel, P. A. R. Ade, F. E. Angilè, P. Ashton, S. J. Benton, M. J. Devlin, B. Dober, Y. Fukui, N. Galitzki, N. N. Gandilo, J. Klein, A. L. Korotkov, Z.-Y. Li, P. G. Martin, T. G. Matthews, L. Moncelsi, F. Nakamura, C. B. Netterfield, G. Novak, E. Pascale, F. Poidevin, F. P. Santos, G. Savini, D. Scott, J. A. Shariff, J. Diego Soler, N. E. Thomas, C. E. Tucker, G. S. Tucker, and D. Ward-Thompson, “Balloon-Borne Submillimeter Polarimetry of the Vela C Molecular Cloud: Systematic Dependence of Polarization Fraction on Column Density and Local Polarization-Angle Dispersion,” *Ap. J.* **824** (June, 2016) 134, [arXiv:1509.05298](#).
- [462] Planck Collaboration Int. XXXV, “Planck intermediate results. XXXV. Probing the role of the magnetic field in the formation of structure in molecular clouds,” *A & A* **586** (2016) A138, [arXiv:1502.04123](#).
- [463] P. K. King, L. M. Fissel, C.-Y. Chen, and Z.-Y. Li, “Modelling dust polarization observations of molecular clouds through MHD simulations,” *Mon. Not. Roy. Astron. Soc.* **474** (Mar., 2018) 5122–5142, [arXiv:1709.08641](#).
- [464] G. Vedrenne and J.-L. Atteia, *Gamma-Ray Bursts*. 2009.
- [465] G. Ghirlanda, D. Burlon, G. Ghisellini, R. Salvaterra, M. G. Bernardini, S. Campana, S. Covino, P. D’Avanzo, V. D’Elia, A. Melandri, T. Murphy, L. Nava, S. D. Vergani, and G. Tagliaferri, “GRB Orphan Afterglows in Present and Future Radio Transient Surveys,” *PASA* **31** (May, 2014) e022, [arXiv:1402.6338](#) [[astro-ph.HE](#)].
- [466] G. P. Lamb, M. Tanaka, and S. Kobayashi, “Transient survey rates for orphan afterglows from compact merger jets,” *Mon. Not. Roy. Astron. Soc.* **476** (June, 2018) 4435–4441, [arXiv:1712.00418](#) [[astro-ph.HE](#)].
- [467] G. Ghirlanda, G. Ghisellini, R. Salvaterra, L. Nava, D. Burlon, G. Tagliaferri, S. Campana, P. D’Avanzo, and A. Melandri, “The faster the narrower: characteristic bulk velocities and jet opening angles of gamma-ray bursts,” *Mon. Not. Roy. Astron. Soc.* **428** (Jan., 2013) 1410–1423, [arXiv:1210.1215](#) [[astro-ph.HE](#)].
- [468] B. D. Metzger, P. K. G. Williams, and E. Berger, “Extragalactic Synchrotron Transients in the Era of Wide-field Radio Surveys. I. Detection Rates and Light Curve Characteristics,” *Ap. J.* **806** (June, 2015) 224, [arXiv:1502.01350](#) [[astro-ph.HE](#)].
- [469] D. Macpherson and D. Coward, “Multiwavelength detectability of Pop III GRBs from afterglow simulations,” *Mon. Not. Roy. Astron. Soc.* **467** (May, 2017) 2476–2493.
- [470] E. F. Keane and E. Petroff, “Fast radio bursts: search sensitivities and completeness,” *Mon. Not. Roy. Astron. Soc.* **447** (Mar., 2015) 2852–2856, [arXiv:1409.6125](#) [[astro-ph.HE](#)].
- [471] P. C. Boyle and Chime/Frb Collaboration, “First detection of fast radio bursts between 400 and 800 MHz by CHIME/FRB,” *The Astronomer’s Telegram* **11901** (Aug., 2018) .
- [472] V. Ravi, “The observed properties of fast radio bursts,” *Mon. Not. Roy. Astron. Soc.* **482** (Jan., 2019) 1966–1978, [arXiv:1710.08026](#) [[astro-ph.HE](#)].

- [473] D. R. Lorimer, M. Bailes, M. A. McLaughlin, D. J. Narkevic, and F. Crawford, “A Bright Millisecond Radio Burst of Extragalactic Origin,” *Science* **318** (Nov., 2007) 777, [arXiv:0709.4301](#).
- [474] V. Ravi, R. M. Shannon, M. Bailes, K. Bannister, S. Bhandari, N. D. R. Bhat, S. Burke-Spolaor, M. Caleb, C. Flynn, A. Jameson, S. Johnston, E. F. Keane, M. Kerr, C. Tiburzi, A. V. Tuntsov, and H. K. Vedantham, “The magnetic field and turbulence of the cosmic web measured using a brilliant fast radio burst,” *Science* **354** (Dec., 2016) 1249–1252, [arXiv:1611.05758](#) [astro-ph.HE].
- [475] L. G. Spitler, P. Scholz, J. W. T. Hessels, S. Bogdanov, A. Brazier, F. Camilo, S. Chatterjee, J. M. Cordes, F. Crawford, J. Deneva, R. D. Ferdman, P. C. C. Freire, V. M. Kaspi, P. Lazarus, R. Lynch, E. C. Madsen, M. A. McLaughlin, C. Patel, S. M. Ransom, A. Seymour, I. H. Stairs, B. W. Stappers, J. van Leeuwen, and W. W. Zhu, “A repeating fast radio burst,” *Nat.* **531** (Mar., 2016) 202–205, [arXiv:1603.00581](#) [astro-ph.HE].
- [476] CHIME/FRB Collaboration, “A second source of repeating fast radio bursts,” *Nat.* **566** (Jan., 2019) 235–238, [arXiv:1901.04525](#) [astro-ph.HE].
- [477] P. J. Armitage, “Physical processes in protoplanetary disks,” *ArXiv e-prints* (Sept., 2015) , [arXiv:1509.06382](#) [astro-ph.SR].
- [478] D. Johnstone, G. J. Herczeg, S. Mairs, J. Hatchell, G. C. Bower, H. Kirk, J. Lane, G. S. Bell, S. Graves, Y. Aikawa, H.-R. V. Chen, W.-P. Chen, M. Kang, S.-J. Kang, J.-E. Lee, O. Morata, A. Pon, P. Scicluna, A. Scholz, S. Takahashi, H. Yoo, and The JCMT Transient Team, “The JCMT Transient Survey: Stochastic and Secular Variability of Protostars and Disks In the Submillimeter Region Observed over 18 Months,” *Ap. J.* **854** (Feb., 2018) 31, [arXiv:1801.03537](#) [astro-ph.SR].
