



Reconstructing high fidelity maps of the microwave sky from multi-kilo pixel experiments data for high-precision cosmology

Hamza El Bouhargani

► To cite this version:

Hamza El Bouhargani. Reconstructing high fidelity maps of the microwave sky from multi-kilo pixel experiments data for high-precision cosmology. Astrophysics [astro-ph]. Université Paris Cité, 2021. English. NNT : 2021UNIP7238 . tel-04011234

HAL Id: tel-04011234

<https://theses.hal.science/tel-04011234>

Submitted on 2 Mar 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Université de Paris
École doctorale 560 STEP'UP
Laboratoire ASTROPARTICULE & COSMOLOGIE

Reconstructing high-fidelity maps of the microwave sky from multi-kilo pixel experiments data for high-precision cosmology

Thèse de doctorat de **Physique de l'Univers**

par
Hamza EL BOUHARGANI
dirigée par
Dr. Radek STOMPOR

présentée et soutenue publiquement le
12 octobre 2021

Devant un jury composé de

Pr. Simona MEI (Université de Paris, APC)
Pr. Michael BROWN (University of Manchester)
Pr. Davide MAINO (Università degli Studi di Milano)
Pr. Joanna DUNKLEY (Princeton University)
Dr. Sophie HENROT-VERSILLÉ (CNRS, IJCLab)
Dr. Radek STOMPOR (APC)
Dr. Josquin ERRARD (APC)

Présidente du jury
Rapporteur
Rapporteur
Examinatrice
Examinatrice
Directeur de thèse
Co-encadrant de thèse

*"I am tormented with an everlasting itch for things remote. I love to sail forbidden seas,
and land on barbarous coasts."*

– Herman Melville, *Moby-Dick; or, The Whale*

Acknowledgments

Before all else, I would like to thank my two advisors Josquin Errard and Radek Stompor, not only for giving me the opportunity to work on this exciting field of research, but also for their constant support, advice and supervision throughout this thesis. I have learned so much from you in these three years even beyond the realm of science, and for that I am immensely grateful.

I also would like to thank the members of my thesis monitoring committee, Sophie Henrot-Versillé and Guillaume Patanchon, for their helpful comments and discussion of my work in our yearly meetings, as well as for their valuable advice and support.

My warmest thanks to Michael Brown and Davide Maino for agreeing to review my manuscript and for their very insightful comments which helped improve this thesis, and to all the remaining members of the jury: Simona Mei, Joanna Dunkley and Sophie Henrot-Versillé.

My special thanks go to fellow members of our group at APC and POLARBEAR Europe: Dominic Beck, Carlo Baccigalupi, Giulio Fabbian, Aygul Jamal, Baptiste Jost, Nicoletta Krachmalnicoff, Maude Le Jeune, Clément Leloup, Anto I. Lonappan, Magdy Morshed, Julien Peloton, Davide Poletti, Giuseppe Puglisi, Arianna Rizzieri and Clara Vergès. It is always a pleasure interacting with you!

I thank all members of the POLARBEAR and Simons Observatory collaborations. In particular, I would like to address my heartfelt thanks to my close collaborators at the Computational Cosmology Center at LBNL, Julian Borrill, Reijo Kesitalo and Ted Kisner for their invaluable help with the simulation and data analysis tools and for hosting me during my visits there. I also thank fellow collaborators who welcomed me at Berkeley: Shawn Beckman, Yuji Chinone, Kevin Crowley, Neil Goeckner-Wald, John Groh, Charlie Hill, Akito Kusaka, Adrian Lee, Aashrita Mangu, and Ben Westbrook. I thank Simone Aiola, Matthew Haselfield, Sigurd Naess, and fellow members of the Simons Observatory Pipeline Working Group for their insightful comments and discussion during our telecons.

Thanks a lot to my teammates of the ANR-B3DCMB project: Nicolas Chopin, Thibault Cimis, Gabriel Ducrocq, Julien Grain, Laura Grigori, Jan Papež and Matthieu Tristram. Our fruitful collaboration led to many of the results presented in this thesis.

Many thanks to my colleagues at APC, and to the administration for their consistent help during my PhD. I am also incredibly grateful for the great time I spent together with the fellow and former graduate students: Alexandre, Bastien, Calum, Dimitri, Edwan, Gabriel M., Guillaume, Hugo, Hoang, Jean-Baptiste, Jie, Julien M., Konstantin, Leandro, Louise, Makarim, Marc, Marion, Philippe, Pierre, Raphaël, Romain, Thibault V., Thomas and Tuan.

Thanks to Jean-Luc Atteia, Bertrand Cordier and Stéphane Schanne from whom I learned a lot during my first research experience with astrophysics, and who encouraged me to pursue a PhD.

My deepest appreciation to my first mentors who stirred up my interest for physics and astronomy: my high school physics teacher Mr. El Arib, and my father through nightlong discussions under the summer starry nights in our back-

yard.

Finally, I would like to express my deepest gratitude to my parents and my sisters. You have blessed me with your unconditional love and support and always believed in me, sparing no effort to help me pursue my ambitions. Thanks to you I am able to write these words today.

Abstract

In the last two decades, cosmology has been propelled by numerous theoretical and experimental developments allowing it to enter into the high precision era. New experiments are currently being prepared to start operating in the coming years and continue pushing the frontiers of knowledge and improve our theoretical understanding of the inner workings of the Universe. In particular, the next generation of cosmic microwave background (CMB) experiments will enable the mapping of the CMB temperature and polarization anisotropies with unprecedented precision, from the largest degree scale patterns on the sky down to arcminute scales. They will hence open a unique window on both the primordial Universe and its late time evolution. These measurements will encode key insights on the nature of dark matter, dark energy and the physics of neutrinos, and may reveal for the first time the presence of primordial gravitational waves, thought to be generated during the inflationary phase at energies close to the scale of the Grand Unified Theories in the very early Universe.

Reaching these ambitious scientific goals entails building experiments with exceptional complexity. The extremely faint signals we are attempting to measure require both very high sensitivities, and hence the collection of huge volumes of data, as well as a precise control of all systematic effects. This makes the scientific exploitation and analysis of the data uniquely challenging from both a computational and methodological point of view.

This work focuses on tackling these issues in the early stages of the analysis where we deal with the full size of the time-domain data. I generalize the map-making procedure, to allow for efficient mitigation of some of the major time-domain systematic effects concurrently with the unbiased estimation of sky maps. I address both computational and methodological challenges, in the context of stage-3 CMB ground experiments such as Simons Array and the Simons Observatory, developing a massively parallel framework tested and validated on state-of-the-art supercomputing architectures. The framework provides robust methods to solve the map-making problem in the presence of numerous time-domain systematic effects, which need to be modeled with sufficient accuracy for a robust measurement of cosmological signals of interest such as B-modes from inflationary gravitational waves.

Keywords Cosmology, Cosmic Microwave Background, Data analysis, Map-making, Numerical methods, Linear systems solvers, High performance computing.

Résumé

Durant les deux dernières décennies, la cosmologie a été propulsée par de nombreux développements théoriques et expérimentaux lui permettant d'entrer dans l'ère de la haute précision. De nouvelles expériences sont actuellement en cours de préparation pour débiter leurs opérations dans les prochaines années et continuer de pousser les frontières du savoir et améliorer notre compréhension théorique des rouages internes de l'Univers. En particulier, la prochaine génération d'observatoires du fond diffus cosmologique (*Cosmic Microwave Background* - CMB) permettra la cartographie des anisotropies de température et de polarisation du CMB avec une précision inédite, sur des échelles allant du degré sur le ciel jusqu'aux minutes d'arc. Ces observatoires vont ainsi ouvrir une fenêtre unique sur l'Univers primordial ainsi que sur son évolution récente. Ces mesures contiennent des informations clés sur la nature de la matière noire, l'énergie noire, ainsi que la physique des neutrinos, et pourraient également révéler l'existence des ondes gravitationnelles primordiales, qui seraient générées durant la phase inflationnaire à des énergies proches de l'échelle des théories de grande unification aux premiers instant de l'Univers.

Atteindre ces ambitieux objectifs scientifiques nécessite la construction d'expériences d'une complexité exceptionnelle. Les signaux extrêmement faibles qu'on tente de mesurer requièrent de très grandes sensibilités, et donc la collection de volumes énormes de données, ainsi qu'un contrôle précis des effets systématiques. Ceci rend l'analyse et l'exploitation scientifique des données particulièrement difficile sur le plan, à la fois, computationnel et méthodologique.

Le travail présenté dans cette thèse se donne pour objectif de s'attaquer à ces problèmes dans les premières étapes de l'analyse où l'on traite tout le volume des données dans le domaine temporel. Je généralise la procédure de cartographie afin de permettre une mitigation efficace de certains des effets systématiques majeurs dans le domaine temporel simultanément avec l'estimation non biaisée des cartes du ciel. J'aborde à la fois les problèmes computationnels et méthodologiques, posés lors du développement d'observatoires du CMB au sol de type "stage-3" comme Simons Array et Simons Observatory, en développant un environnement massivement parallèle testé et validé sur des supercalculateurs à l'état de l'art mondial. L'environnement offre des méthodes robustes pour la résolution du problème de cartographie en cas de présence de nombreux effets systématiques dans le domaine temporel, qui doivent être modélisés avec une précision suffisante pour permettre une mesure robuste des signaux cosmologiques d'intérêt comme les modes-B issus des ondes gravitationnelles produites lors de l'inflation.

Mots-clés Cosmologie, Fond diffus cosmologique, Analyse de données, Cartographie, Méthodes numériques, Solveurs de systèmes linéaires, Calcul haute performance.

Contents

1	The Standard Cosmological Model	1
1.1	The homogeneous and isotropic Universe	3
1.1.1	The cosmological principle	3
1.1.2	The Friedmann-Lemaître-Robertson-Walker metric	4
1.1.3	Cosmological redshift	5
1.1.4	Distances	6
1.2	Dynamics of expansion	8
1.2.1	Friedmann equations	8
1.2.2	Evolution of the scale factor	9
1.2.3	Cosmological parameters	10
1.3	The Cosmic inventory	10
1.4	Physics of the Hot Big Bang	19
1.4.1	Thermodynamics	19
1.4.2	Thermal history of the Universe	23
1.5	Open questions of the standard model	31
2	Inflation	34
2.1	Addressing the cosmic puzzles	35
2.2	Inflationary dynamics	36
2.3	Seeding structure in the Universe	39
2.3.1	Metric perturbations	40
2.3.2	Matter perturbations	41
2.3.3	Quantum origin of structures	44
2.4	Observational evidence and constraints	46
3	The Cosmic Microwave Background	51
3.1	CMB anisotropies on the sphere	52
3.1.1	Temperature anisotropies	52
3.1.2	Polarization anisotropies	53
3.2	Evolution of perturbations	56
3.2.1	Boltzmann equation	57
3.2.2	Line of sight integral	58
3.2.3	Acoustic oscillations	59
3.3	Secondary anisotropies	62
3.3.1	Scattering effects	62
3.3.2	Gravitational effects	64
3.4	Cosmological parameters and observational constraints	69
3.5	Beyond anisotropies	74

3.5.1	Primordial non-Gaussianity	74
3.5.2	Spectral distortions	76
3.5.3	Sunyaev-Zel'dovich Clusters	79
3.5.4	Cosmic birefringence	81
4	Cosmology from multiple probes	83
4.1	Galaxy surveys	84
4.1.1	Baryon Acoustic Oscillations	85
4.1.2	Redshift-space distortions	86
4.1.3	Weak lensing and cosmic shear	88
4.2	Lyman- α forest	89
4.3	21-cm cosmology	90
4.4	Type-Ia supernovae	92
4.4.1	Hubble diagram	92
4.4.2	Cosmic distance ladder and the H_0 tension	94
4.4.3	Future cosmological targets	94
4.5	Gravitational waves	95
5	Extracting cosmology with current and future CMB experiments	98
5.1	The experimental landscape	99
5.1.1	Overview of CMB experiments	99
5.1.2	Selected Chilean experiments	102
5.1.3	Future experiments	110
5.2	Challenges of CMB observations	111
5.2.1	Systematics	112
5.2.2	Astrophysical foregrounds	115
5.2.3	Delensing	118
5.3	Data analysis pipelines for modern CMB experiments	119
6	Unbiased map-making in the presence of time-domain systematics	127
6.1	The map-making problem	128
6.1.1	Data model	128
6.1.2	Map-making solution	129
6.1.3	Pointing matrix	131
6.2	Maximum likelihood map-making	133
6.3	Templates marginalization map-making	135
6.3.1	Atmosphere and $1/f$ instrumental noise	136
6.3.2	Scan synchronous signal	136
6.3.3	Half-wave plate synchronous signal	137
6.4	Degeneracy issues	138
6.4.1	Scanning strategy sourced degeneracies	138
6.4.2	Templates linear dependencies	139
6.4.3	Templates and sky degeneracies	139
7	Massively parallel mapping and linear solvers techniques	140
7.1	Short introduction to parallel computing	142
7.2	Overview of the MAPPRaiser framework	143
7.3	MIDAPACK: the microwave data analysis package	144
7.3.1	MAPMAT: pointing algebra package	145

7.3.2	Toeplitz algebra package	146
7.3.3	Templates algebra package	149
7.4	MAPPRAISER: the MIDAPACK parallel iterative sky estimator . .	151
7.4.1	Mapping methods	151
7.4.2	MAPPRAISER solvers	152
7.4.3	Preconditioners for PCG	154
7.5	Tests and validation	159
7.5.1	Simulations description	160
7.5.2	Strong scaling tests and collective communication perfor- mance	161
7.5.3	Large run demonstration	163
7.5.4	Solvers and preconditioners performance	166
8	Map-making applications on satellite and ground-based experiments	176
8.1	Satellite experiments: LiteBIRD case	177
8.1.1	Simulation description	177
8.1.2	Full-sky maximum likelihood maps	178
8.1.3	Full-sky templates marginalization maps	180
8.2	Ground-based experiments: Simons Observatory, Small Aperture Telescopes case	182
8.2.1	Simulation	183
8.2.2	Validation: Scan & HWP synchronous signals mitigation . .	183
8.2.3	Atmospheric effects	187
	Conclusion	197
A	Block-wise derivation of the unbiased map estimator	199
B	Customized collective communication operations	202
B.1	Standard MPI_Allreduce	202
B.2	Ring scheme	203
B.3	Butterfly scheme	204
C	Analytical calculation of the standard IQU map-making and the pair- differencing case	206
D	Résumé détaillé de la thèse en français	211
D.1	Introduction générale	211
D.2	Contenu par chapitre	212
	Bibliography	216

List of Figures

- 1.1 A composite view of the Local Universe derived from the 2MASS Extended Source Catalog (XSC), showing about 1 million galaxies. The map is shown in a galactic Aitoff (equal-area) projection (Milky Way at center). The 2MASS Photometric Redshift catalogue has been used to perform statistical tests to detect possible deviations from the statistical homogeneity and isotropy in the large-scale structure of the Universe [55]. *Credit:* T. Jarrett (IPAC/Caltech). 4
- 1.2 The Hubble diagram for type Ia supernovæ compiled by Jha [265]. The error bars correspond to statistical distance errors of $< 10\%$ per object. The small red square in the bottom left of the plot indicates the range of Hubble’s original 1929 diagram. *Credit:* Robert P. Kirshner [283]. 5
- 1.3 A summary chart of the Standard Model particles and their properties. On the bottom right of the diagram, the hypothetical graviton is represented. It is supposed to be the mediator of gravitational interactions and is not currently in the Standard Model. *Credit:* CERN. 11
- 1.4 Illustration of the two neutrino mass orderings. The colors represent the relative probability of finding the different flavour eigenstates in the corresponding mass eigenstates. $\Delta m_{atm}^2 \equiv |\Delta m_{32}^2|$ refers to the “atmospheric” mass splitting, and $\Delta m_{sol}^2 \equiv \Delta m_{21}^2$ refers to the “solar” mass splitting. *Credit:* JUNO collaboration [368]. 15
- 1.5 *Left:* Evolution of the abundances of light elements as a function of the age of the Universe (in minutes) and of temperature. *Credit:* S. Burles et al. [115]. *Right:* Standard model BBN predictions of the abundances of ^4He , D, ^3He , and ^7Li - the bands show the 95% CL range. The yellow boxes represent the observed abundances. The vertical bands correspond to the 95% CL of respectively the CMB measurement of the baryon-to-photon ratio (narrow band) and the BBN D+ ^4He concordance range (wider band). *Credit:* Particle Data Group [219]. 28
- 1.6 Evolution of the ionization fraction as a function of redshift and temperature. *Credit:* D. Baumann [2]. 30

1.7	Conformal diagram illustrating the horizon problem in the framework of the standard Hot Big Bang cosmology. The past light cones of two regions on the last scattering surface, separated by 90° , do not intersect. This implies that the two regions are causally disconnected, which is at odds with the observed homogeneity of the CMB. <i>Credit:</i> D. Baumann [83].	33
2.1	At early times, regions separated by a given cosmological scale of interest, indicated by the blue horizontal line, are causally connected. This scale exits the horizon during the decreasing phase of the comoving Hubble radius. As the standard Big Bang expansion unfolds, the comoving Hubble radius increases until it reaches again this scale at late times. <i>Credit:</i> D. Baumann [83].	36
2.2	An illustration of the motion of the inflaton field rolling a generic potential $V(\phi)$. The field starts with a dominating potential energy, as the field rolls down the potential, the Universe is inflating, and the kinetic energy grows until it becomes no longer negligible $\frac{1}{2}\dot{\phi}^2 \approx V$. Inflation stops and the inflaton decays converting its energy density into radiation in a process called <i>reheating</i> . <i>Credit:</i> D. Baumann [83].	37
2.3	Illustration of a one-to-one mapping, ψ , between the FLRW background spacetime, $\overline{\mathcal{M}}$, and the perturbed spacetime \mathcal{M} . <i>Credit:</i> Patrick Peter & Jean-Philippe Uzan [369].	41
2.4	Illustration of the spacetime distortion induced by the two polarization states $+$ (upper panel) and \times (lower panel) of a propagating gravitational wave in the plane perpendicular to its direction. <i>Credit:</i> Scott Dodelson & Fabian Schmidt [165].	43
2.5	The Planck 2018 angular power spectra of the CMB (TT and TE). The solid blue line is the best-fit Λ CDM model to temperature and polarization data. <i>Credit:</i> Planck collaboration.	47
2.6	Evolution of monopole and dipole Fourier modes with arbitrary initial phases and different amplitudes, until the time of recombination τ_{rec} . <i>Credit:</i> Adapted version from [166] by Daniel Baumann [83].	47
2.7	Evolution of monopole and dipole Fourier modes with coherent initial phases and different amplitudes, until the time of recombination τ_{rec} . <i>Credit:</i> Adapted version from [166] by Daniel Baumann [83].	48
2.8	Constraints from <i>Planck 2018</i> data, and their combination with BICEP2/Keck array (BK15) and BAO data, on n_s and r compared to predictions of selected inflationary models (Cf. [377]). The contours represent the 68% and 95% CL regions. <i>Credit:</i> Planck collaboration [377]	49
3.1	Illustration of E-modes and B-modes polarization patterns.	55
3.2	<i>Left:</i> Impact of the optical depth to reionization on the temperature power spectrum. <i>Left:</i> Impact of the optical depth to reionization on the E-mode polarization power spectrum. <i>Credit:</i> [391]	63

3.3	Geometrical representation of a single deflection induced by gravitational lensing.	65
3.4	<i>Left:</i> The cumulative contributions to the angular lensing power spectrum as a function of maximum redshifts. The black curve indicates the total lensing power spectrum. <i>Right:</i> A comparison between the angular lensing power spectrum in the linear regime and non-linear corrections from HALOFIT [422]. <i>Credit:</i> [311] . . .	67
3.5	<i>Left:</i> TT and TE lensed and unlensed spectra and their differences in all-sky and flat-sky approximation as well as the error induced by the latter. <i>Right:</i> EE and BB lensed and unlensed spectra with their differences. Here we assume no tensor modes, and hence the unlensed BB vanishes. <i>Credit:</i> [251]	68
3.6	A compilation of state of the art measurements of CMB angular power spectra and the lensing power spectrum from multiple ground and satellite experiments. <i>Credit:</i> Planck & ACT collaborations [143].	71
3.7	<i>Left:</i> Ratio of C_ℓ^{TT} for different values of $\Delta N_{\text{eff}} = N_{\text{eff}} - 3.045$ over a reference model with $\Delta N_{\text{eff}} = 0$. <i>Right:</i> Ratio of C_ℓ^{TT} for different values of the total mass of neutrinos over a reference model with massless neutrinos. <i>Credit:</i> Particle Data Group [219].	72
3.8	Lensing and primordial B-modes for different values of r . <i>Credit:</i> Dominic Beck.	73
3.9	Spectral distortions of the CMB after energy injection at different redshifts z_h . <i>Credit:</i> Jens Chluba [4].	77
3.10	Maps of the Abell 2319 cluster in seven frequencies obtained by the Planck satellite. We see the local spectral distortion of the CMB caused by the SZ effect which results in a reduced intensity at low frequencies and an increased intensity at high frequencies. The transition happens at 217 GHz where the corresponding map shows no signal. <i>Credit:</i> Planck Collaboration.	80
3.11	<i>Left:</i> SZE surface brightness of RXJ1347.5-1145 by ALMA overlaid with X-ray contours from <i>Chandra</i> . <i>Right:</i> Residual tSZ signal revealed after subtracting the mean profile. Lensing contours from HST are overlaid, revealing the location of dark matter. <i>Credit:</i> S. Ueda et al. [457].	80
4.1	2D slice in the 3D distribution map of galaxies measured by the Sloan Digital Sky Survey (SDSS). Each point represents a galaxy. <i>Credit:</i> M. Blanton and the Sloan Digital Sky Survey.	84
4.2	The BAO peak in the galaxy correlation function measured from the LRG sample of eBOSS for redshifts $0.6 < z < 1$, representing an effective volume of 2.7 Gpc^3 . <i>Credit:</i> Bautista et al. [86].	85

4.3	<i>Left:</i> Constraints in the $\Omega_m - \Omega_k$ plane, on a model with cosmological constant and free curvature (o Λ CDM), from the combination of CMB temperature and polarization data alone (gray), CMB + SNe Ia data (red), and CMB + BAO (blue). <i>Right:</i> Constraints in the $\Omega_m - w$ plane, on a flat w CDM model, from CMB data alone (gray), SNe Ia data alone (red), and BAO data only (blue). <i>Credit:</i> SDSS [53].	86
4.4	The two-point correlation function ξ_g as a function of transverse separation $\sigma = cz\Delta\theta/H_0$ and radial separation $\pi = c\Delta z/H_0$. Solid lines represent contours of constant ξ_g . The dashed curves represent the best-fitting model. <i>Credit:</i> Hawkins et al. [239].	87
4.5	<i>Left:</i> $\Omega_m - \sigma_8$ 68% and 95% constraints. Blue contours represent constraints from RSD data alone, red contours correspond to weak lensing from DES, dark blue to CMB weak lensing from Planck, light purple to the combination of RSD, and lensing from both DES and Planck, while gray contours correspond to Planck temperature and polarization data. <i>Right:</i> Constraints (68% and 95%) on $\Sigma_0 - \mu_0$ from RSD data alone (blue), weak lensing from DES and CMB lensing from Planck (red), and their combination (gray). <i>Credit:</i> SDSS [53].	88
4.6	Simulated images of distant galaxies, shown in light blue, lensed by the dark matter shown in the red and white filaments. Weak gravitational lensing stretches galaxy images along directions which are parallel to the dark matter filaments. <i>Credit:</i> S. Colombi, IAP / Canada-France-Hawaii Telescope (CFHT) Corporation. . . .	89
4.7	Illustration of the evolution of the spectrum of light emitted from a distant quasar (<i>left</i>) as it travels through the IGM (<i>middle</i>) and near a galaxy (<i>right</i>) towards our telescopes. When travelling near a galaxy, saturated absorption lines are produced which are distinct from the low-density absorptions in the IGM, in addition metal absorption lines also appear at larger wavelengths than the Ly α emission line. <i>Credit:</i> Solène Chabanier [132] / Andrew Pontzen [6].	90
4.8	Best-fitting absorption profiles detected by EDGES for different hardware configurations and data analysis pipelines. <i>Credit:</i> J. D. Bowman et al. [104].	91
4.9	Hubble diagram constructed with the “JLA” sample comprising 740 SNe Ia, the black line is the best-fit Λ CDM model for a fixed $H_0 = 70$ km/s/Mpc. <i>Credit:</i> M. Betoule et al. [94].	93
4.10	Compilation of the different recent measurements of H_0 giving rise to a 4σ to 6σ tension between early-time and late-time measurements. <i>Credit:</i> Vivien Bonvin.	95
4.11	The spectrum of possible frequencies for gravitational waves sources and their corresponding detectors. <i>Credit:</i> NASA Goddard Space Flight Center.	96
4.12	Standard-siren measurement of H_0 compared with the other measurements from Planck[373] and SH0ES [395]. <i>Credit:</i> LIGO and Virgo collaborations [17].	97

5.1	History of some of the major CMB experiments along with their important milestones. Satellite experiments are shown in green, balloon-borne in purple, and ground-based in orange. <i>Credit:</i> Dominic Beck [87].	99
5.2	10 deg ² sky patches from the full-sky maps delivered by the three generation of satellites, COBE, WMAP, and Planck (left to right). <i>Credit:</i> NASA/JPL-Caltech/ESA.	100
5.3	Illustration of some of the main operating ground-based CMB experiments at the South Pole and at the Atacama desert in Chile. <i>Top left:</i> SPT. <i>Top right:</i> BICEP. <i>Bottom left:</i> POLARBEAR. <i>Bottom right:</i> ACT. <i>Credit:</i> SPT / BICEP/Keck / POLARBEAR / ACT collaborations.	101
5.4	Diagram illustrating the path of incoming light towards the polarization sensitive detectors of the POLARBEAR receiver. <i>Credit:</i> Neil Goeckner-Wald.	102
5.5	The four sky patches observed by the POLARBEAR experiment. The blue regions correspond to the small patches named RA4.5, RA12, and RA23 after their right ascension coordinates. The orange region corresponds to the large patch. <i>Credit:</i> Dominic Beck [87]	103
5.6	Illustration of the cross-linking resulting from scanning the same sky patch at different moments of the day at different elevations. The solid lines trace the attack angle of the scan occurring at the current position of the sky patch, while the dotted lines trace the attack angles of previous scans.	104
5.7	<i>Left:</i> Microscopic image of sinuous antenna and TES detector of POLARBEAR-2a. <i>Right:</i> Venus map observed with 150 GHz TESs in one detector module. <i>Credit:</i> POLARBEAR collaboration [273].	106
5.8	The expected Simons Array sky survey shown on top of the Planck 353 GHz map. The survey will target the low dust regions to minimize the effect of foregrounds. <i>Credit:</i> Aashrita Mangu.	107
5.9	Illustration of the Simons Observatory instruments designs.	108
5.10	Simons Observatory SATs and LAT surveys footprints on top of galactic dust emission. <i>Credit:</i> Simons Observatory collaboration [36].	108
5.11	Global design of the LiteBIRD payload module, and the onboard instruments and focal plane units. Figure compiled from 3D models given in [471, 339].	111
5.12	Illustration of the different signals that intervene along the line-of-sight of a typical CMB ground-based experiment. Note that we mention only dust and synchrotron foregrounds, which are the most relevant diffused and polarized contaminants. Other contributions exist, however, at the smallest scales as well as for temperature data.	112
5.13	Simons Array bandpasses on top of the atmospheric transmission profile at the Chajnantor plateau during good weather conditions. <i>Credit:</i> John Groh.	113
5.14	Caption for LOF	114

5.15	Realistic TOAST simulation of the HWPSS signal expected in the Simons Observatory Small Aperture Telescopes, based on [403]. The different panels show the HWPSS template as a function of the HWP angle for the 2f, 4f and 6f harmonics as well as their sum for two detectors of an orthogonal pair. Each panel represents either the total spurious signal, or the component corresponding to either transmission, reflection or emission (see titles above each panel). <i>Credit:</i> Reijo Kesitalo.	115
5.16	<i>Left:</i> Expected B-mode foregrounds contamination (dust+synchrontron) on the cleanest 1-90% of the sky at 100 and 200 GHz. <i>Right:</i> RMS of the polarization amplitudes of foregrounds components as a function of frequency, compared with CMB. <i>Credit:</i> Josquin Errard / Planck Collaboration [376]. . .	116
5.17	Elongated dust grains aggregates, aligning preferentially along the perpendicular direction to the magnetic field of the galaxy, absorb unpolarized starlight and re-emit polarized emission aligned with their long axis.	117
5.18	The green curve shows the lensing induced B-modes which coincides on large scales with the gray dashed line representing $5\mu\text{K}$ -arcmin white noise. Delensing can significantly suppress this effect as shown in the green dashed line, resulting in lower effective noise in the B-mode map. <i>Credit:</i> CMB-S4 Science Book [16].	118
7.1	Moore's law for CMB experiments and supercomputer performance. <i>Credit:</i> Julian Borill [1].	141
7.2	Overview of the software top-level architecture.	143
7.3	Distribution scheme of the TOD depending on the processed data set size. Each cell in the data matrix corresponds to one detector time stream for the full duration of a scan. Depending on the size of the data set, each MPI process represented by the different colors in the diagram, will then handle a subset of detectors for a single scan or all detectors for a subset of scans. The lines separating the detector scans in a single process are dashed, since they actually hold a single contiguous block of memory.	145
7.4	Illustration of the shift and overlap algorithm implemented in the Toeplitz algebra package. <i>Credit:</i> MIDAPACK v1.1b documentation [7].	148
7.5	The Cori machine at NERSC. <i>Credit:</i> NERSC.	159
7.6	The Joliot-Curie machine at TGCC. <i>Credit:</i> Atos.	160
7.7	Strong scaling of the mean time per iteration comparing 5 sets of runs: (1) runs executed at the Cori machine at NERSC using the <code>MPI_Allreduce</code> communication scheme; (2) runs executed at the Cori machine at NERSC using the Butterfly communication scheme; (3) runs executed at the Cori machine at NERSC using the Ring communication scheme; (4) runs executed at the Joliot-Curie machine at TGCC using the <code>MPI_Allreduce</code> communication scheme; (5) runs executed at the Joliot-Curie machine at TGCC using the Butterfly communication scheme.	162