- [479] S. Mairs, D. Johnstone, H. Kirk, J. Lane, G. S. Bell, S. Graves, G. J. Herczeg, P. Scicluna, G. C. Bower, H.-R. V. Chen, J. Hatchell, Y. Aikawa, W.-P. Chen, M. Kang, S.-J. Kang, J.-E. Lee, O. Morata, A. Pon, A. Scholz, S. Takahashi, H. Yoo, and The JCMT Transient Team, “The JCMT Transient Survey: Identifying Submillimeter Continuum Variability over Several Year Timescales Using Archival JCMT Gould Belt Survey Observations,” *Ap. J.* **849** (Nov., 2017) 107, [arXiv:1710.03761](#) [astro-ph.SR].
- [480] J. Forbrich, M. J. Reid, K. M. Menten, V. M. Rivilla, S. J. Wolk, U. Rau, and C. J. Chandler, “Extreme Radio Flares and Associated X-Ray Variability from Young Stellar Objects in the Orion Nebula Cluster,” *Ap. J.* **844** (Aug., 2017) 109, [arXiv:1706.05562](#) [astro-ph.HE].
- [481] R. J. Ivison, D. H. Hughes, H. M. Lloyd, M. K. Bang, and M. F. Bode, “Millimetre and Submillimetre Continuum Observations of Nova Cygni 1992 - a New Test of Mass Ejection Models,” *Mon. Not. Roy. Astron. Soc.* **263** (Aug., 1993) L43.
- [482] L. Chomiuk, J. D. Linford, J. Yang, T. J. O’Brien, Z. Paragi, A. J. Mioduszewski, R. J. Beswick, C. C. Cheung, K. Mukai, T. Nelson, V. A. R. M. Ribeiro, M. P. Rupen, J. L. Sokoloski, J. Weston, Y. Zheng, M. F. Bode, S. Eyres, N. Roy, and G. B. Taylor, “Binary orbits as the driver of γ -ray emission and mass ejection in classical novae,” *Nat.* **514** (Oct., 2014) 339–342, [arXiv:1410.3473](#) [astro-ph.HE].
- [483] J. M. Paredes, J. Martí, M. Peracaula, G. Pooley, and I. F. Mirabel, “Exploring the high frequency emission of radio loud X-ray binaries,” *A & A* **357** (May, 2000) 507–514, [arXiv:astro-ph/0005517](#) [astro-ph].
- [484] J. R. Masiero, E. Redwing, A. K. Mainzer, J. M. Bauer, R. M. Cutri, T. Grav, E. Kramer, C. R. Nugent, S. Sonnett, and E. L. Wright, “Small and Nearby NEOs Observed by NEOWISE During the First Three Years of Survey: Physical Properties,” *Astron. J.* **156** (Aug., 2018) 60, [arXiv:1806.04759](#) [astro-ph.EP].

- [485] T. Müller, A. Marciniak, C. Kiss, R. Duffard, V. Ali-Lagoa, P. Bartczak, M. Butkiewicz-Bak, G. Dudzinski, E. Fernandez-Valenzuela, G. Marton, N. Morales, J.-L. Ortiz, D. Oszkiewicz, T. Santana-Ros, P. Santos-Sanz, R. Szakats, A. Takacsne Farkas, and E. Varga-Verebelyi, “Small Bodies Near and Far (SBNAF): Characterization of asteroids and TNOs,” *European Planetary Science Congress* **11** (Sept., 2017) EPSC2017–474.
- [486] K. Batygin and M. E. Brown, “Evidence for a Distant Giant Planet in the Solar System,” *Astron. J.* **151** (Feb., 2016) 22, [arXiv:1601.05438](#) [astro-ph.EP].
- [487] IceCube Collaboration, M. G. Aartsen, M. Ackermann, J. Adams, J. A. Aguilar, M. Ahlers, M. Ahrens, I. Al Samarai, D. Altmann, K. Andeen, and et al., “Multimessenger observations of a flaring blazar coincident with high-energy neutrino IceCube-170922A,” *Science* **361** (July, 2018) eaat1378, [arXiv:1807.08816](#) [astro-ph.HE].
- [488] BICEP 2, Keck Array and Planck Collaborations, “Joint Analysis of BICEP2/Keck Array and Planck Data,” *Phys. Rev. Lett.* **114** (2015) 101301, [arXiv:1502.00612](#) [astro-ph.CO].
- [489] M. Kamionkowski, A. Kosowsky, and A. Stebbins, “A Probe of primordial gravity waves and vorticity,” *Phys.Rev.Lett.* **78** (1997) 2058–2061, [arXiv:astro-ph/9609132](#) [astro-ph].
- [490] U. Seljak and M. Zaldarriaga, “Signature of gravity waves in polarization of the microwave background,” *Phys.Rev.Lett.* **78** (1997) 2054–2057, [arXiv:astro-ph/9609169](#) [astro-ph].
- [491] M. Zaldarriaga, D. N. Spergel, and U. Seljak, “Microwave background constraints on cosmological parameters,” *Astrophys. J.* **488** (1997) 1–13, [arXiv:astro-ph/9702157](#) [astro-ph].
- [492] L. Knox and Y.-S. Song, “A Limit on the detectability of the energy scale of inflation,” *Phys. Rev. Lett.* **89** (2002) 011303, [arXiv:astro-ph/0202286](#) [astro-ph].
- [493] M. Kesden, A. Cooray, and M. Kamionkowski, “Separation of gravitational wave and cosmic shear contributions to cosmic microwave background polarization,” *Phys. Rev. Lett.* **89** (2002) 011304, [arXiv:astro-ph/0202434](#) [astro-ph].
- [494] A. de Ugarte Postigo, A. Lundgren, S. Martín, D. Garcia-Appadoo, I. de Gregorio Monsalvo, A. Peck, M. J. Michałowski, C. C. Thöne, S. Campana, J. Gorosabel, N. R. Tanvir, K. Wiersema, A. J. Castro-Tirado, S. Schulze, C. De Breuck, G. Petitpas, J. Hjorth, P. Jakobsson, S. Covino, J. P. U. Fynbo, J. M. Winters, M. Bremer, A. J. Levan, A. Llorente, R. Sánchez-Ramírez, J. C. Tello, and R. Salvaterra, “Pre-ALMA observations of GRBs in the mm/submm range,” *A & A* **538** (Feb, 2012) A44, [arXiv:1108.1797](#) [astro-ph.CO].
- [495] T. Laskar, K. D. Alexander, E. Berger, C. Guidorzi, R. Margutti, W.-f. Fong, C. D. Kilpatrick, P. Milne, M. R. Drout, C. G. Mundell, S. Kobayashi, R. Lunnan, R. Barniol Duran, K. M. Menten, K. Ioka, and P. K. G. Williams, “First ALMA Light Curve Constrains Refreshed Reverse Shocks and Jet Magnetization in GRB 161219B,” *Ap. J.* **862** no. 2, (Aug, 2018) 94, [arXiv:1808.09476](#) [astro-ph.HE].
- [496] R. W. Ogburn, P. A. R. Ade, R. W. Aikin, M. Amiri, S. J. Benton, C. A. Bischoff, J. J. Bock, J. A. Bonetti, J. A. Brevik, E. Bullock, B. Burger, G. Davis, C. D. Dowell, L. Duband, J. P. Filippini, S. Fliescher, S. R. Golwala, M. Gordon, M. Halpern, M. Hasselfield, G. Hilton, V. V. Hristov, H. Hui, K. Irwin, J. P. Kaufman, B. G. Keating, S. A. Kernasovskiy, J. M. Kovac, C. L. Kuo, E. M. Leitch, M. Lueker, T. Montroy, C. B. Netterfield, H. T. Nguyen, R. O’Brien, A. Orlando, C. L. Pryke, C. Reintsema, S. Richter, J. E. Ruhl, M. C. Runyan, R. Schwarz, C. D. Sheehy, Z. K. Staniszewski, R. V. Sudiwala, G. P. Teply, K. Thompson, J. E. Tolan, A. D. Turner, A. G. Vieregg,

- D. V. Wiebe, P. Wilson, and C. L. Wong, “BICEP2 and Keck array operational overview and status of observations,” in *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy VI*, vol. 8452 of *Proc. SPIE*, p. 84521A. Sept., 2012. arXiv:1208.0638 [astro-ph.IM].
- [497] J. E. Carlstrom, P. A. R. Ade, K. A. Aird, B. A. Benson, L. E. Bleem, S. Buseti, C. L. Chang, E. Chauvin, H.-M. Cho, T. M. Crawford, A. T. Crites, M. A. Dobbs, N. W. Halverson, S. Heimsath, W. L. Holzapfel, J. D. Hrubes, M. Joy, R. Keisler, T. M. Lanting, A. T. Lee, E. M. Leitch, J. Leong, W. Lu, M. Lueker, D. Luong-Van, J. J. McMahon, J. Mehl, S. S. Meyer, J. J. Mohr, T. E. Montroy, S. Padin, T. Plagge, C. Pryke, J. E. Ruhl, K. K. Schaffer, D. Schwan, E. Shirokoff, H. G. Spieler, Z. Staniszewski, A. A. Stark, C. Tucker, K. Vanderlinde, J. D. Vieira, and R. Williamson, “The 10 Meter South Pole Telescope,” *PASP* **123** (May, 2011) 568, arXiv:0907.4445 [astro-ph.IM].
- [498] J. M. Fowler *et al.*, “Optical Design of the Atacama Cosmology Telescope and the Millimeter Bolometric Array Camera,” *Appl. Opt.* **46** (2007) 3444–3454, arXiv:astro-ph/0701020 [astro-ph].
- [499] D. Barron, P. Ade, A. Anthony, K. Arnold, D. Boettger, J. Borrill, S. Chapman, Y. Chinone, M. Dobbs, J. Edwards, J. Errard, G. Fabbian, D. Flanigan, G. Fuller, A. Ghribi, W. Grainger, N. Halverson, M. Hasegawa, K. Hattori, M. Hazumi, W. Holzapfel, J. Howard, P. Hyland, G. Jaehnig, A. Jaffe, B. Keating, Z. Kermish, R. Keskitalo, T. Kisner, A. T. Lee, M. Le Jeune, E. Linder, M. Lungu, F. Matsuda, T. Matsumura, X. Meng, N. J. Miller, H. Morii, S. Moyerman, M. Myers, H. Nishino, H. Paar, J. Peloton, E. Quealy, G. Rebeiz, C. L. Reichardt, P. L. Richards, C. Ross, A. Shimizu, C. Shimmmin, M. Shimon, M. Sholl, P. Siritanasak, H. Spieler, N. Stebor, B. Steinbach, R. Stompor, A. Suzuki, T. Tomaru, C. Tucker, A. Yadav, and O. Zahn, “The POLARBEAR Cosmic Microwave Background Polarization Experiment,” *Journal of Low Temperature Physics* **176** (Sept., 2014) 726–732.
- [500] C.-L. Kuo, “Assessments of Ali, Dome A, and Summit Camp for mm-wave Observations Using MERRA-2 Reanalysis,” *Ap. J.* **848** (Oct., 2017) 64, arXiv:1707.08400 [astro-ph.IM].
- [501] A. P. Lane, “Submillimeter Transmission at South Pole,” in *Astrophysics From Antarctica*, G. Novak and R. Landsberg, eds., vol. 141 of *Astronomical Society of the Pacific Conference Series*, pp. 289–295. 1998.
- [502] D. A. Erasmus, “An Analysis of CLOUD Cover and Water Vapor for the ALMA Project,” *A Report to European Southern Observatories* (Dec, 2002). https://www.eso.org/gen-fac/pubs/astclim/espas/radioseeing/pdf/ALMAFnlRep_Erasmus_Dec2002.pdf.
- [503] J. W. Fowler, M. D. Niemack, S. R. Dicker, A. M. Aboobaker, P. A. R. Ade, E. S. Battistelli, M. J. Devlin, R. P. Fisher, M. Halpern, P. C. Hargrave, A. D. Hincks, M. Kaul, J. Klein, J. M. Lau, M. Limon, T. A. Marriage, P. D. Mauskopf, L. Page, S. T. Staggs, D. S. Swetz, E. R. Switzer, R. J. Thornton, and C. E. Tucker, “Optical design of the atacama cosmology telescope and the millimeter bolometric array camera,” *Appl. Opt.* **46** no. 17, (Jun, 2007) 3444–3454. <http://ao.osa.org/abstract.cfm?URI=ao-46-17-3444>.