7.8	Strong scaling efficiency comparing 4 sets of runs: (1) runs executed at the Cori machine at NERSC using the <code>MPI_Allreduce</code> communication scheme; (2) runs executed at the Cori machine at NERSC using the Butterfly communication scheme; (3) runs executed at the Joliot-Curie machine at TGCC using the <code>MPI_Allreduce</code> communication scheme; (4) runs executed at the Joliot-Curie machine at TGCC using the Butterfly communication scheme.	163
7.9	Convergence curve showing the relative residuals as a function of the wall clock time.	164
7.10	Hits map of a $5 \text{ deg} \times 5 \text{ deg}$ sky patch located at 45 deg longitude and -42 deg latitude.	164
7.11	Q and U polarization maps obtained from the large run demonstration.	165
7.12	Scaling of ECG iterations with the enlarging factor, using 1h of SO-SAT simulated data.	166
7.13	Convergence in terms of the number of iterations of the BD preconditioned CG, compared with the <i>a priori</i> two-level preconditioned CG, for different dimensions of the deflation subspace.	168
7.14	Convergence in terms of the number of iterations of the BD preconditioned CG, compared with the <i>a posteriori</i> two-level preconditioned CG, for different dimensions of the deflation subspace. . .	169
7.15	<i>Left</i> : Precomputation cost of the <i>a posteriori</i> preconditioner as a function of the dimension of the deflation subspace matrix $\text{dim } \mathbf{Z}$. In practice this dimension is the number of iterations performed by the Lanczos procedure. <i>Right</i> : Average iteration cost with the same preconditioner as a function of the dimension of the deflation subspace matrix $\text{dim } \mathbf{Z}$	170
7.16	Convergence in wall clock time of the BD preconditioned CG, compared with the <i>a posteriori</i> two-level preconditioned CG, for different dimensions of the deflation subspace. The precomputation time is not included here.	170
7.17	Temperature angular power spectra of the maps at different stages of convergence of the BD preconditioned CG.	171
7.18	EE angular power spectra of the maps at different stages of convergence of the BD preconditioned CG.	172
7.19	Temperature angular power spectra of the maps at different stages of convergence of the 2lvl- $\text{dim } \mathbf{Z} = 128$ preconditioned CG.	172
7.20	EE angular power spectra of the maps at different stages of convergence of the 2lvl- $\text{dim } \mathbf{Z} = 128$ preconditioned CG.	173
7.21	Temperature angular power spectra of the maps at different stages of convergence of the 2lvl- $\text{dim } \mathbf{Z} = 512$ preconditioned CG.	173
7.22	EE angular power spectra of the maps at different stages of convergence of the 2lvl- $\text{dim } \mathbf{Z} = 512$ preconditioned CG.	174
7.23	Temperature angular power spectra of the maps at different stages of convergence of the 2lvl- $\text{dim } \mathbf{Z} = 2048$ preconditioned CG.	174
7.24	EE angular power spectra of the maps at different stages of convergence of the 2lvl- $\text{dim } \mathbf{Z} = 2048$ preconditioned CG.	175

8.1	Caption for LOF	177
8.2	Full-sky maximum likelihood maps from 1 year of simulated observations of the LiteBIRD space mission. The <i>top left</i> map corresponds to the LFT 40 GHz channel, the <i>top right</i> one to the MFT 140 GHz channel, and the <i>bottom</i> map corresponds to the HFT 235 GHz channel.	179
8.3	Comparison of the noise power spectra obtained via the different map-making methods available in MAPPRASER. The baseline for the Template marginalization run is set to 120 s. The maps are computed from the simulated data of the LFT 40 GHz channel. . . .	179
8.4	The <i>top left</i> map shows the difference between the binned temperature noise map and the maximum likelihood one. The <i>top right</i> and <i>bottom</i> maps shows the differences between the templates marginalization temperature noise maps, respectively, for a baseline of 120 s and 20 s, and the maximum likelihood one. The maps are computed from the simulated data of the LFT 40 GHz channel. . . .	181
8.5	The temperature noise power spectra obtained via the templates marginalization approach for different baselines. The <i>left</i> figure shows the $\ell \in [2, 30]$ range, and the <i>right</i> figure shows the $\ell \in [30, 10^3]$ range. I also include the noise power spectra from the binned and maximum likelihood maps for comparison. The maps are computed from the simulated data of the LFT 40 GHz channel. . . .	181
8.6	Hits map of the S2 simulation.	183
8.7	Temperature maps of CMB+SSS mapped with respectively, a binned map-maker, a 10 azimuth bins, and 30 azimuth bins SSS templates map-makers (top to bottom).	185
8.8	Temperature maps of CMB+HWPSS mapped with respectively, a binned map-maker, a fourth, and sixth order HWPSS templates map-makers (top to bottom).	186
8.9	Q polarization maps of CMB+HWPSS mapped with respectively, a binned map-maker, a fourth, and sixth order HWPSS templates map-makers (top to bottom).	187
8.10	<i>Left</i> : Distribution of the knee frequencies induced by atmospheric noise for all detectors in the focal plane and for all the scans of the S2 simulation. The binning in frequency is $\Delta\nu = 5 \cdot 10^{-3}$ Hz. <i>Right</i> : Inverse periodogram computed from the noise timestream of a single detector during one CES, overlaid with the best-fit to our PSD model and the resulting PSDs after apodization with Gaussian windows corresponding to different bandwidths of the Toeplitz blocks.	189
8.11	Convergence curves for different values of the half bandwidth in the SAT simulation.	190
8.12	EE and BB noise power spectra for different values of the half-bandwidth, compared with the C_ℓ of the input sky signal (solid black curve), in the SAT simulation.	190
8.13	Relative difference of the EE and BB noise power spectra for different values of the half-bandwidth, with the case $\lambda = 4096$, in the SAT simulation.	191

8.14	Convergence curves for different orders of the polynomial filters and different values of the half-bandwidth in the SAT simulation. .	191
8.15	EE and BB noise power spectra for different orders of the polynomial filters, compared with the C_ℓ of the input sky signal (solid black curve), in the SAT simulation.	192
8.16	Relative differences of the EE and BB noise power spectra for the pair differencing case along with the different orders of the polynomial filters, relative to $n_{poly} = 6$, in the SAT simulation.	193
8.17	Distribution of the knee frequencies of orthogonal pairs in the S2 simulation for different levels of gain mismatch. The binning in frequency is $\Delta\nu = 2 \cdot 10^{-3}$ Hz. The residual atmosphere signal causes the right tail of the distribution to get stronger, increasing the effective knee frequency.	194
8.18	Q and U polarization maps obtained with different levels of gain mismatch and no HWP.	195
8.19	EE and BB noise power spectra for different levels of gain mismatch between orthogonal pairs, in the SAT simulation with no HWP.	195
8.20	Q and U polarization maps obtained with different levels of gain mismatch in the presence of a HWP rotating at 2 Hz.	196
8.21	EE and BB noise power spectra for different levels of gain mismatch between orthogonal pairs, in the SAT simulation with HWP.	196
B.1	Ring communication pattern. <i>Credit:</i> Adapted from Pierre Cargemel.	203
B.2	Butterfly communication pattern. <i>Credit:</i> Adapted from Pierre Cargemel.	205

Note on units

Throughout this thesis, and unless I explicitly specify otherwise, the choice of units will be set so that the speed of light in vacuum c , the Boltzmann constant k_B , and the reduced Planck constant \hbar are all equal to 1.

CHAPTER 1

The Standard Cosmological Model

Contents

1.1	The homogeneous and isotropic Universe	3
1.1.1	The cosmological principle	3
1.1.2	The Friedmann-Lemaître-Robertson-Walker metric	4
1.1.3	Cosmological redshift	5
1.1.4	Distances	6
1.2	Dynamics of expansion	8
1.2.1	Friedmann equations	8
1.2.2	Evolution of the scale factor	9
1.2.3	Cosmological parameters	10
1.3	The Cosmic inventory	10
1.4	Physics of the Hot Big Bang	19
1.4.1	Thermodynamics	19
1.4.2	Thermal history of the Universe	23
1.5	Open questions of the standard model	31

It is truly an astonishing fact that only a century ago, Shapley and Curtis, both prominent astronomers at the time, were debating, in what came to be known later as *The Great Debate* [412, 247], whether or not the entire Universe was made of a single galaxy, the Milky Way, comprised of all the stars and nebulae we observe in the sky. This reveals how far away the contemporary scientific understanding was from grasping the full scope and complexity of cosmic structure as it is understood and discussed in the present-day scientific literature. The controversy, which largely revolved around the nature of the so-called *spiral nebulae*, was settled shortly after when Edwin Hubble provided the definite evidence of their extragalactic nature using Cepheid variable stars to determine their distances [256, 257]. These were in fact independent spiral galaxies. The following

years led to theoretical and observational breakthroughs which completely revolutionized our understanding of the true scale and dynamics of the Universe, paving the way to the emergence of *Cosmology* as a new scientific discipline concerned with the study of the physical origin, structure and evolution of the Universe, thereby bringing fundamental questions which were previously confined to the realms of philosophy and mythology under the rigorous scrutiny of the scientific method.

The foundation of cosmology lies in the theory of general relativity which opened up the possibility to quantitatively study world models by providing a mathematical framework for how space-time behaves depending on its matter-energy content. By assuming a homogeneous and isotropic Universe at large scales, analytical solutions can be derived. In fact, Einstein himself was the first to realize this and proposed his first relativistic world model, a static universe [177], immediately after formulating the final version of his theory of gravity. It was soon followed by many alternative proposals, most importantly from De Sitter [160], Friedmann [204, 205], and Lemaître [310]. These works not only demonstrated the non-uniqueness of Einstein’s model, but also pointed to the possibility of a dynamic universe. Subsequent work showed that in fact Einstein’s static universe was unstable [173]. Hubble’s observations [258] constituted yet again a turning point in cosmology by providing the convincing demonstration that the Universe was indeed expanding. The natural extrapolation of this fact into the past ultimately led to the realization that the Universe went through a very hot and dense state in its early history: the Big Bang.

In addition to the Universe’s expansion, two more pieces of evidence served to consolidate the emerging picture of a very hot and dense early universe as suggested by the Big Bang scenario. The first important piece was a direct consequence of theoretical work conducted by Gamow and Alpher to explain the origin of chemical elements in the hot early Universe [57]. This eventually led to the emergence of the so called Big Bang Nucleosynthesis Theory (BBN), which provided the theoretical basis for explaining the origin and abundances of light elements, as opposed to heavy elements which are synthesized in stars, and achieved excellent agreement with observations. The second piece was a by-product of the first work, but of no less importance. It was the prediction, by Alpher and Herman [58] of relic radiation from the hot primordial Universe coming from all directions of the sky with a blackbody spectrum with temperature of about 5 K. This radiation is observable in the microwave range of frequencies and would later be known as the *Cosmic Microwave Background* (CMB). It was first accidentally observed by Penzias and Wilson [365] roughly two decades later. In the 1990s, the Far InfraRed Absolute Spectrometer (FIRAS) onboard the Cosmic Background Explorer (COBE) space mission performed a high precision measurement of the CMB blackbody spectrum giving a temperature of 2.725 ± 0.002 K [200, 329]. In addition to these considerations, it was also understood, that the early Universe should feature small departures from homogeneity that would ultimately grow through gravitational collapse to form the large scale structures that comprise our observable universe. These inhomogeneities can be characterized through the observation of anisotropies in the CMB. In fact, this is one of the major fields of study in modern cosmology, and it is at the core of the research conducted in this thesis. In Chapter 2 we will discuss a possible scenario for

the origin of these fluctuations, followed by a more thorough discussion of the physics of CMB anisotropies in Chapter 3.

Simultaneously to these developments, astronomical observations were accumulating ample evidence for the existence of a non-luminous form of matter, the presence of which can only be determined through its gravitational effects, for example through the estimation of galaxy cluster mass [481] or through the measurement of the rotation curves of spiral galaxies [399]. These observations established that this *dark matter* is in fact the dominant form of mass in galaxies and galaxy clusters. An even more striking discovery came about by the end of the 1990s after studies of type Ia supernovæ allowed to extend the Hubble diagram to very large distances thereby allowing to probe the expansion history of the Universe, and determined that it was recently undergoing an accelerated expansion [396, 366]. This is pointing in two possible directions: either a new form of energy, so called “dark energy”, the precise nature of which is yet to be determined, dominated the Universe in its recent past and is driving the accelerated expansion, or Einstein’s theory of gravity needs to be corrected. In its simplest form, dark energy is assumed to be equivalent to a cosmological constant (Λ), that is a constant correction term in the field equations of general relativity, that is usually interpreted as the energy density of space itself, or vacuum energy. Historically, it was introduced by Einstein himself in his static universe model. By the beginning of the 21st century, a concordance model of cosmology, capturing all of this rich physics in a total of six parameters, began to emerge. It was named Λ CDM after the components of the dark sector which form 95% of the energy content of the Universe today, and demonstrated spectacular agreement with observations. As of today, observations have allowed cosmologists to determine the parameters with sub-percent precision levels [373, 172]. It is worth mentioning however, that an increasingly significant tension [473] on the expansion rate of the Universe, between the early-time model dependent measurements and the direct late-time measurements, could be pointing to new physics beyond Λ CDM and its simplest extensions [164].

In this first chapter, we will explore more quantitatively how the different elements laid out in the previous paragraphs are assembled together to form the standard cosmological model, and what reasons may drive us to consider new physical ingredients in our description of the Universe. Large parts of this chapter are following lecture notes of Lesgourgues and Baumann [3, 2].

1.1 The homogeneous and isotropic Universe

1.1.1 The cosmological principle

The starting point of any cosmology is the Copernican principle, which states that we do not occupy a special place in the Universe, or with a stronger formulation: “The Universe should look the same for all observers”. Two immediate consequences of this principle are the homogeneity and isotropy of the Universe at large scales, which together form the basis of the so-called *cosmological principle*.

Today the statistical isotropy is supported with strong evidence from observations of the distribution of galaxies at large scales (see for instance figure 1.1) and the CMB which exhibits anisotropies only at the 10^{-5} level. By virtue of the

Copernican principle, the isotropy can be generalized to all observers wherever they are located. The immediate consequence is that the Universe should be homogeneous as well.