- [504] K. Arnold, N. Stebor, P. A. R. Ade, Y. Akiba, A. E. Anthony, M. Atlas, D. Barron, A. Bender, D. Boettger, J. Borrill, S. Chapman, Y. Chinone, A. Cukierman, M. Dobbs, T. Elleflot, J. Errard, G. Fabbian, C. Feng, A. Gilbert, N. Goeckner-Wald, N. W. Halverson, M. Hasegawa, K. Hattori, M. Hazumi, W. L. Holzapfel, Y. Hori, Y. Inoue, G. C. Jaehnig, A. H. Jaffe, N. Katayama, B. Keating, Z. Kermish, R. Keskitalo, T. Kisner, M. L. Jeune, A. T. Lee, E. M. Leitch, E. Linder, F. Matsuda, T. Matsumura, X. Meng, N. J. Miller, H. Morii, M. J. Myers, M. Navaroli, H. Nishino, T. Okamura, H. Paar, J. Peloton, D. Poletti, C. Raum, G. Rebeiz, C. L. Reichardt, P. L. Richards, C. Ross, K. M. Rotermund, D. E. Schenck, B. D. Sherwin, I. Shirley, M. Sholl, P. Siritanasak, G. Smecher,

- B. Steinbach, R. Stompor, A. Suzuki, J. Suzuki, S. Takada, S. Takakura, T. Tomaru, B. Wilson, A. Yadav, and O. Zahn, “The simons array: expanding polarbear to three multi-chroic telescopes,” vol. 9153, p. 91531F. Aug., 2014. <https://doi.org/10.1117/12.2057332>.
- [505] K. Harington, T. Marriage, A. Ali, C. L. Appel, John W. and Bennett, F. Boone, M. Brewer, M. Chan, D. T. Chuss, F. Colazo, S. Dahal, K. Denis, R. Dunner, J. Eimer, T. Essinger-Hileman, P. Fluxa, M. Halpern, G. Hilton, G. F. Hinshaw, J. Hubmayr, J. Iuliano, J. Karakla, J. McMahon, N. T. Miller, S. H. Miseley, G. Palma, L. Parker, M. Petroff, B. Pradenas, K. Rostern, M. Sagliocca, D. Valle, D. Watts, E. Wollack, Z. Xu, and L. Zeng, “The Cosmology Large Angular Scale Surveyor,” vol. 9914. 2016. <https://doi.org/10.1117/12.2233125>.
- [506] W. Hu, M. M. Hedman, and M. Zaldarriaga, “Benchmark parameters for CMB polarization experiments,” *Phys.Rev.* **D67** (2003) 043004, [arXiv:astro-ph/0210096](https://arxiv.org/abs/astro-ph/0210096) [astro-ph].
- [507] S. C. Parshley, M. Niemack, R. Hills, *et al.*, “The optical design of the six-meter CCAT-prime and Simons Observatory telescopes,” in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, vol. 10700 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, p. 1070041. July, 2018. [arXiv:1807.06678](https://arxiv.org/abs/1807.06678) [astro-ph.IM].
- [508] N. Galitzki *et al.*, “The Simons Observatory: Instrument Overview,” *Proc. SPIE Int. Soc. Opt. Eng.* **10708** (2018) 1070804, [arXiv:1808.04493](https://arxiv.org/abs/1808.04493) [astro-ph.IM].
- [509] C. Dragone, “Offset multireflector antennas with perfect pattern symmetry and polarization discrimination,” *AT T Technical Journal* **57** (Sept., 1978) 2663–2684.
- [510] C. Bischoff, A. Brizius, I. Buder, Y. Chinone, K. Cleary, R. N. Dumoulin, A. Kusaka, R. Monsalve, S. K. Naess, L. B. Newburgh, G. Nixon, R. Reeves, K. M. Smith, K. Vanderlinde, I. K. Wehus, M. Bogdan, R. Bustos, S. E. Church, R. Davis, C. Dickinson, H. K. Eriksen, T. Gaier, J. O. Gundersen, M. Hasegawa, M. Hazumi, C. Holler, K. M. Huffenberger, W. A. Imbriale, K. Ishidoshiro, M. E. Jones, P. Kangaslahti, D. J. Kapner, C. R. Lawrence, E. M. Leitch, M. Limon, J. J. McMahon, A. D. Miller, M. Nagai, H. Nguyen, T. J. Pearson, L. Piccirillo, S. J. E. Radford, A. C. S. Readhead, J. L. Richards, D. Samtleben, M. Seiffert, M. C. Shepherd, S. T. Staggs, O. Tajima, K. L. Thompson, R. Williamson, B. Winstein, E. J. Wollack, and J. T. L. Zwart, “The Q/U Imaging Experiment Instrument,” *Ap. J.* **768** (May, 2013) 9, [arXiv:1207.5562](https://arxiv.org/abs/1207.5562) [astro-ph.IM].
- [511] A. Kusaka, J. Appel, T. Essinger-Hileman, J. A. Beall, L. E. Campusano, H.-M. Cho, S. K. Choi, K. Crowley, J. W. Fowler, P. Gallardo, M. Hasselfield, G. Hilton, S.-P. P. Ho, K. Irwin, N. Jarosik, M. D. Niemack, G. W. Nixon, M. Nolta, J. Page, Lyman A., G. A. Palma, L. Parker, S. Raghunathan, C. D. Reintsema, J. Sievers, S. M. Simon, S. T. Staggs, K. Visnjic, and K.-W. Yoon, “Results from the Atacama B-mode Search (ABS) experiment,” *Journal of Cosmology and Astro-Particle Physics* **2018** (Sept., 2018) 005.
- [512] M. D. Niemack, “Designs for a large-aperture telescope to map the CMB 10X faster,” *Appl. Opt.* **55** (2016) 1688–96, [arXiv:1511.04506](https://arxiv.org/abs/1511.04506) [astro-ph.IM].
- [513] C. Dragone, “First-order correction of aberrations in Cassegrainian and Gregorian antennas,” *IEEE Transactions on Antennas and Propagation* **31** (Sept., 1983) 764–775.
- [514] S. C. Parshley, J. Kronshage, *et al.*, “CCAT-prime: a novel telescope for sub-millimeter astronomy,” in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, vol. 10700 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, p. 107005X. July, 2018. [arXiv:1807.06675](https://arxiv.org/abs/1807.06675) [astro-ph.IM].

- [515] H. Tran, A. Lee, S. Hanany, M. Milligan, and T. Renbarger, “Comparison of the crossed and the Gregorian Mizuguchi-Dragone for wide-field millimeter-wave astronomy,” *Applied Optics* **47** (Jan., 2008) 103–109.
- [516] P. A. Gallardo *et al.*, “Studies of Systematic Uncertainties for Simons Observatory: Optical Effects and Sensitivity Considerations,” *Proc. SPIE Int. Soc. Opt. Eng.* **10708** (2018) 107083Y, [arXiv:1808.05152](#) [astro-ph.IM].
- [517] S. R. Dicker *et al.*, “Cold optical design for the large aperture Simons’ Observatory telescope,” *Proc. SPIE Int. Soc. Opt. Eng.* **10700** (2018) 107003E, [arXiv:1808.05058](#) [astro-ph.IM].
- [518] C. A. Hill *et al.*, “BoloCalc: a sensitivity calculator for the design of Simons Observatory,” *Proc. SPIE Int. Soc. Opt. Eng.* **10708** (2018) 1070842, [arXiv:1806.04316](#) [astro-ph.IM]. [Astronomy9,1070842(2018)].
- [519] P. A. Gallardo, N. F. Cothard, R. Puddu, R. Dunner, B. J. Koopman, M. D. Niemack, S. M. Simon, and E. J. Wollack, “Far Sidelobes from Baffles and Telescope Support Structures in the Atacama Cosmology Telescope,” 2018. [arXiv:1808.05101](#) [astro-ph.IM].
- [520] P. A. F. Rojas *et al.*, “Far sidelobe effects from panel gaps of the Atacama Cosmology Telescope,” *Proc. SPIE Int. Soc. Opt. Eng.* **9914** (2016) 99142Q.
- [521] J. G. Mangum, J. W. M. Baars, A. Greve, R. Lucas, R. C. Snel, P. Wallace, and M. Holdaway, “Evaluation of the ALMA Prototype Antennas,” *Publ. Astron. Soc. Pac.* **118** (2006) 1257, [arXiv:astro-ph/0609329](#) [astro-ph].
- [522] B. Nikolic, R. E. Hills, and J. S. Richer, “Measurement of antenna surfaces from in- and out-of-focus beam maps using astronomical sources,” *A & A* **465** (Apr., 2007) 679–683, [astro-ph/0612241](#).
- [523] N. Zhu, J. L. Orlowski-Scherer, Z. Xu, A. Ali, K. S. Arnold, P. C. Ashton, G. Coppi, M. J. Devlin, S. Dicker, N. Galitzki, P. A. Gallardo, S. W. Henderson, S.-P. P. Ho, J. Hubmayr, B. Keating, A. T. Lee, M. Limon, M. Lungu, P. D. Mauskopf, A. J. May, J. McMahon, M. D. Niemack, L. Piccirillo, G. Puglisi, M. Sathyanarayana Rao, M. Salatino, M. Silva-Feaver, S. M. Simon, S. Staggs, R. Thornton, J. N. Ullom, E. M. Vavagiakis, B. Westbrook, and E. J. Wollack, “Simons Observatory large aperture telescope receiver design overview,” in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, vol. 10708 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, p. 1070829. July, 2018. [arXiv:1808.10037](#) [astro-ph.IM].
- [524] K. N. Shukla, “Heat Pipe for Aerospace Applications—An Overview,” *Journal of Electronics Cooling and Thermal Control* **5** (2015) 1–14.
- [525] G. Coppi *et al.*, “Cooldown strategies and transient thermal simulations for the Simons Observatory,” *Proc. SPIE Int. Soc. Opt. Eng.* **10708** (2018) 1070827, [arXiv:1808.07896](#) [astro-ph.IM].
- [526] J. L. Orlowski-Scherer, N. Zhu, Z. Xu, A. Ali, K. S. Arnold, P. C. Ashton, G. Coppi, M. Devlin, S. Dicker, N. Galitzki, P. A. Gallardo, B. Keating, A. T. Lee, M. Limon, M. Lungu, A. May, J. McMahon, M. D. Niemack, L. Piccirillo, G. Puglisi, M. Salatino, M. Silva-Feaver, S. M. Simon, R. Thornton, and E. M. Vavagiakis, “Simons Observatory large aperture receiver simulation overview,” in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, vol. 10708 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, p. 107083X. July, 2018. [arXiv:1808.06648](#) [astro-ph.IM].

- [527] M. Crumrine, P. A. R. Ade, Z. Ahmed, R. W. Aikin, K. D. Alexander, D. Barkats, S. J. Benton, C. A. Bischoff, J. J. Bock, R. Bowens-Rubin, J. A. Brevik, I. Buder, E. Bullock, V. Buza, J. Connors, J. Cornelison, B. P. Crill, M. Dierickx, L. Duband, C. Dvorkin, J. P. Filippini, S. Fliescher, J. A. Grayson, G. Hall, M. Halpern, S. A. Harrison, S. R. Hildebrandt, G. C. Hilton, H. Hui, K. D. Irwin, J. H. Kang, K. S. Karkare, E. Karpel, J. P. Kaufman, B. G. Keating, S. Kefeli, S. A. Kernasovskiy, J. M. Kovac, C. L. Kuo, N. A. Larsen, K. Lau, E. M. Leitch, M. V. Lueker, K. G. Megerian, L. Moncelsi, T. Namikawa, C. B. Netterfield, H. T. Nguyen, R. O’Brien, R. W. Ogburn, S. Palladino, C. Pryke, B. Racine, S. Richter, R. Schwarz, A. Schillaci, C. D. Sheehy, A. Soliman, T. St. Germaine, Z. K. Staniszewski, B. Steinbach, R. V. Sudiwala, G. P. Teply, K. L. Thompson, J. E. Tolán, C. E. Tucker, A. D. Turner, C. Umiltà, A. G. Viereg, A. Wandui, A. C. Weber, D. V. Wiebe, J. Willmert, C. L. Wong, W. L. K. Wu, E. Yang, K. W. Yoon, and C. Zhang, “BICEP Array cryostat and mount design,” in *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy IX*, vol. 10708 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, p. 107082D. July, 2018. [arXiv:1808.00569](https://arxiv.org/abs/1808.00569) [astro-ph.IM].
- [528] BICEP2 Collaboration, P. A. R. Ade *et al.*, “BICEP2 III: Instrumental Systematics,” *Astrophys. J.* **814** no. 2, (2015) 110, [arXiv:1502.00608](https://arxiv.org/abs/1502.00608) [astro-ph.IM].
- [529] K. Rostem, A. Ali, J. W. Appel, C. L. Bennett, A. Brown, M.-P. Chang, D. T. Chuss, F. A. Colazo, N. Costen, K. L. Denis, T. Essinger-Hileman, R. Hu, T. A. Marriage, S. H. Moseley, T. R. Stevenson, K. U-Yen, E. J. Wollack, and Z. Xu, “Silicon-based antenna-coupled polarization-sensitive millimeter-wave bolometer arrays for cosmic microwave background instruments,” in *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy VIII*, vol. 9914 of *Proc. SPIE*, p. 99140D. July, 2016. [arXiv:1608.08891](https://arxiv.org/abs/1608.08891) [astro-ph.IM].