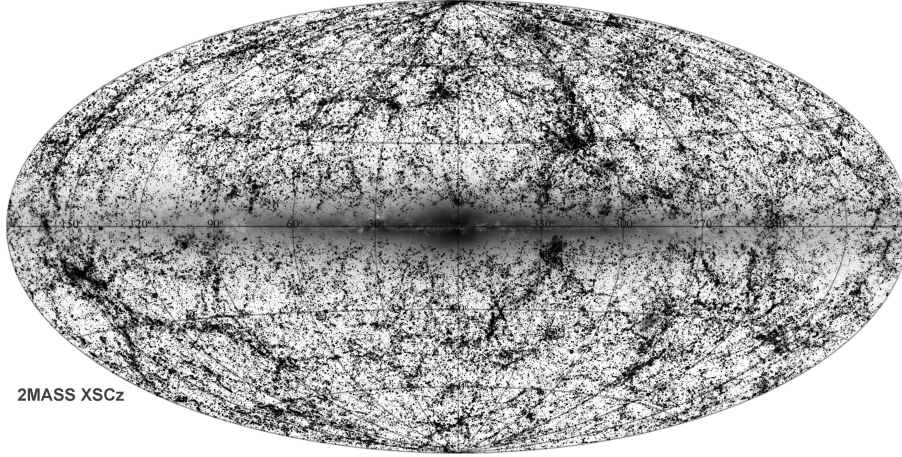


Figure 1.1 – A composite view of the Local Universe derived from the 2MASS Extended Source Catalog (XSC), showing about 1 million galaxies. The map is shown in a galactic Aitoff (equal-area) projection (Milky Way at center). The 2MASS Photometric Redshift catalogue has been used to perform statistical tests to detect possible deviations from the statistical homogeneity and isotropy in the large-scale structure of the Universe [55]. *Credit: T. Jarrett (IPAC/Caltech).*

1.1.2 The Friedmann-Lemaître-Robertson-Walker metric

Under the assumptions of homogeneity and isotropy, the most general form of the metric is the Friedmann–Lemaître–Robertson–Walker (FLRW) metric, which in comoving spherical coordinates (r, θ, ϕ) reads

$$\begin{aligned} ds^2 &= g_{\mu\nu} dx^\mu dx^\nu \\ &= -dt^2 + a^2(t) \left[\frac{dr^2}{1 - kr^2} + r^2(d\theta^2 + \sin^2 \theta d\phi^2) \right], \end{aligned} \quad (1.1)$$

The comoving coordinates define a space grid which follows the Hubble flow, i.e. the expansion of space. If we consider distant galaxies at rest in the comoving frame, the physical distances between the galaxies would continuously increase. The time dependent parameter $a(t)$ in Eq. (1.1), called *the scale factor*, encodes this expansion dynamic: the physical distance d between two galaxies is given by the product of the comoving distance χ and the scale factor a . It follows that the relative velocity is given by

$$v(t) = \dot{a}(t)\chi \equiv H(t)d(t), \quad (1.2)$$

where the *Hubble parameter*, $H(t)$, is defined as the relative time variation of the scale factor

$$H(t) \equiv \frac{\dot{a}}{a}(t). \quad (1.3)$$

By convention, the scale factor is taken to be equal to unity at the present time t_0 , $a(t_0) = 1$, and the current value of the Hubble parameter is denoted $H_0 \equiv$

$H(t_0)$. The linear relation between the recession velocities and distances between galaxies, given in Eq. (1.2), is the *Hubble-Lemaître law*. It was first demonstrated by Hubble using Cepheid variables in his famous 1929 paper [258]. Since then, the analysis of observations of type Ia supernovae allowed us to prove that the law holds true far beyond the range of distances studied by Hubble’s original work, as can be seen in figure 1.2.

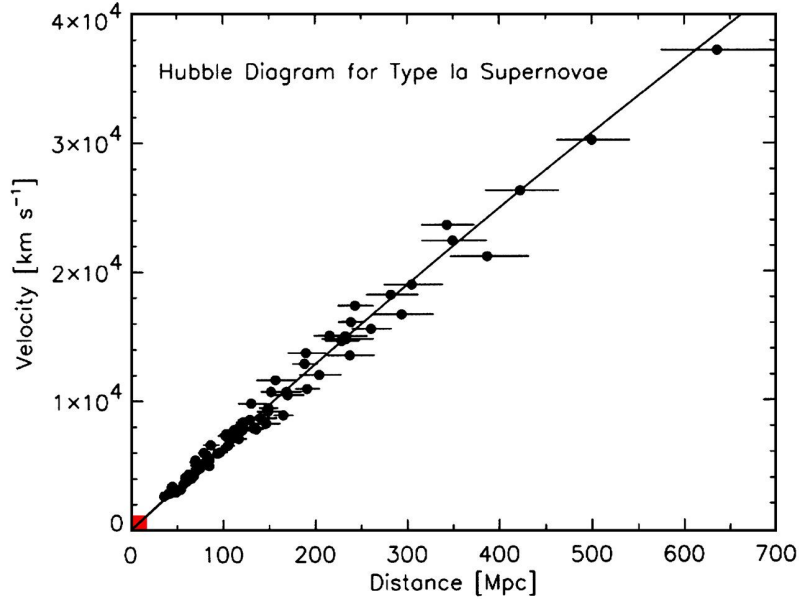


Figure 1.2 – The Hubble diagram for type Ia supernovae compiled by Jha [265]. The error bars correspond to statistical distance errors of $< 10\%$ per object. The small red square in the bottom left of the plot indicates the range of Hubble’s original 1929 diagram. *Credit: Robert P. Kirshner [283].*

The Hubble parameter is also used to derive two characteristic scales of the FLRW metric. The time scale of expansion, also called the *Hubble time*, $t_H \equiv H^{-1}$, and the *Hubble length*, $l_H \equiv ct_H = H^{-1}$ ($c = 1$ in natural units).

The constant parameter k in Eq. (1.1) describes the spatial curvature of the Universe. We can consider three separate cases: $k = 0$ for a “flat universe” (i.e. Euclidean space), $k > 0$ for a “closed universe” (i.e. positively curved space), and $k < 0$ for an “open universe” (i.e. negatively curved space). We can introduce a second length scale associated with the FLRW metric, the radius of curvature which reads $R_{\text{curv}}(t) = a(t)/\sqrt{|k|}$.

1.1.3 Cosmological redshift

Redshift is a key observable in cosmology, as it is essential to define measures of distance. Consider a source of light $S(r_S, \theta_S, \phi_S)$ and an observer O at the origin of coordinates, both with no peculiar motion. Light leaves the source at time t_1 with wavelength λ_1 and reaches the observer at time t_0 with wavelength λ_0 . We denote δt_1 and δt_0 the corresponding periods of the signal at emission and reception respectively.

The photons leaving the source will propagate in straight lines in space ($d\theta =$

$d\phi = 0$) following null geodesics

$$ds^2 = -dt^2 + a^2(t) \left[\frac{dr^2}{1 - kr^2} \right] = 0. \quad (1.4)$$

Integrating the above relation, we can write the following for two successive wave crests (S and O are at rest)

$$\int_0^{r_S} \frac{dr}{\sqrt{1 - kr^2}} = \int_{t_1}^{t_0} \frac{dt}{a(t)} = \int_{t_1 + \delta t_1}^{t_0 + \delta t_0} \frac{dt}{a(t)}. \quad (1.5)$$

It follows from the second equality that

$$\int_{t_1}^{t_1 + \delta t_1} \frac{dt}{a(t)} = \int_{t_0}^{t_0 + \delta t_0} \frac{dt}{a(t)} \implies \frac{\delta t_1}{a(t_1)} = \frac{\delta t_0}{a(t_0)}. \quad (1.6)$$

Hence we can derive a relation between the emission and reception wavelengths

$$\frac{\lambda_0}{\lambda_1} = \frac{\delta t_0}{\delta t_1} = \frac{a(t_0)}{a(t_1)} \quad (1.7)$$

The wavelengths are therefore stretched by the ratio of the scale factor between reception and emission, that is the factor by which space has expanded since the light signal was emitted. Qualitatively this corresponds to a shift towards red wavelengths. This shift is parametrized by a quantity unsurprisingly called *redshift*, z , given by

$$z \equiv \frac{\Delta\lambda}{\lambda} = \frac{\lambda_0 - \lambda_1}{\lambda_1} = \frac{a(t_0)}{a(t_1)} - 1 \quad (1.8)$$

If t_0 corresponds to the present day, and t_1 to a random instant in the past t , then the previous relation can be written as

$$1 + z = \frac{1}{a}. \quad (1.9)$$

A measurement of the redshift of a given distant galaxy therefore allows us to know the scale factor at the time t when light was emitted, and therefore provides a measure of how much smaller the Universe was at the time of emission.

1.1.4 Distances

In a curved spacetime, like the FLRW Universe, the notion of physical distance is ambiguous, as different measures of distance are not equivalent, and agree only locally for small redshifts ($z \ll 1$), where spacetime curvature is not noticeable. In the following, we will consider two measures of distance commonly used in cosmology, angular diameter distance and luminosity distance, and derive their relation to redshift. We start off by deriving an expression of the comoving distance as a function of redshift.

Comoving distance

As we have seen before, the proper distance to a comoving object with spherical coordinates (r, θ, ϕ) at the present time t_0 is equal to the comoving distance $\chi(r)$ up to a normalization with the scale factor a . The two quantities coincide with the convention $a(t_0) = 1$. Their expression is given by

$$\begin{aligned}\chi(r) &= \int_0^r \frac{dr'}{\sqrt{1 - kr'^2}} \\ &= \begin{cases} k^{-1/2} \sin^{-1}(k^{1/2}r) & \text{if } k > 0, \\ r & \text{if } k = 0, \\ |k|^{-1/2} \sinh^{-1}(|k|^{1/2}r) & \text{if } k < 0. \end{cases}\end{aligned}\quad (1.10)$$

We define the function

$$S_k(x) \equiv \begin{cases} k^{-1/2} \sin(k^{1/2}x) & \text{if } k > 0, \\ x & \text{if } k = 0, \\ |k|^{-1/2} \sinh(|k|^{1/2}x) & \text{if } k < 0, \end{cases}\quad (1.11)$$

such that $r = S_k(\chi)$.

We can relate this quantity to redshift by using the integrated relation over null geodesics (see Eq. (1.5)) which gives $\chi(t_e) = \int_{t_e}^{t_0} \frac{dt}{a(t)}$, where t_e is the time of light emission. We can then use the redshift relation Eq. (1.9), and the definition of the Hubble parameter in Eq. (1.3) to find the expression

$$\chi(z) = \int_0^z \frac{dz'}{H(z')}. \quad (1.12)$$

Angular diameter distance

The physical size of some astronomical objects, called *standard rulers* can be known in advance. In these cases, a measure of distance called *angular diameter distance*, d_A , can be inferred from the ratio of their physical transverse size ds to the observed angular diameter $d\theta$, $d_A = ds/d\theta$.

From the FLRW metric, the physical transverse size of an object at the time when light is emitted, t_e , is related to its observed angular diameter $d\theta$ through $ds = a(t_e) r_e d\theta$, with r_e its comoving radial coordinate. Hence the angular diameter distance as a function of redshift is given by

$$d_A(z) = \frac{S_k(\chi)}{1+z} = (1+z)^{-1} S_k \left(\int_0^z \frac{dz'}{H(z')} \right). \quad (1.13)$$

Luminosity distance

Astronomical objects called *standard candles* such as Cepheids and type Ia supernovæ for which the absolute luminosity L can be determined independently of their distance and measured flux S , allow for another measure of distance called *luminosity distance*, d_L , given by

$$d_L \equiv \sqrt{\frac{L}{4\pi S}}. \quad (1.14)$$

To relate this distance to redshift, we consider a source of light at comoving coordinate r . The photons emitted spread over a sphere of surface area $4\pi r^2$ ($a(t_0) = 1$). Moreover, the photons wavelengths are redshifted by a factor $(1 + z)$ leading to a drop in their corresponding energy by the same factor. The same effect of cosmological time dilation that stretches photons wavelengths, also slows the rate at which they are received by the same factor $(1 + z)$. Hence we can write how the measured flux relates to the absolute luminosity in these terms

$$S = \frac{L}{4\pi r^2(1 + z)^2} \quad (1.15)$$

We can now derive the expression of the luminosity distance as a function of redshift

$$d_L = (1 + z)S_k(\chi) = (1 + z)S_k \left(\int_0^z \frac{dz'}{H(z')} \right) = (1 + z)^2 d_A(z). \quad (1.16)$$

1.2 Dynamics of expansion

1.2.1 Friedmann equations

General Relativity allows us to study the dynamics of expansion by relating the spacetime metric to the matter-energy content of the Universe. This relation is given by Einstein's field equations

$$G_{\mu\nu} + \Lambda g_{\mu\nu} \equiv R_{\mu\nu} - \frac{1}{2}g_{\mu\nu}\mathcal{R} + \Lambda g_{\mu\nu} = 8\pi G T_{\mu\nu}, \quad (1.17)$$

where $R_{\mu\nu}$ is the Ricci tensor, and $\mathcal{R} \equiv g^{\mu\nu}R_{\mu\nu}$ is the Ricci scalar. They can be computed using Christoffel symbols, and their combination in the first two terms of the left hand side of Eq. (1.17) gives the Einstein tensor $G_{\mu\nu}$ which, for the FLRW metric, is found to be diagonal ($G_{0i} = G_{i \neq j} = 0$) and isotropic ($G_{11} = G_{22} = G_{33}$).

$T_{\mu\nu}$ is the energy-momentum tensor which under the assumption of homogeneity and isotropy takes the following form of a perfect fluid

$$T^\mu_\nu = \begin{pmatrix} -\rho & 0 & 0 & 0 \\ 0 & p & 0 & 0 \\ 0 & 0 & p & 0 \\ 0 & 0 & 0 & p \end{pmatrix}, \quad (1.18)$$

where ρ and p refer respectively to the energy density and the pressure of the cosmic fluid.

Λ is the cosmological constant (it does not vary with time or depends on space). As stated previously, it was first introduced by Einstein who observed that its addition doesn't violate energy conservation and allowed for a static universe. The static universe solution was dismissed long ago, but the term has regained its relevance since the discovery of accelerated expansion. In fact, by moving the term to the right hand side of Eq. (1.17), the term can be interpreted as the contribution to the energy-momentum tensor, of a homogeneous fluid with equation of state $\rho = -p = \frac{\Lambda}{8\pi G}$.

In the FLRW metric, Einstein's field equations reduce to two non-linear and coupled ordinary differential equations, called *Friedmann equations*

$$\left(\frac{\dot{a}}{a}\right)^2 + \frac{k}{a^2} = \frac{8\pi G}{3}\rho + \frac{\Lambda}{3}, \quad (1.19)$$

$$\frac{\ddot{a}}{a} = -\frac{4\pi G}{3}(\rho + 3p) + \frac{\Lambda}{3}. \quad (1.20)$$

1.2.2 Evolution of the scale factor

From the above equations, one can derive the *continuity equation* which also stems from the local energy-momentum conservation law ($\nabla_\mu T^\mu_\nu = 0$)

$$\dot{\rho} + 3\frac{\dot{a}}{a}(\rho + p) = 0. \quad (1.21)$$

Hence the scaling law of the energy density ρ as a function of a , relies on the equation of state of the $p(\rho)$. We can consider three limiting cases depending on what matter-energy component dominates the Universe at each epoch:

- **Matter:** in the common terminology used in cosmology, the term "matter" refers to the strongly non-relativistic matter, such that its kinetic energy is negligible leading to a vanishing pressure, i.e. an equation of state $p = 0$.
- **Radiation:** the term "radiation" in cosmology, refers to the ultra-relativistic particles such as photons or other particles of vanishing rest mass, such that their thermal velocity is equal or very close to c . In this case the equation of state reads $p = \rho/3$.
- **Cosmological constant:** As stated above, the cosmological constant is equivalent to a homogeneous fluid with equation of state $p = -\rho$.