- [530] S. M. Simon, J. Austermann, J. A. Beall, S. K. Choi, K. P. Coughlin, S. M. Duff, P. A. Gallardo, S. W. Henderson, F. B. Hills, S.-P. P. Ho, J. Hubmayr, A. Josaitis, B. J. Koopman, J. J. McMahon, F. Nati, L. Newburgh, M. D. Niemack, M. Salatino, A. Schillaci, B. L. Schmitt, S. T. Staggs, E. M. Vavagiakis, J. Ward, and E. J. Wollack, “The design and characterization of wideband spline-profiled feedhorns for Advanced ACTPol,” in *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy VIII*, vol. 9914 of *Proc. SPIE*, p. 991416. July, 2016.
- [531] G. Yassin, P. Kittara, A. Jiralucksanawong, S. Wangsuya, J. Leech, and M. Jones, “A High Performance Horn for Large Format Focal Plane Arrays,” in *Eighteenth International Symposium on Space Terahertz Technology*, A. Karpov, ed., p. 199. 2007.
- [532] L. Zeng, C. L. Bennett, D. T. Chuss, and E. J. Wollack, “A Low Cross-Polarization Smooth-Walled Horn With Improved Bandwidth,” *IEEE Transactions on Antennas and Propagation* **58** (Apr., 2010) 1383–1387.
- [533] C. Granet, G. L. James, R. Bolton, and G. Moorey, “A Smooth-Walled Spline-Profile Horn as an Alternative to the Corrugated Horn for Wide Band Millimeter-Wave Applications,” *IEEE Transactions on Antennas and Propagation* **52** (Mar., 2004) 848–854.
- [534] S. M. Simon, J. E. Golec, A. Ali, J. Austermann, J. A. Beall, S. M. M. Bruno, S. K. Choi, K. T. Crowley, S. Dicker, B. Dober, S. M. Duff, E. Healy, C. A. Hill, S.-P. P. Ho, J. Hubmayr, Y. Li, M. Lungu, J. M. an dJohn Orlowski-Scherer, M. Salatino, S. Staggs, , E. J. Wollack, Z. Xu, and N. Zhu, “Feedhorn development and scalability for simons observatory and beyond,” in *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy VIII*, vol. 10708 of *Proc. SPIE*, pp. 10708 – 10708 – 12. 2018. <https://doi.org/10.1117/12.2313405>.
- [535] E. J. Wollack, K. U-yen, and D. T. Chuss, “Photonic choke-joints for dual-polarization waveguides,”.

- [536] B. J. Dober, P. A. R. Ade, P. Ashton, F. E. Angilè, J. A. Beall, D. Becker, K. J. Bradford, G. Che, H.-M. Cho, M. J. Devlin, L. M. Fissel, Y. Fukui, N. Galitzki, J. Gao, C. E. Groppi, S. Hillbrand, G. C. Hilton, J. Hubmayr, K. D. Irwin, J. Klein, J. Van Lanen, D. Li, Z.-Y. Li, N. P. Lourie, H. Mani, P. G. Martin, P. Mauskopf, F. Nakamura, G. Novak, D. P. Pappas, E. Pascale, F. P. Santos, G. Savini, D. Scott, S. Stanchfield, J. N. Ullom, M. Underhill, M. R. Vissers, and D. Ward-Thompson, “The next-generation BLASTPol experiment,” in *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy VII*, vol. 9153 of *Proc.SPIE Int.Soc.Opt.Eng.*, p. 91530H. July, 2014. [arXiv:1407.3756](#) [astro-ph.IM].
- [537] **BICEP3** Collaboration, H. Hui *et al.*, “BICEP3 focal plane design and detector performance,” *Proc. SPIE Int. Soc. Opt. Eng.* **9914** (2016) 99140T, [arXiv:1607.06861](#) [astro-ph.IM].
- [538] W. J.T. *et al.*, “Mechanical design and development of tes bolometer detector arrays for the Advanced ACTPol experiment,” *Proc. SPIE* **9914** (2016) 991437. [astro-ph/1607.05754](#).
- [539] P. G. *et al.*, “High-Density Superconducting Cables for Advanced ACTPol,” *JLTP* **184** no. 1-2, (2016) 473–479.
- [540] G. K. Lockwood, D. Hazen, Q. Koziol, R. S. Canon, K. Antypas, J. Balewski, and *et al.*, “Storage 2020: A Vision for the Future of HPC Storage,” Tech. Rep. LBNL-2001072, Lawrence Berkeley National Laboratory, 2017. <https://escholarship.org/uc/item/744479dp>.
- [541] J. T. Ward, D. Alonso, J. Errard, M. J. Devlin, and M. Hasselfield, “The Effects of Bandpass Variations on Foreground Removal Forecasts for Future CMB Experiments,” *Astrophys. J.* **861** no. 2, (2018) 82, [arXiv:1803.07630](#) [astro-ph.IM].
- [542] F. Nati, M. J. Devlin, M. Gerbino, B. R. Johnson, B. Keating, L. Pagano, and G. Teply, “POLOCALC: A Novel Method to Measure the Absolute Polarization Orientation of the Cosmic Microwave Background,” *Journal of Astronomical Instrumentation* **6** no. 2, (Jun, 2017) 1740008, [arXiv:1704.02704](#) [astro-ph.IM].
- [543] W. Hu and T. Okamoto, “Mass reconstruction with cmb polarization,” *Astrophys.J.* **574** (2002) 566–574, [arXiv:astro-ph/0111606](#) [astro-ph].
- [544] T. Okamoto and W. Hu, “CMB lensing reconstruction on the full sky,” *Phys. Rev.* **D67** (2003) 083002, [arXiv:astro-ph/0301031](#) [astro-ph].
- [545] M. Millea, E. Anderes, B. D. Wandelt, and M. Millea, “Bayesian delensing of CMB temperature and polarization,” [arXiv:1708.06753](#) [astro-ph.CO].
- [546] J. Caldeira, W. L. K. Wu, B. Nord, C. Avestruz, S. Trivedi, and K. T. Story, “DeepCMB: Lensing Reconstruction of the Cosmic Microwave Background with Deep Neural Networks,” [arXiv:1810.01483](#) [astro-ph.CO].
- [547] M. S. Madhavacheril and J. C. Hill, “Mitigating Foreground Biases in CMB Lensing Reconstruction Using Cleaned Gradients,” *Phys. Rev.* **D98** no. 2, (2018) 023534, [arXiv:1802.08230](#) [astro-ph.CO].
- [548] A. Manzotti, “Future cosmic microwave background delensing with galaxy surveys,” *Phys. Rev.* **D97** no. 4, (2018) 043527, [arXiv:1710.11038](#) [astro-ph.CO].
- [549] E. L. Wright, “Scanning and mapping strategies for CMB experiments,” [arXiv:astro-ph/9612006](#) [astro-ph].
- [550] J. R. Bond, A. H. Jaffe, and L. Knox, “Estimating the power spectrum of the cosmic microwave background,” *Phys. Rev.* **D57** (1998) 2117–2137, [arXiv:astro-ph/9708203](#) [astro-ph].

- [551] J. Borrill, “The challenge of data analysis for future CMB observations,” *AIP Conf. Proc.* **476** (1999) 277–282, [arXiv:astro-ph/9903204](#) [astro-ph].
- [552] E. Hivon, K. M. Gorski, C. B. Netterfield, B. P. Crill, S. Prunet, and F. Hansen, “Master of the cosmic microwave background anisotropy power spectrum: a fast method for statistical analysis of large and complex cosmic microwave background data sets,” *Astrophys. J.* **567** (2002) 2, [arXiv:astro-ph/0105302](#) [astro-ph].
- [553] **BICEP2, Keck Array** Collaboration, P. A. R. Ade *et al.*, “BICEP2 / Keck Array X: Constraints on Primordial Gravitational Waves using Planck, WMAP, and New BICEP2/Keck Observations through the 2015 Season,” *Submitted to: Phys. Rev. Lett.* (2018) , [arXiv:1810.05216](#) [astro-ph.CO].
- [554] **BICEP2** Collaboration, P. A. R. Ade *et al.*, “Detection of *B*-Mode Polarization at Degree Angular Scales by BICEP2,” *Phys. Rev. Lett.* **112** no. 24, (2014) 241101, [arXiv:1403.3985](#) [astro-ph.CO].
- [555] **DES, SPT** Collaboration, T. M. C. Abbott *et al.*, “Dark Energy Survey Year 1 Results: Joint Analysis of Galaxy Clustering, Galaxy Lensing, and CMB Lensing Two-point Functions,” [arXiv:1810.02322](#) [astro-ph.CO].
- [556] **DES** Collaboration, E. J. Baxter *et al.*, “Dark Energy Survey Year 1 Results: Methodology and Projections for Joint Analysis of Galaxy Clustering, Galaxy Lensing, and CMB Lensing Two-point Functions,” [arXiv:1802.05257](#) [astro-ph.CO].
- [557] A. Lewis and S. Bridle, “Cosmological parameters from CMB and other data: a Monte- Carlo approach,” *Phys. Rev.* **D66** (2002) 103511, [astro-ph/0205436](#).
- [558] J.-B. Melin, J. G. Bartlett, and J. Delabrouille, “Catalog extraction in SZ cluster surveys: a matched filter approach,” *A & A* **459** (Nov., 2006) 341–352, [astro-ph/0602424](#).
- [559] Planck Collaboration XXIX, “Planck 2013 results. XXIX. The Planck catalogue of Sunyaev-Zeldovich sources,” *Astron. Astrophys.* **571** (2014) A29, [arXiv:1303.5089](#) [astro-ph.CO].
- [560] E. Pierpaoli, S. Anthoine, K. M. Huffenberger, and I. Daubechies, “Reconstructing Sunyaev-Zeldovich clusters in future CMB experiments,” *Mon. Not. Roy. Astron. Soc.* **359** (2005) 261–271, [arXiv:astro-ph/0412197](#) [astro-ph].
- [561] T. C.-S. C. et. al., “CMB-S4 Performance-Based Constraints on Primordial Gravitational Waves,” .
- [562] K. M. Smith, D. Hanson, M. LoVerde, C. M. Hirata, and O. Zahn, “Delensing CMB Polarization with External Datasets,” *JCAP* **1206** (2012) 014, [arXiv:1010.0048](#) [astro-ph.CO].
- [563] S. Paine, “The am atmospheric model,” Mar., 2018. <https://doi.org/10.5281/zenodo.1193646>.
- [564] B. Thorne, J. Dunkley, D. Alonso, and S. Naess, “The Python Sky Model: software for simulating the Galactic microwave sky,” [arXiv:1608.02841](#) [astro-ph.CO].
- [565] B. Hensley, *On the nature of interstellar grains*. PhD thesis, Princeton University, 2015.
- [566] T. Ghosh, F. Boulanger, P. G. Martin, A. Bracco, F. Vansyngel, J. Aumont, J. J. Bock, O. Doré, U. Haud, P. M. W. Kalberla, and P. Serra, “Modelling and simulation of large-scale polarized dust emission over the southern Galactic cap using the GASS Hi data,” *A & A* **601** (May, 2017) A71, [arXiv:1611.02418](#).
- [567] Planck Collaboration, “Planck intermediate results. L. Evidence of spatial variation of the polarized thermal dust spectral energy distribution and implications for CMB B-mode analysis,” *A & A* **599** (Mar., 2017) A51, [arXiv:1606.07335](#).

- [568] C. Sheehy and A. Slosar, “No evidence for dust B-mode decorrelation in Planck data,” *ArXiv e-prints* (Sept., 2017) , [arXiv:1709.09729](#).
- [569] A. G. Kritsuk, S. D. Ustyugov, and M. L. Norman, “The structure and statistics of interstellar turbulence,” *New Journal of Physics* **19** no. 6, (June, 2017) 065003, [arXiv:1705.01912](#).
- [570] H. K. Eriksen, A. J. Banday, K. M. Gorski, and P. B. Lilje, “Foreground removal by an internal linear combination method: Limitations and implications,” *Astrophys. J.* **612** (2004) 633–646, [arXiv:astro-ph/0403098](#) [astro-ph].
- [571] M. Tegmark and G. Efstathiou, “A method for subtracting foregrounds from multi-frequency cmb sky maps,” *Mon. Not. Roy. Astron. Soc.* **281** (1996) 1297, [arXiv:astro-ph/9507009](#) [astro-ph].
- [572] J. Dunkley, E. Calabrese, J. Sievers, G. E. Addison, N. Battaglia, E. S. Battistelli, J. R. Bond, S. Das, M. J. Devlin, R. Dünner, J. W. Fowler, M. Gralla, A. Hajian, M. Halpern, M. Hasselfield, A. D. Hincks, R. Hlozek, J. P. Hughes, K. D. Irwin, A. Kosowsky, T. Louis, T. A. Marriage, D. Marsden, F. Menanteau, K. Moodley, M. Niemack, M. R. Nolta, L. A. Page, B. Partridge, N. Sehgal, D. N. Spergel, S. T. Staggs, E. R. Switzer, H. Trac, and E. Wollack, “The Atacama Cosmology Telescope: likelihood for small-scale CMB data,” *JCAP* **7** (July, 2013) 025, [arXiv:1301.0776](#).