We can write a general equation of state parametrized by a constant w whose value depends on the considered component

$$p = w\rho. \quad (1.22)$$

The integration of the continuity equation then leads to the following scaling law

$$\rho \propto a^{-3(1+w)}, \quad (1.23)$$

which also allows us to derive the time evolution of the scale factor

$$a(t) \propto \begin{cases} t^{\frac{2}{3(1+w)}} & \text{if } w \neq -1, \\ e^{Ht} & \text{if } w = -1. \end{cases} \quad (1.24)$$

1.2.3 Cosmological parameters

It is generally useful to introduce a parametrization that explicitly includes the relative energy densities of the different species in the Universe. For this purpose we introduce the *critical energy density*, ρ_c , that is the energy density required for the Universe to be spatially flat. It is defined as

$$\rho_c = \frac{3H_0^2}{8\pi G} \quad (1.25)$$

We define the density parameters, Ω_i , of the different species as the relative ratio of their present energy density to the critical density

$$\Omega_i \equiv \frac{\rho_i(t_0)}{\rho_c} \quad (1.26)$$

We can now rewrite the first Friedmann equation (1.19) using this parametrization

$$\left(\frac{H}{H_0}\right)^2 = \sum_i \Omega_i a^{-3(1+w_i)} + \Omega_k a^{-2}, \quad (1.27)$$

where we have introduced the curvature energy density, $\Omega_k \equiv -k/H_0^2$. The evaluation of the above equation at the present day, ($a_0 = 1$), gives

$$\sum_i \Omega_i + \Omega_k = 1. \quad (1.28)$$

As expected, for a flat universe ($\Omega_k = 0$), the total energy density corresponds to the critical density, $\Omega_{tot} = 1$.

It is important to note that given a set of cosmological parameters $\{H_0, \Omega_r, \Omega_m, \Omega_\Lambda, \Omega_k\}$, one can explicitly compute $H(z)$ from Eq. (1.27), and use it to derive a model dependent value of the cosmological distances we have defined in the previous section. Conversely, we can use observations of standard candles or standard rulers to fit for these parameters.

1.3 The Cosmic inventory

In this section, we will take a closer look at the different matter-energy components of the Universe insofar as we understand them. For completeness, we will first start with a brief and qualitative introduction to the Standard Model of particle physics, before we consider how the elements fit in the cosmological arena. Modern cosmology has a strong role to play in advances in particle physics, as cosmological observables encode information that can be key to unlocking some of the mysteries that are still lurking in our current understanding of the subatomic world. We will review some of the possible solutions to deviations from the Standard Model picture, and how they can be enlightened by cosmological observations.

A lightning introduction to the Standard Model of particle physics

The Standard Model is the best description we have to date of the physics of (known) fundamental particles and their interactions, gravity excluded. Its formulation has not significantly changed since the mid-1970s. It incorporates the *Electroweak theory* which describes the weak and electromagnetic interactions, and *Quantum chromodynamics* (QCD) which describes the strong interaction. Elementary particles are classified in two main families, *fermions* and *bosons*. Fermions are the matter particles with spin- $\frac{1}{2}$ following the Fermi-Dirac statistics. Bosons are integer spin particles and follow the Bose-Einstein statistics. All of the particles have anti-particles of the same mass but opposite charge (some particles are their own anti-particle, such as the photon).

Fermions are divided in two groups, *quarks* and *leptons*. Each comes in six flavours as shown in figure 1.3. In the Standard Model, quarks cannot exist independently. They form bound states known as *hadrons*, essentially triplets or doublets. Triplets are called *baryons*, the most familiar of which are the nucleons, i.e. protons and neutrons. Doublets (quark/anti-quark pairs) are called *mesons*. More exotic forms of hadrons such as *tetraquarks* and *pentaquarks* have been discovered and can help improve our understanding of the strong interaction. Charged leptons include the *electron* and its more massive cousins, the *muon* and the *tau*. For each of them, there is an associated neutral lepton, a *neutrino*. Neutrinos interact only weakly with the other particles making them extremely difficult to detect. There are three generations of fermions. The fermions of the first generation

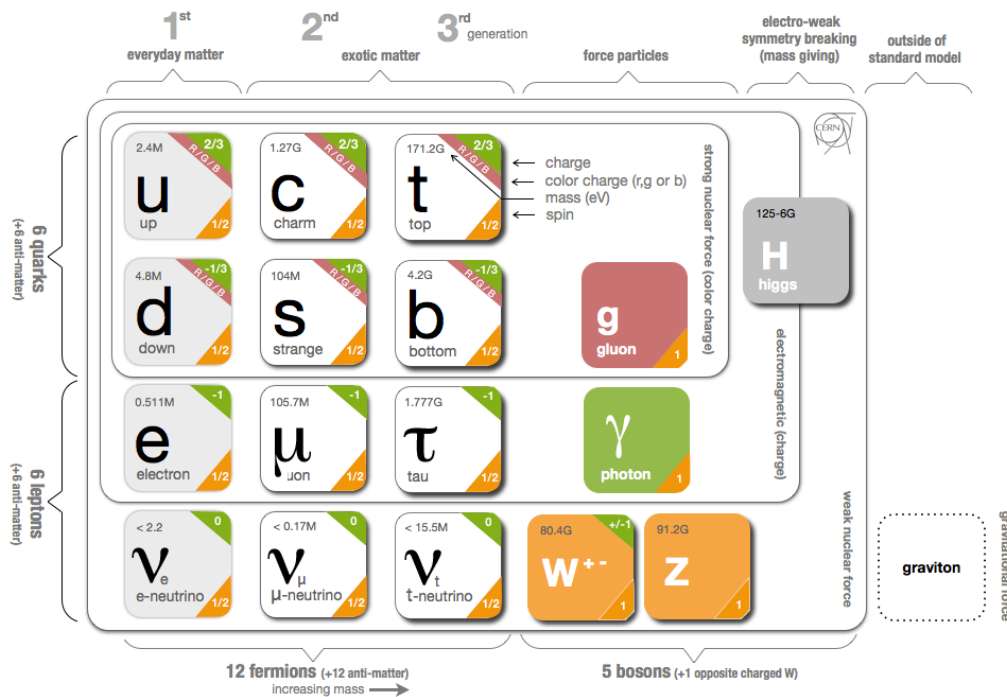


Figure 1.3 – A summary chart of the Standard Model particles and their properties. On the bottom right of the diagram, the hypothetical graviton is represented. It is supposed to be the mediator of gravitational interactions and is not currently in the Standard Model. *Credit: CERN.*

shown in the leftmost column of figure 1.3 are very stable and essentially combine to form most of the ordinary matter we are familiar with, while those of the second and third generations are unstable due to their excessive masses resulting in very short lifespans.

The force carriers of the three fundamental interactions of the Standard Model are the so-called *gauge bosons*. The *photon* is a massless, neutral and spin-1 particle that acts as the mediator of the electromagnetic interaction coupling to all the charged fermions. The *gluon* is also a massless, electrically neutral and spin-1 particle. In addition it carries another property called *color charge* related to strong interactions in QCD. Gluons act as mediators of the strong force and couple to quarks as well as other gluons through the color charge. The electrically charged W^\pm bosons, as well as the electrically neutral Z boson are the mediators of the weak interaction. Both have spin-1 but unlike the two previous gauge bosons, these are massive particles, about 100 times heavier than a proton, making the weak interaction only effective at a very short range ($\sim 10^{-18}$ m) at energy scales below the force carriers masses. In the Standard Model, particles acquire mass through the *Brout-Englert-Higgs mechanism*, which describes their interaction with a scalar field, the so-called *Higgs field*. The particle associated with this field, is a scalar boson (spin-0) called the *Higgs boson*, and was only discovered recently by the ATLAS and CMS experiments at the LHC [13, 136].

Although the Standard Model has been spectacularly successful in describing the world of particles, it is still incomplete. In particular, we still do not know how gravity fits in the picture. The nature of dark matter is also left unexplained. In addition, interpreting dark energy as the zero-point energy of the vacuum gives a large discrepancy with the observed value [468]. The model also provides no explanation for the observed matter-antimatter asymmetry. Furthermore, many more questions remain open, while some have possible theoretical solutions which are yet to be experimentally verified.

In this context, cosmology can bring unique insights on some of these problems, and help complement the picture we are building of the physics governing our Universe from the smallest scales to the largest ones. In the next paragraphs, we will go through the most relevant matter-energy components in cosmology, and briefly explore how cosmological observations can help improve our understanding of their properties.

Photons

Photons of the CMB are by far the most abundant species in the Universe with a number density around ~ 410 per cm^3 , and a blackbody spectrum measured with very high precision with an estimated absolute temperature $T = 2.72548 \pm 0.00057$ K [198]. Despite the exceeding abundance of photons, their very low temperature, at the present day, results in a marginal contribution to the total energy density budget of the Universe:

$$\Omega_\gamma h^2 \approx 2.47 \times 10^{-5}, \quad (1.29)$$

where h is the reduced Hubble constant defined as

$$h \equiv \frac{H_0}{100 \text{ km/s/Mpc}}. \quad (1.30)$$

Baryons

In cosmology, the term *baryons* is used in a loose sense meaning both nuclei and electrons, since most of the mass is in the former. This is the ordinary matter that compose the stars, and the diffuse gas in the interstellar as well as the intergalactic medium. The baryon density can be estimated either through the measurement of CMB anisotropies or inferred from the predictions of Big Bang Nucleosynthesis (BBN) when compared to the observed abundances of light elements. We will explore the two methods in greater details in the following section for BBN, and in Chapter 3 for CMB. As of the date of writing this thesis, the tightest constraints on the baryon density come from the Planck mission [373]:

$$\Omega_b h^2 = 0.02237 \pm 0.00015, \quad (1.31)$$

This amounts to roughly 5% of the total energy density content in the Universe. It is worth mentioning however that, for a long time, the observed amount of baryonic matter was adding up to only half of the value inferred from BBN and CMB, this was called the *missing baryon problem*. Theoretical results from cosmological hydrodynamic simulations indicated that the missing matter may be hiding in a web of hot and diffuse gas that is hard to detect, the so-called *warm-hot intergalactic medium* (WHIM) [131, 155]. Many independent observations using diverse techniques such as quasar spectroscopy [351], Sunyaev-Zel'dovich analyses [442, 158], or more recently the dispersion of localized fast radio bursts [323] provide strong support to the WHIM solution, with a baryon density measurement that is consistent with BBN and CMB estimates.

Neutrinos

Neutrinos were first postulated by W. Pauli in 1930 as an attempt to explain the continuous energy spectrum of the beta decay process which seemed at odds with energy conservation [109]. In the Standard Model, neutrinos interact via the weak interactions, and can couple to the electrically charged W^\pm bosons in the charged current (CC) interactions, or to the neutral Z bosons in the neutral current (NC) interactions. Each neutrino flavour, is defined by the specific charged lepton (electron, muon, or tau) that it produces or is produced with in the CC interactions. The number of flavours of these neutrinos can be constrained in the studies of Z bosons production via e^+e^- annihilation. The Large Electron-Positron collider (LEP) at CERN, and the Stanford Linear Collider (SLC) at SLAC, constrain this number, based on data from millions of Z boson decays, to $N_\nu = 2.9840 \pm 0.0082$ in agreement with the three observed generations of fermions [11]. Neutrinos are considered to be massless fermions in the minimal Standard Model. However observations of neutrino oscillations [206, 47] revealed that they have non-zero mass. The experimental bounds on their masses suggest that they are significantly low compared to the other fermions [219], which may warrant some explanation. Possible extensions of the Standard Model, could allow for the existence of the so-called, *sterile neutrinos*, which do not interact via any of the three interactions of the Standard Model. If established, their existence could be the key to solving many pending problems in neutrino physics, and more generally in particle physics. Most notably, it could explain the origin

of neutrino mass and the above noted disparity between the scale of neutrino masses and that of the other fermions, for instance, by invoking the *seesaw mechanism* associated with physics beyond the Standard Model [219]. Sterile neutrinos are also an essential ingredient in theoretical scenarios explaining the matter-antimatter asymmetry problem, such as *Leptogenesis* [207]. They also serve as a good dark matter candidate [167]. A related and equally important open question is whether neutrinos are *Dirac* or *Majorana* fermions. Contrary to other fermions of the Standard Model, neutrinos carry no electric charge and hence it is possible that they are their own antiparticle, in this case they are called Majorana fermions and described by only two degrees of freedom instead of the four required for Dirac fermions. If neutrinos are Majorana particles, lepton number would not be a conserved quantity, hence one way of determining the nature of neutrinos is to search for lepton number violating processes such as the neutrinoless double beta decay ($0\nu\beta\beta$) [219, 168], where a parent nucleus (A, Z) transforms into a daughter nucleus $(A, Z + 2)$ creating a pair of free electrons in the process: $(A, Z) \rightarrow (A, Z + 2) + 2e^-$, which is a forbidden nuclear transition in the Standard Model.

Any of these open questions in neutrino physics could provide a possible port of entry to physics beyond the Standard Model, so there are many experimental efforts carried out to push the frontiers of knowledge in these areas leveraging strong synergies between particle physics, nuclear physics, astrophysics and cosmology. In particular, determining neutrino masses is one such question where there is an excellent complementarity between cosmology and laboratory probes. Neutrinos do not conserve their flavours while propagating as a direct consequence of the presence of a mixing between their flavour eigenstates (ν_e, ν_μ, ν_τ) and their mass eigenstates (ν_1, ν_2, ν_3). More neutrino species, e.g. sterile neutrinos, could potentially be involved in the mixing but this possibility is not currently experimentally substantiated. Formally, the mixing is described as the following,

$$|\nu_l\rangle = \sum_{i=1}^3 U_{li} |\nu_i\rangle \quad \text{with } l = (e, \mu, \tau), \quad (1.32)$$

where U_{li} are the coefficients of the so-called *Pontecorvo-Maki-Nakagawa-Sakata* (PMNS) matrix [380]. The matrix is parametrized by three mixing angles ($\theta_{12}, \theta_{23}, \theta_{13}$) and one (δ) or three phases ($\delta, \alpha_1, \alpha_3$) depending on the Dirac or Majorana nature of the neutrinos. The full form of the PMNS matrix is given by

$$U = \begin{pmatrix} c_{12}c_{13} & s_{12}c_{13} & s_{13}e^{-i\delta} \\ -s_{12}c_{23} - c_{12}s_{23}s_{13}e^{i\delta} & c_{12}c_{23} - s_{12}s_{23}s_{13}e^{i\delta} & s_{23}c_{13} \\ s_{12}s_{23} - c_{12}c_{23}s_{13}e^{i\delta} & -c_{12}s_{23} - s_{12}c_{23}s_{13}e^{i\delta} & c_{23}c_{13} \end{pmatrix} \begin{pmatrix} 1 & 0 & 0 \\ 0 & e^{i\alpha_1/2} & 0 \\ 0 & 0 & e^{i\alpha_2/2} \end{pmatrix},$$

where $c_{ij} \equiv \cos \theta_{ij}$ and $s_{ij} \equiv \sin \theta_{ij}$. Neutrino oscillations experiments are able to set constraints on some of the parameters of the PMNS matrix as well as the mass splittings, that is the squared differences of neutrino masses $\Delta m_{ij}^2 \equiv m_i^2 - m_j^2$. The numbering of the neutrino mass eigenstates is arbitrary. By convention it is chosen such that we have $|\Delta m_{31}^2| > \Delta m_{21}^2 > 0$. The ordering of the spectrum of neutrino masses is not yet determined, and two non-equivalent configurations can be considered as shown in figure 1.4. In addition, a third ordering called *quasi-degenerate* (QD) where $m_1 \approx m_2 \approx m_3$ is also possible. The data shows

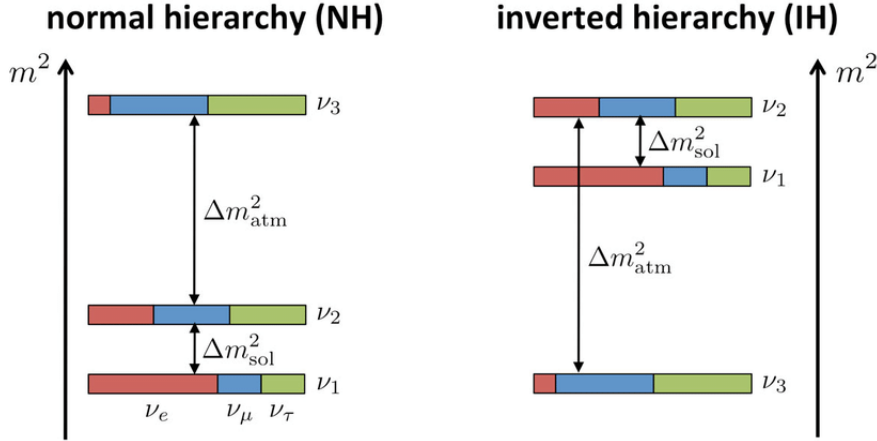


Figure 1.4 – Illustration of the two neutrino mass orderings. The colors represent the relative probability of finding the different flavour eigenstates in the corresponding mass eigenstates. $\Delta m_{\text{atm}}^2 \equiv |\Delta m_{32}^2|$ refers to the “atmospheric” mass splitting, and $\Delta m_{\text{sol}}^2 \equiv \Delta m_{21}^2$ refers to the “solar” mass splitting. *Credit: JUNO collaboration [368].*

that the mass splittings are ordered as $\Delta m_{21}^2 \ll |\Delta m_{31}^2| \simeq |\Delta m_{32}^2|$. By setting the lightest neutrino mass to zero, we are able to derive lower bounds on the neutrino masses from the data depending on the ordering scheme, as the following [219]

- **Normal Ordering (NO):** ($m_1 < m_2 \ll m_3$)
 $(m_1, m_2, m_3) \geq (0, \sqrt{\Delta m_{21}^2}, \sqrt{\Delta m_{32}^2 + \Delta m_{21}^2}) \sim (0, 8.6 \times 10^{-3}, 0.05) \text{ eV}.$
- **Inverted Ordering (IO):** ($m_3 \ll m_1 < m_2$)
 $(m_3, m_1, m_2) \geq (0, \sqrt{|\Delta m_{32}^2 + \Delta m_{21}^2|}, \sqrt{|\Delta m_{32}^2|}) \sim (0, 0.0492, 0.05) \text{ eV}.$
- **Quasi-degenerate spectrum (QD):** ($m_1 \approx m_2 \approx m_3$)
 $m_1 \simeq m_2 \simeq m_3 \gg \sqrt{|\Delta m_{32}^2|} \sim 0.05 \text{ eV}.$

To set upper bounds on the neutrino masses and therefore determine their mass scale, laboratory probes can measure the effective neutrino mass in beta decays given by $m_\beta \equiv \sqrt{\sum_i |U_{ei}|^2 m_i^2}$ or—if neutrinos are Majorana—the effective Majorana neutrino mass in experimental searches for neutrinoless double beta decays given by $m_{\beta\beta} = |\sum_i m_i U_{ei}^2|$. Furthermore, cosmological data is sensitive to the total mass of neutrinos $M_\nu \equiv \sum_i m_i$ and can bring stringent constraints with the potential caveat that the measurements are model dependent. In Chapter 3, we will explore in greater details how the total mass of neutrinos affects the observed CMB anisotropies. Currently the strongest constraints on m_β comes from Tritium beta decay in the KATRIN experiment which sets an upper limit [49, 50]

$$m_\beta < 1.1 \text{ eV} \quad (90\% \text{ CL}). \quad (1.33)$$

The experiment aims at reaching a sensitivity of $m_\beta < 0.2 \text{ eV}$ at the 90% confidence level in the future. The strongest limit from the neutrinoless double-beta decay searches comes from the KamLAND-Zen experiment [210], which sets a

lower limit on the $^{136}\text{Xe } 0\nu\beta\beta$ decay half-life $T_{1/2}^{0\nu} > 1.07 \times 10^{26}$ yr at the 90% confidence level and derives a range of the upper bound on the effective Majorana mass

$$\langle m_{\beta\beta} \rangle < 61 - 165 \text{ meV} \quad (90\% \text{ CL}). \quad (1.34)$$

The previously mentioned lower bounds from neutrino oscillations set a lower limit on the total neutrino mass, M_ν , of ~ 59 meV for the normal ordering and of ~ 100 meV for the inverted ordering. Data from Planck combined with Baryon Acoustic Oscillations (BAO) and Lyman- α [356] gives a stringent upper bound putting a marginal tension on the inverted mass ordering

$$M_\nu < 0.09 \text{ eV} \quad (95\% \text{ CL}). \quad (1.35)$$

The combination of data from large scale structures surveys such as DESI [185], Euclid [188], LSST [186], and SKA [9], or future CMB experiments such as CMB-S4 [187] combined with constraints from LiteBIRD [5], should allow a $3-4\sigma$ detection of the total neutrino mass even at the conservative lower bound allowed by neutrino oscillations, $M_\nu \sim 59$ meV, as concluded by multiple forecasting studies using different data combinations [54, 360, 68, 226, 108].

At early times, neutrinos behave as radiation and their contribution to the total radiation energy density is characterized by the *effective number of relativistic species*, N_{eff} , and is given by

$$\frac{\rho_\nu}{\rho_\gamma} = \frac{7}{8} N_{\text{eff}} \left(\frac{4}{11} \right)^{4/3}. \quad (1.36)$$

In the idealized case where neutrinos decouple instantaneously from the photon-baryon fluid, $N_{\text{eff}} = 3$. More careful analysis of the decoupling process taking into account flavour oscillations yields a small correction to this value, $N_{\text{eff}} = 3.045$ [159]. As we will see in more details in Chapter 3, N_{eff} has distinctive effects on cosmological observables, and on the CMB in particular. Current constraints from Planck and BAO give $N_{\text{eff}} = 2.99^{+0.34}_{-0.33}$ (95% CL) in agreement with theoretical expectations. In fact, departures from the standard value of N_{eff} would be the sign of additional free streaming radiation in the early Universe, and as such it serves as a good probe for light relics such as axions [274, 85], sterile neutrinos [15, 105], dark photons [25, 152] and many more [63, 469].

At late times, neutrinos become non-relativistic. We know from the temperature of the neutrino background today (~ 1.95 K as we will see in the next section) and the lower bounds on neutrino masses, that at least two neutrino species are non-relativistic today (the lightest neutrino could still be relativistic, but its contribution to the energy density would be negligible). As such they behave as matter, and their density parameter is given by

$$\Omega_\nu h^2 = \frac{M_\nu}{93.14 \text{ eV}}. \quad (1.37)$$

Their total average number density is $n_\nu^0 \approx 339 \text{ cm}^{-3}$.

Dark matter

Historically, dark matter was first hypothesized by the Swiss astronomer Fritz Zwicky in 1933, after observing a discrepancy between the mass of the Coma

galaxy cluster as derived from the galaxy velocities and that expected from their luminosity [481]. The hypothesis was later corroborated, first by the measurement of the rotation curves of spiral galaxies [399], and since then by a wealth of evidence from diverse observations (e.g. CMB, gravitational lensing, X-ray emission). Estimations of the baryon density from the CMB anisotropies and inferred from predictions of primordial nucleosynthesis show that it is much smaller than the total matter density and hence strongly point to the non-baryonic nature of dark matter. Furthermore, the search for Massive Compact Halo Objects (MACHOs) in the Milky Way through gravitational microlensing revealed that astrophysical compact objects do not significantly contribute to our galaxy's dark matter halo [337, 431]. Dark matter models can be classified in three categories according to their thermal velocities: hot (HDM), warm (WDM), and cold dark matter (CDM). HDM models are ruled out by galaxy clustering observations [203]. Below we list examples of some fairly popular dark matter candidates which are being considered:

- **Weakly Interacting Massive Particles (WIMPs):** these are particles that arise in theories that address issues with the Standard Model. The WIMPs framework has gained great popularity since it also provides an explanation of the observed relic abundance via the so-called “WIMP miracle”, namely that a particle with electroweak-scale mass and interacting via the weak interaction which are the expected properties of WIMP candidates in theoretical extensions of the Standard Model, gives the correct dark matter abundance. Several experiments are aiming at a direct detection of WIMPs through nuclear recoil induced by a WIMP elastically scattering off a target nucleus, but no confirmed detection has been made to date. Leading experiments include, among many others, the LUX experiment [51], the XENON1T experiment [62], the PandaX-II experiment [149], and the PICO-60 experiment [60].
- **Axions:** these particles were first introduced in the *Peccei–Quinn theory* [363, 362] which was postulated as a solution to the *strong CP problem*. The problem refers to the fact that QCD naturally gives rise to a term which breaks CP-symmetry, while experimental bounds on the neutron electric dipole moment do not indicate any CP violating effects and constrain the parameter characterizing the strength of the CP-violating term to very small values, thereby posing a fine-tuning problem [73, 148]. The theory proposes a new symmetry which suppresses the CP-violating terms. Axion-like particles (ALPs) also arise in many models of physics beyond the Standard Model such as string theory [437, 66]. Many experimental searches for ALPs rely on the *Primakoff effect* [384, 415] converting dark matter axions to photons in the presence of strong magnetic fields. These include *haloscopes* such as ADMX [106], *helioscopes* such as CAST [12] which target solar axions, and *light-shining through a wall* (LSW) experiments such as ALPS [175]. Other experiments like CASPER [71] are using nuclear magnetic resonance (NMR) techniques. Constraints are also drawn from a wide variety of astrophysical and cosmological probes [326]. One example is the cosmic birefringence effect through the Chern-Simons mechanism (more details are presented in the last section of Chapter 3).

- **Sterile neutrinos:** as was mentioned before inert neutrinos which do not couple with the Standard Model particles, are theoretically well motivated and also stand as good dark matter candidates. Hints of a possible sterile neutrino signal were reported by the Liquid Scintillator Neutrino Detector (LSND) [45], and the Mini boosted Neutrino Experiment (MiNiBooNE) [46], in addition to an observed reactor antineutrino anomaly [331]. However these findings are not totally understood and are at odds with other results such as the ones from the IceCube Neutrino Observatory which show no evidence for sterile neutrinos [145].
- **Primordial black holes (PBHs):** these black holes would be generated from overdensities in the early Universe and would therefore span a mass range that is not accessible by stellar black holes. The possible mass windows of PBHs are highly constrained, but they still stir some interest because they could potentially be the progenitors of some of the LIGO/Virgo mergers and contribute at least a fraction of dark matter [126]. Gravitational waves detectors can help set constraints on their parameter space [78, 140, 139]. Constraints can also be derived from CMB non-Gaussianity [476] and spectral distortions [358, 346].

Progress is being made to explore the parameter space of all these models, but the true nature of dark matter is yet to be determined. The only robust cosmological measurement of a dark matter parameter that can be made to date is its density. Assuming a CDM model, the Planck combined with BAO estimate [373] reads

$$\Omega_{\text{CDM}} h^2 = 0.11933 \pm 0.00091. \quad (1.38)$$

Dark energy

Very little is known about this last component. As was mentioned before, it is the manifestation of some repulsive gravity effect driving the observed accelerated expansion of the Universe [396, 366]. In the standard cosmological model (Λ CDM), it takes the form of Einstein's cosmological constant (Λ) [177], and the associated density parameter as measured by Planck and BAO is

$$\Omega_{\Lambda} = 0.6889 \pm 0.0056. \quad (1.39)$$

However this is merely one possibility in a wider class of models. Two main approaches are adopted to address the dark energy problem. The first one describes dark energy in terms of a time-varying scalar field, this class of models is called *quintessence dark energy* [388, 120]. The second approach is to consider that the accelerated expansion is the manifestation of gravitational effects which are not captured by general relativity (GR) at the largest scales. This led to the formulation of a class of modified gravity theories [298]. Recently gravitational waves detectors have been able to conduct stringent tests on these theories [19, 20, 22], e.g., setting constraints on the mass of the graviton. Bounds on the graviton mass can also be derived from solar system experiments [472, 93]. So far no evidence have been found of new physics beyond general relativity.

When searching for deviations from Λ CDM at the background level, a useful parametrization of dark energy is to allow for a time-varying equation of state as

in the following:

$$w(a) = w_0 + (1 - a)w_a, \quad (1.40)$$

where w_0 and w_a are assumed to be constants. In the latest results from Planck, such deviations were investigated in combination with external data and the values $w_0 = -1$ and $w_a = 0$, corresponding to a cosmological constant, are found to be close to the best fit. Furthermore, investigations of deviations from GR at the perturbation level also find consistent results with Λ CDM [373]. Future experiments will continue to push the parameter space limits on quintessence and modified gravity models using a variety of probes such as large-scale structures, type Ia supernovae, CMB, gravitational waves, time-delays from lensed quasar systems, galaxy clusters, redshift-space distortions and many more [419].

1.4 Physics of the Hot Big Bang

1.4.1 Thermodynamics

Before giving a thorough description of the thermal history of the Universe, we will need to introduce some basic concepts and equations that we will be relying on to understand the physics governing the sea of particles filling our Universe, both in conditions of thermal equilibrium and out of equilibrium.