- [573] M. Béthermin, E. Daddi, G. Magdis, M. T. Sargent, Y. Hezaveh, D. Elbaz, D. Le Borgne, J. Mullaney, M. Pannella, V. Buat, V. Charmandaris, G. Lagache, and D. Scott, “A Unified Empirical Model for Infrared Galaxy Counts Based on the Observed Physical Evolution of Distant Galaxies,” *Ap. J.* **757** no. 2, (Oct, 2012) L23, [arXiv:1208.6512](#) [astro-ph.CO].
- [574] Planck Collaboration, P. A. R. Ade, N. Aghanim, C. Armitage-Caplan, M. Arnaud, M. Ashdown, F. Atrio-Barandela, J. Aumont, C. Baccigalupi, A. J. Banday, R. B. Barreiro, J. G. Bartlett, E. Battaner, K. Benabed, A. Benoît, A. Benoit-Lévy, J. P. Bernard, M. Bersanelli, M. Bethermin, P. Bielewicz, K. Blagrove, J. Bobin, J. J. Bock, A. Bonaldi, J. R. Bond, J. Borrill, F. R. Bouchet, F. Boulanger, M. Bridges, M. Bucher, C. Burigana, R. C. Butler, J. F. Cardoso, A. Catalano, A. Challinor, A. Chamballu, X. Chen, H. C. Chiang, L. Y. Chiang, P. R. Christensen, S. Church, D. L. Clements, S. Colombi, L. P. L. Colombo, F. Couchot, A. Coulais, B. P. Crill, A. Curto, F. Cuttaia, L. Danese, R. D. Davies, R. J. Davis, P. de Bernardis, A. de Rosa, G. de Zotti, J. Delabrouille, J. M. Delouis, F. X. Désert, C. Dickinson, J. M. Diego, H. Dole, S. Donzelli, O. Doré, M. Douspis, X. Dupac, G. Efstathiou, T. A. Enßlin, H. K. Eriksen, F. Finelli, O. Forni, M. Frailis, E. Franceschi, S. Galeotta, K. Ganga, T. Ghosh, M. Giard, Y. Giraud-Héraud, J. González-Nuevo, K. M. Górski, S. Gratton, A. Gregorio, A. Gruppuso, F. K. Hansen, D. Hanson, D. Harrison, G. Helou, S. Henrot-Versillé, C. Hernández-Monteagudo, D. Herranz, S. R. Hildebrandt, E. Hivon, M. Hobson, W. A. Holmes, A. Hornstrup, W. Hovest, K. M. Huffenberger, A. H. Jaffe, T. R. Jaffe, W. C. Jones, M. Juvela, P. Kalberla, E. Keihänen, J. Kerp, R. Keskitalo, T. S. Kisner, R. Kneissl, J. Knoche, L. Knox, M. Kunz, H. Kurki-Suonio, F. Lacasa, G. Lagache, A. Lähteenmäki, J. M. Lamarre, M. Langer, A. Lasenby, R. J. Laureijs, C. R. Lawrence, R. Leonardi, J. León-Tavares, J. Lesgourgues, M. Liguori, P. B. Lilje, M. Linden-Vørnle, M. López-Caniego, P. M. Lubin, J. F. Macías-Pérez, B. Maffei, D. Maino, N. Mandolese, M. Maris, D. J. Marshall, P. G. Martin, E. Martínez-González, S. Masi, M. Massardi, S. Matarrese, F. Matthai, P. Mazzotta, A. Melchiorri, L. Mendes, A. Mennella, M. Migliaccio, S. Mitra, M. A. Miville-Deschênes, A. Moneti, L. Montier, G. Morgante, D. Mortlock, D. Munshi, J. A. Murphy, P. Naselsky, F. Nati, P. Natoli, C. B. Netterfield, H. U. Nørgaard-Nielsen, F. Noviello, D. Novikov, I. Novikov, S. Osborne, C. A. Oxborrow, F. Paci, L. Pagano, F. Pajot, R. Paladini, D. Paoletti, B. Partridge, F. Pasian, G. Patanchon, O. Perdereau, L. Perotto, F. Perrotta, F. Piacentini, M. Piat, E. Pierpaoli, D. Pietrobon, S. Plaszczynski, E. Pointecouteau, G. Polenta, N. Ponthieu, L. Popa, T. Poutanen, G. W. Pratt, G. Prézeau, S. Prunet, J. L. Puget, J. P. Rachen, W. T. Reach, R. Rebolo, M. Reinecke, M. Remazeilles, C. Renault, S. Ricciardi, T. Riller, I. Ristorcelli, G. Rocha, C. Rosset,

- G. Roudier, M. Rowan-Robinson, J. A. Rubiño-Martín, B. Rusholme, M. Sandri, D. Santos, G. Savini, D. Scott, M. D. Seiffert, P. Serra, E. P. S. Shellard, L. D. Spencer, J. L. Starck, V. Stolyarov, R. Stompor, R. Sudiwala, R. Sunyaev, F. Sureau, D. Sutton, A. S. Suur-Uski, J. F. Sygnet, J. A. Tauber, D. Tavagnacco, L. Terenzi, L. Toffolatti, M. Tomasi, M. Tristram, M. Tucci, J. Tuovinen, M. Türler, L. Valenziano, J. Valiviita, B. Van Tent, P. Vielva, F. Villa, N. Vittorio, L. A. Wade, B. D. Wandelt, N. Welikala, M. White, S. D. M. White, B. Winkel, D. Yvon, A. Zacchei, and A. Zonca, “Planck 2013 results. XXX. Cosmic infrared background measurements and implications for star formation,” *A & A* **571** (Nov, 2014) A30, [arXiv:1309.0382](#) [[astro-ph.CO](#)].
- [575] C. M. Hirata and U. Seljak, “Reconstruction of lensing from the cosmic microwave background polarization,” *Phys. Rev.* **D68** (2003) 083002, [arXiv:astro-ph/0306354](#) [[astro-ph](#)].
- [576] D. Green, J. Meyers, and A. van Engelen, “CMB Delensing Beyond the B Modes,” [arXiv:1609.08143](#) [[astro-ph.CO](#)].
- [577] M. Arnaud, G. W. Pratt, R. Piffaretti, H. Böhringer, J. H. Croston, and E. Pointecouteau, “The universal galaxy cluster pressure profile from a representative sample of nearby systems (REXCESS) and the $Y_{SZ} - M_{500}$ relation,” *A & A* **517** (Jul, 2010) A92, [arXiv:0910.1234](#) [[astro-ph.CO](#)].