Decoupling

A central element in understanding the thermal history of the Universe lies in the comparison between the rate of particle interactions Γ and the rate of expansion H . When $\Gamma \gg H$, particles interact much faster than they are able to experience the effect of expansion, hence they are able to reach local thermal equilibrium. However, as the Universe expands and cools the rate of interactions drops low enough that they become inefficient ($\Gamma \sim H$) and particles decouple from the thermal bath. As the rate of interactions depends on the specific properties of each particle, different particles decouple at different times.

Local thermal equilibrium

A solid observational argument that the early Universe was in a state of local thermal equilibrium is given by the perfect blackbody spectrum measured by COBE FIRAS [200]. In such conditions, the phase space distribution of particles is given by the Fermi-Dirac or Bose-Einstein statistics depending, respectively, on the fermion or boson nature of the particle

$$f(\mathbf{p}_i, T_i) = \frac{1}{e^{(E(\mathbf{p}_i) - \mu_i)/T_i} \pm 1}, \quad (1.41)$$

where the + sign is for fermions and the - sign is for bosons. \mathbf{p}_i are the momenta of the particle species i , T_i its temperature, μ_i its chemical potential which may also be temperature dependent, and its energy $E(\mathbf{p}_i)$ is given by $E(\mathbf{p}_i) = \sqrt{m_i^2 + \mathbf{p}_i^2}$, m_i being its mass.

The chemical potential is defined as the relative change of free energy U due to a change in the number of particles N , at fixed entropy S and volume V ,

$$\mu = \left(\frac{\partial U}{\partial N} \right)_{S,V}. \quad (1.42)$$

In a given reaction, chemical equilibrium is reached when the sum of chemical potentials of reacting particles is equal to the sum of the chemical potentials of the products. Whenever photon number is not conserved their chemical potential is zero, $\mu_\gamma = 0$. In the following, we consider that this is the case. We discuss deviations from $\mu_\gamma = 0$ when covering spectral distortions in Chapter 3. It follows then from considering particle-antiparticle annihilation that the chemical potential of a particle X and that of its antiparticle \bar{X} are related via $\mu_{\bar{X}} = -\mu_X$.

We can express the number density n_i , energy density ρ_i , pressure P_i and entropy density s_i of the different particle species as the following:

$$n_i = \frac{g_i}{(2\pi)^3} \int d^3\mathbf{p}_i f(\mathbf{p}_i), \quad \rho_i = \int dn_i E(\mathbf{p}_i), \quad P_i = \int dn_i \frac{\mathbf{p}_i^2}{3E(\mathbf{p}_i)} \quad (1.43)$$

$$\text{and } s_i = \frac{\rho_i + P_i - \mu_i n_i}{T_i}$$

where $dn_i = \frac{g_i}{(2\pi)^3} f(\mathbf{p}_i) d^3\mathbf{p}_i$ is the number density of the particle species i per unit volume of momentum space, and g_i is its number of degrees of freedom.

Relativistic Limit:

Particles are relativistic when their temperatures are much higher than their mass scales $T_i \gg m_i$. In addition, we will neglect chemical potentials $\mu_i \ll T_i$. This can be justified by relating the ratio, μ_i/T_i , to the net particle number: the inequality then holds for baryons because of the low baryon-to-photon ratio, and is also valid for electrons because of the electric neutrality of the Universe, however for neutrinos this is only a hypothesis which still needs to be proven. We can then simplify the integrals in Eq. (1.43), which gives

$$n_i = \frac{\zeta(3)}{\pi^2} g_i T_i^3 \quad \left(\times \frac{3}{4} \text{ for fermions} \right), \quad \rho_i = \frac{\pi^2}{30} g_i T_i^4 \quad \left(\times \frac{7}{8} \text{ for fermions} \right),$$

$$P_i = \frac{1}{3} \rho_i, \quad s_i = \frac{4}{3T_i} \rho_i. \quad (1.44)$$

We obtain the total radiation density by summing over the energy densities of all relativistic species, which we put in the following form

$$\rho_r = \frac{\pi^2}{30} \left(\sum_{i \in \text{bosons}} g_i \left(\frac{T_i}{T} \right)^4 + \frac{7}{8} \sum_{i \in \text{fermions}} g_i \left(\frac{T_i}{T} \right)^4 \right) T^4 \equiv \frac{\pi^2}{30} g_\star(T) T^4, \quad (1.45)$$

where T is the temperature of the photon gas, and $g_\star(T)$ is the *effective number of relativistic degrees of freedom* at the temperature T . Note that decoupled species could have different temperatures than the photon gas, $T_i \neq T$. When particle species become non-relativistic they are no longer counted in the sum.

Similarly, we can write the total entropy density in the radiation dominated era as

$$s = \frac{2\pi^2}{45} \left(\sum_{i \in \text{bosons}} g_i \left(\frac{T_i}{T} \right)^3 + \frac{7}{8} \sum_{i \in \text{fermions}} g_i \left(\frac{T_i}{T} \right)^3 \right) T^3 \equiv \frac{\pi^2}{30} g_{*S}(T) T^3, \quad (1.46)$$

where $g_{*S}(T)$ is the *effective number of relativistic degrees of freedom in entropy* at temperature T . Note that due to the different temperature dependence, g_* and g_{*S} are only equal when all relativistic species have the same temperature. As we will see later in this section, this condition is no longer satisfied after $t \sim 1\text{s}$ when electron-positron pair annihilation injects energy to the photon bath from which neutrinos have decoupled shortly before. A quick way to understand this process is to look at the conservation of entropy which implies that $g_{*S}(T) T^3 a^3 = \text{const.}$ As long as no particle undergoes a non-relativistic transition, g_{*S} remains constant and the temperature of the thermal bath scales as $T \propto a^{-1}$. Close to a non-relativistic transition however, g_{*S} varies and the thermal bath temperature scales as $T \propto g_{*S}^{-1/3} a^{-1}$, therefore decreasing slightly more slowly to compensate for the lost degrees of freedom.

Non-Relativistic Limit:

As the Universe expands and the particle species mean kinetic energies drop below their mass scales, they transition to the non-relativistic regime. In the limit $m_i \gg T_i$, we can write the number density, energy density, pressure and entropy density as the following

$$n_i = g_i \left(\frac{m_i T_i}{2\pi} \right)^{3/2} \exp \left[-\frac{(m_i - \mu_i)}{T_i} \right], \quad \rho_i = m_i n_i, \quad P_i = n_i T_i (\ll \rho_i) \quad (1.47)$$

and $s_i = \frac{(m_i + T_i - \mu_i) n_i}{T_i}$

If we neglect chemical potentials, then the above equations imply that the thermodynamic quantities of the non-relativistic particle species in thermal equilibrium are exponentially suppressed by the factor $\exp(-m/T)$ with respect to the relativistic species. This is the result of pair annihilation that is no longer balanced by pair production at low temperatures. We will see next that this suppression stops once equilibrium is disrupted, thereby freezing out the number densities of the particle species.

A non zero chemical potential would introduce an asymmetry between the number densities of the particles and their antiparticle counterparts. It is associated with the conservation of some quantum number, e.g. charge, baryon number, or lepton number. At early times, μ/T is roughly constant and of very small value given the tiny asymmetries in our Universe (the baryon-to-photon density ratio is of order $\sim 10^{-10}$) as temperature drops and annihilation proceeds to suppress the densities, the chemical potential would grow very quickly to balance the suppression in order to preserve the conservation of the associated quantum number.

Boltzmann equation and freeze out

In order to properly describe the statistical behavior of a species during decoupling, we should follow the evolution of its phase space distribution function,

$f(\mathbf{p}, T)$, which is no longer given by its equilibrium form (Eq. (1.41)). This evolution is governed by the *Boltzmann equation* which reads

$$\hat{\mathbf{L}}[f] = \mathbf{C}[f], \quad (1.48)$$

where $\hat{\mathbf{L}}$ is the *Liouville operator* and \mathbf{C} is the *collision operator*. In the FLRW Universe, spatial homogeneity and isotropy allows an equivalent parametrization of the phase space distribution function by energy and time, $f(E, t)$, and the Liouville operator is written as

$$\hat{\mathbf{L}}[f(E, t)] = E \frac{\partial f}{\partial t} - H \mathbf{p}^2 \frac{\partial f}{\partial E}. \quad (1.49)$$

Integrating over momentum space yields the following form of the Boltzmann equation, tracing the evolution of the number density

$$\frac{dn}{dt} + 3Hn = \frac{g}{(2\pi)^3} \int \mathbf{C}[f] \frac{d^3\mathbf{p}}{E}. \quad (1.50)$$

To simplify the collision term, we will consider only the interactions involving one single particle or two particles (given the rarity of interactions between three or more particles). These can be single-particle decays or two-particle scatterings or annihilation. Consider the following process

$$1 + 2 \rightleftharpoons 3 + 4 \quad (1.51)$$

The variations of the number density n_1 of species 1 will be driven by the difference between its annihilation rate with particle 2, $\langle\sigma v\rangle n_1 n_2$ (where $\langle\sigma v\rangle$ is the *thermally averaged cross section*) and its creation rate by particles 3 and 4, set such that the collision term vanishes in chemical equilibrium, given by $\langle\sigma v\rangle \left(\frac{n_1 n_2}{n_3 n_4}\right)_{eq} n_3 n_4$.

The Boltzmann equation is then given by the following

$$\frac{dn}{dt} + 3Hn = -\langle\sigma v\rangle \left[n_1 n_2 - \left(\frac{n_1 n_2}{n_3 n_4}\right)_{eq} n_3 n_4 \right]. \quad (1.52)$$

To qualitatively understand how the Boltzmann equation above drives the evolution of species beyond-equilibrium, we can rewrite it as a function of the number of particles in a comoving volume

$$\frac{d \ln N_1}{d \ln a} = -\frac{\Gamma_1}{H} \left[1 - \left(\frac{N_1 N_2}{N_3 N_4}\right)_{eq} \frac{N_3 N_4}{N_1 N_2} \right], \quad (1.53)$$

where $\Gamma_1 \equiv n_2 \langle\sigma v\rangle$. One can observe that when interactions are very efficient, $\Gamma_1 \gg H$, the system is driven towards the chemical equilibrium state. For instance, starting with $N_1 \gg N_1^{eq}$ ($\ll N_1^{eq}$) - keeping the numbers of the other particles at their equilibrium values -, the r.h.s. of the above equation is negative (positive) and the particle species 1 are destroyed (created) such that N_1 is driven towards N_1^{eq} . However, when the interaction rate drops below the expansion rate, $\Gamma_1 < H$, the r.h.s. of the equation vanishes and the number of particles in a comoving volume approaches a constant relic density, $N_1 = \text{const.}$ This means that reactions effectively cease and the number densities of particles are only diluted with the expansion of space. This process is known as *freeze out*, and it is central in understanding the abundances of the species populating our Universe today.

1.4.2 Thermal history of the Universe

In the following paragraphs, we apply the previously explored notions to derive a consistent picture of the evolution of the Universe throughout its 13.8 billion years history, as depicted in the standard model of cosmology. We will also extrapolate towards the possible scenarios that may apply in the very far future.

Planck Era

This is the earliest stage in the Universe history that we can identify. It lasts up to $\sim 10^{-43}$ s after the initial Big Bang singularity (as predicted by classical general relativity), and reaches temperatures higher than 10^{19} GeV. In this regime, quantum gravity effects are believed to be important, and our classical understanding of spacetime as provided by general relativity is no longer reliable. Hence, the attempts at describing this phase are all highly speculative. A proper description would require a complete quantum theory of gravity which we do not have.

Grand Unified Theory Era

This epoch lasts from the end of the Planck era at $\sim 10^{-43}$ s to $\sim 10^{-38}$ s, with temperatures beyond 10^{16} GeV. In this regime, the strong and electroweak interactions are indistinguishable. A class of theories called Grand Unified Theories (GUTs) attempt to describe the unification of these forces in a single framework. This is also highly speculative.

Electroweak phase transition

At around 100 GeV particles acquire mass after spontaneous symmetry breaking through the Higgs mechanism. W^\pm and Z bosons become massive, and the electroweak interaction splits into the electromagnetic and weak interactions.

QCD phase transition

When temperatures drop below ~ 100 MeV, the strong interactions between quarks and gluons become important, and they transition from a quark-gluon plasma to form hadrons, i.e. doublet (mesons) or triplet (baryons) bound states.

Neutrino decoupling and the $C\nu B$

Neutrinos are maintained in equilibrium with the thermal bath via weak interactions. The corresponding thermally averaged cross sections are of the order of $\langle\sigma v\rangle \sim G_F^2 T^2$, where $G_F \sim 10^{-5}$ GeV $^{-2}$ is the Fermi constant. It follows that the interaction rates are of the order of $\Gamma \sim G_F^2 T^5$. Hence, with the Hubble rate $H \sim T^2/M_{Pl}$, M_{Pl} being the Planck mass, we find

$$\frac{\Gamma}{H} \sim M_{Pl} G_F^2 T^3 \sim \left(\frac{T}{1 \text{ MeV}} \right)^3. \quad (1.54)$$

As the temperature drops below 1 MeV, neutrinos decouple from the thermal bath and freeze out. Today they form a sea of relic neutrinos called the *Cosmic Neutrino Background* ($C\nu B$). Their phase-space distribution at decoupling is given

by $f_D(p) = [\exp(p/T_D) + 1]^{-1}$ (neglecting chemical potentials), where we have defined p as the norm of the momentum, $p \equiv |\mathbf{p}|$, and T_D is the temperature at decoupling. The decoupled species are free falling in the FLRW Universe, and the geodesic equation shows that they are red shifted by the expansion as, $p \propto a^{-1}$. The subsequent evolution of the phase-space distribution is parametrized only by the ratio p/T_ν and is such that the number density is conserved, $dn_\nu \propto a^{-3}$. Consequently, the evolution of the phase-space distribution function $f_\nu(p, T_\nu) = dn/d^3\mathbf{p}$ is strictly equivalent to a shift in temperature as $T_\nu \propto a^{-1}$. Hence, as long as the number of relativistic degrees remain constant, the temperature of the neutrino background and that of the thermal bath scale in the same way and remain equal. However, as we will see next, the non-relativistic transition of electrons and positrons triggering their annihilation process will break this equality by “heating” the photon bath.

Electron-positron annihilation

When the temperature of the thermal bath reaches the electron and positron mass around 0.5 MeV, pair production is no longer balancing annihilation, and the number densities of electrons and positrons are exponentially suppressed until positrons disappear and a small proportion of electrons survive to ensure electric neutrality ($n_e = n_p \sim n_B \sim 10^{-10}n_\gamma$). As stated before, during this transition the temperature of the thermal bath changes as $T \propto g_{\star S}^{-1/3} a^{-1}$ to preserve entropy conservation. In fact, entropy is separately conserved for the decoupled neutrinos and the thermal bath. If we count the number of relativistic degrees of freedom in entropy of the thermal bath we have

$$g_{\star S}^{th} = \begin{cases} 2 + \frac{7}{8} \times 4 = \frac{11}{2} & T \gtrsim m_e, \\ 2 & T < m_e. \end{cases} \quad (1.55)$$

Hence, aT_γ increases after electron-positron annihilation by a factor $(11/4)^{1/3}$, while aT_ν remains constant. Hence the temperature of the neutrino background following the transition is slightly lower than the photon bath temperature as

$$T_\nu = \left(\frac{4}{11} \right)^{1/3} T_\gamma. \quad (1.56)$$

After electron-positron annihilation, photons are the only remaining relativistic species in the thermal bath, hence entropy conservation ensures that aT_γ stays exactly constant. After photon decoupling (which we discuss later in this section), the temperature of photons continues to scale as a^{-1} for the same reasons invoked previously for neutrinos, therefore the relation in Eq. (1.56) holds until today. $T_{\gamma,0} = 2.73$ K is the CMB temperature measured today [198]. We can therefore deduce the CνB temperature today $T_{\nu,0} = 1.95$ K. Plugging these numbers in the respective expressions of their number densities, Eq. (1.44), we can infer the following

$$\begin{aligned} n_{\gamma,0} &= \frac{2\zeta(3)}{\pi^2} T_{\gamma,0}^3 \approx 410 \text{ photons cm}^{-3}, \\ n_{\nu,0} &= \frac{3\zeta(3)}{4\pi^2} \times 6 \times T_{\nu,0}^3 \approx 339 \text{ neutrinos cm}^{-3}. \end{aligned} \quad (1.57)$$

The direct detection of the CνB is extremely challenging not only because neutrinos interact only weakly with ordinary matter, but also because of the very low kinetic energy of these relic neutrinos. Some experimental efforts are being deployed for this purpose however, such as the PTOLEMY experiment [95] which aim at detecting relic neutrinos via their capture on tritium. If successful, the PTOLEMY experiment would allow measurements of the neutrino mass scale and hierarchy, as well as enable important consistency checks of the CMB and BBN results, and provide an independent way by which we can study earlier times in cosmic history.

The relation in Eq. (1.56) also allows us to rewrite the effective number of relativistic degrees of freedom as

$$g_* = 2 + \frac{7}{8} \cdot 2 \cdot N_{\text{eff}} \left(\frac{4}{11} \right)^{4/3}, \quad (1.58)$$

leading to the parametrization of neutrinos, and potentially other light relics, contribution to the radiation energy density by N_{eff} as shown in Eq. (1.36) and discussed in its following paragraph.

Primordial nucleosynthesis

The synthesis of light elements starts shortly after the decoupling of neutrinos at around $t \sim 1$ s, and lasts for some minutes (roughly $t \sim 10^3$ s). The elements produced during this phase are mainly deuterium, D, and helium-4, ^4He , smaller quantities of helium-3, ^3He , and tritium, ^3H , along with traces of beryllium, Be, and lithium, Li. As mentioned before, the model describing the synthesis of these elements is called Big Bang Nucleosynthesis (BBN), and the remarkable agreement of its predictions with the observed abundances of light elements make it one of the pillars of the success of the Hot Big Bang model. In the following paragraphs we give a simplistic account of this process, to derive a prediction of ^4He abundance.

During the BBN phase, the timescales and densities are not sufficient for reactions involving three or more incoming nuclei to occur at any appreciable rate. Hence, the synthesis of deuterium in sufficient quantities is the key step to trigger a chain reaction that produces the heavier nuclei sequentially from lighter nuclei in two-particle reactions. Let us then first consider the reaction of deuterium formation:



In the range of temperatures considered here, the cross-section of the above reaction is sufficiently large to ensure chemical equilibrium, and neutrons, protons and deuterium are all non-relativistic. Since $\mu_\gamma = 0$, we have $\mu_n + \mu_p = \mu_D$, and we can write the following

$$\left(\frac{n_D}{n_n n_p} \right)_{eq} = \frac{3}{4} \left(\frac{m_D}{m_n m_p} \frac{2\pi}{T} \right)^{3/2} e^{B_D/T}, \quad (1.60)$$

where $B_D \equiv m_n + m_p - m_D = 2.22$ MeV is the binding energy of deuterium. By assuming $m_D \approx 2m_n \approx 2m_p$, and approximating the neutron density by the

baryon density, $n_n \sim n_B = \eta_B n_\gamma \sim \eta_B T^3$ with η_B the baryon-to-photon ratio, we can write

$$\left(\frac{n_D}{n_p}\right)_{eq} \sim \eta_B \left(\frac{T}{m_p}\right)^{3/2} e^{B_D/T}. \quad (1.61)$$

Given that $\eta_B \sim 10^{-10}$, the production of deuterium remains insignificant until the temperature drops well below the binding energy so that the $e^{B_D/T}$ term can compensate the small value of η_B . As long as there is no sizable quantity of deuterium produced, the production of other nuclei is also inhibited, this is known as the *deuterium bottleneck*. Thus, at temperatures above ~ 0.1 MeV baryons exist predominantly in the form of neutrons and protons. At around $T_{nuc} \sim 0.06$ MeV, we get $n_D \sim n_p$, this yields a rough estimate of the time at which the nucleosynthesis chain reaction gets started, $t_{nuc} \sim 330$ s. Given that the binding energy of helium is larger than that of deuterium, the reactions producing helium are energetically favored, with ${}^4\text{He}$ being the most stable configuration. Hence, below $T_{nuc} \sim 0.06$ MeV, ${}^4\text{He}$ nuclei are formed, each from two protons and two neutrons. However protons and neutrons are not necessarily present in equal proportions, a key quantity that sets the abundance of ${}^4\text{He}$ in the Universe is the neutron-to-proton ratio at the onset of nucleosynthesis.

To study the relative abundances of neutrons and protons, let us consider the following β -decay and inverse β -decay processes

$$\begin{aligned} n + \nu_e &\rightleftharpoons p^+ + e^-, \\ n + e^+ &\rightleftharpoons p^+ + \bar{\nu}_e. \end{aligned} \quad (1.62)$$

We neglect the chemical potentials of electrons and neutrinos, such that chemical equilibrium implies $\mu_n = \mu_p$. Using the expressions of the number densities at equilibrium in the non-relativistic limit Eq. (1.47) we can write

$$\left(\frac{n_n}{n_p}\right)_{eq} = \left(\frac{m_n}{m_p}\right)^{3/2} e^{-(m_n - m_p)/T} \simeq e^{-\mathcal{Q}/T}. \quad (1.63)$$

Where we have defined $\mathcal{Q} \equiv m_n - m_p = 1.30$ MeV. Therefore at $T \gg 1$ MeV, neutrons and protons are present in equal proportions. For $T < 1$ MeV, the neutrons start getting suppressed. However, since the above reactions are mediated by the weak interactions, this is also the time they start getting inefficient as we have seen previously for neutrino decoupling. Hence equilibrium is broken, and the subsequent evolution of neutron density is governed by the Boltzmann equation (1.52). If we define the neutron fraction as

$$X_n \equiv \frac{n_n}{n_n + n_p}, \quad (1.64)$$

a rough estimate of the freeze-out neutron fraction is [2]

$$X_n^\infty \sim \frac{1}{6}. \quad (1.65)$$

In addition we also need to take into account neutron decay which becomes important at timescales beyond $t \sim 100$ s. The neutron fraction then reads

$$X_n(t) = X_n^\infty e^{-t/\tau_n}, \quad (1.66)$$

where $\tau_n = 879.4 \pm 0.6$ s is the neutron lifetime [219]. Hence by substituting in the above expression $t_{nuc} \sim 330$ s we find $X_n(t_{nuc}) \sim 1/8$.

Finally, if we consider that all remaining neutrons at the onset of nucleosynthesis are processed into ${}^4\text{He}$, then we can derive a prediction for the final ${}^4\text{He}$ abundance. As we need two neutrons for each ${}^4\text{He}$ nucleus, then $n_{\text{He}} = 1/2 n_n(t_{nuc})$. The result is usually expressed as the mass fraction of helium

$$Y_p \equiv \frac{4n_{\text{He}}}{n_H} \simeq \frac{2X_n(t_{nuc})}{1 - X_n(t_{nuc})} \sim 2X_n(t_{nuc}) \sim 0.25. \quad (1.67)$$

The full treatment of primordial nucleosynthesis requires the tracking of a large set of reactions, solving many coupled equations. This is performed by using numerical codes. The plot on the left of figure 1.5 shows the evolution of the abundances of nuclei from hydrogen to lithium. The right plot of figure 1.5 shows a comparison between BBN predictions and observational constraints. The baryon-to-photon ratio is constrained precisely from the observed abundance of deuterium and from CMB measurements, the observed abundance of ${}^4\text{He}$ is in agreement with the previous constraints, and the overall concordance of the constraints is remarkable, except for lithium. The latter discrepancy has been called the *cosmological lithium problem* [151, 219]. It remains unresolved to date, although several possible solutions have been proposed. In general the mismatch can be traced back to three possible sources of errors, it can arise from systematic errors in the astrophysical measurements, errors in the nuclear inputs, or—a more intriguing possibility—new physics could be at work.

Recombination

At the end of primordial nucleosynthesis, the baryonic matter and radiation contents of the Universe are in the form of a plasma of non-relativistic electrons, essentially hydrogen and helium nuclei, and photons in addition to the decoupled neutrinos. Well below the MeV scale, equilibrium between electrons, nuclei and photons is ensured by electromagnetic interactions, i.e. Compton scattering ($e^- + \gamma \rightleftharpoons e^- + \gamma$) for the coupling between electrons and photons, and Coulomb scattering ($e^- + p^+ \rightleftharpoons e^- + p^+$ or $e^- + {}^4\text{He} \rightleftharpoons e^- + {}^4\text{He}$) for the coupling between electrons and nuclei. As temperatures drop sufficiently, electrons begin to combine with nuclei to form neutral atoms, and the density of free electrons start falling very quickly. This process is referred to as *recombination*. The exact treatment of recombination takes into account the fact that neutral atoms can form in different excited states, and then relax to their fundamental states while emitting photons. In the following, we consider a simplified approach where we track the formation of neutral hydrogen only and directly in its fundamental state.

The reaction of neutral hydrogen formation is:



When temperatures are high enough to ensure equilibrium we have $\mu_p + \mu_e = \mu_H$, using the expression of number densities in the non-relativistic limit (Eq. (1.47)) and further assuming the electric neutrality of the universe which gives us $n_e = n_p$, we can write

$$\left(\frac{n_H}{n_e^2} \right)_{eq} = \left(\frac{2\pi}{m_e T} \right)^{3/2} e^{B_H/T}, \quad (1.69)$$

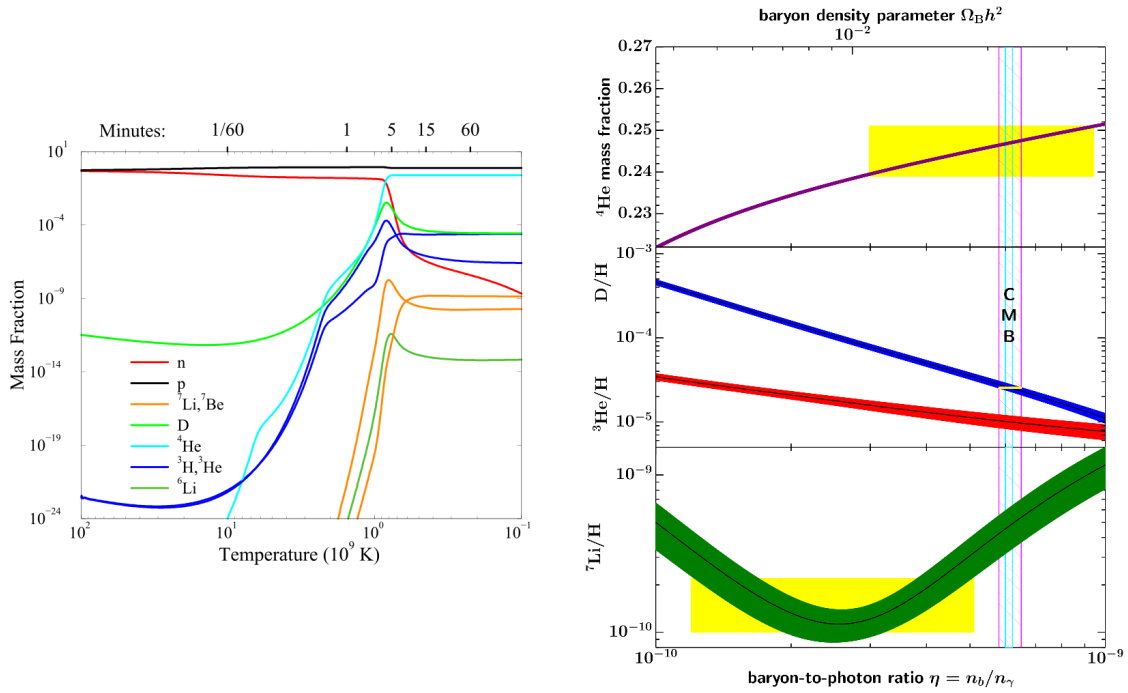


Figure 1.5 – *Left*: Evolution of the abundances of light elements as a function of the age of the Universe (in minutes) and of temperature. *Credit*: S. Burles et al. [115]. *Right*: Standard model BBN predictions of the abundances of ${}^4\text{He}$, D, ${}^3\text{He}$, and ${}^7\text{Li}$ - the bands show the 95% CL range. The yellow boxes represent the observed abundances. The vertical bands correspond to the 95% CL of respectively the CMB measurement of the baryon-to-photon ratio (narrow band) and the BBN D+ ${}^4\text{He}$ concordance range (wider band). *Credit*: Particle Data Group [219].

where $B_H \equiv m_p + m_e - m_H = 13.6$ eV is the binding energy of hydrogen. We define the free electron fraction as

$$X_e \equiv \frac{n_e}{n_B}, \quad (1.70)$$

where n_B is the baryon density. By ignoring other nuclei other than hydrogen we have $n_B \approx n_p + n_H = n_e + n_H$, which yields

$$\frac{1 - X_e}{X_e^2} = \frac{n_H}{n_e^2} n_B = \frac{n_H}{n_e^2} \eta_B n_\gamma. \quad (1.71)$$

Substituting in the above equation, the relation obtained in Eq. (1.69) and the expression of the photon number density we derive the so-called *Saha equation*,

$$\left(\frac{1 - X_e}{X_e^2} \right)_{eq} = \frac{2\zeta(3)}{\pi^2} \eta_B \left(\frac{2\pi T}{m_e} \right)^{3/2} e^{B_H/T}. \quad (1.72)$$

By considering that the process remains in the Saha equilibrium as the fraction of free electrons drops by 90% we find a corresponding temperature

$$T_{rec} \approx 0.3 \text{ eV} \approx 3600 \text{ K}. \quad (1.73)$$

The recombination temperature is substantially lower than the binding energy of hydrogen due to the very low baryon-to-photon ratio. In fact, because photons are much more abundant than baryons, even when the mean kinetic energy of the photon gas is below the binding energy of hydrogen, there are still enough photons in the high energy tail of the distribution to dissociate the newly forming hydrogen atoms. Given the measured CMB temperature today, we can derive the redshift and the time of recombination in the standard cosmological model,

$$\begin{aligned} z_{rec} &\approx 1320, \\ t_{rec} &\approx 290\,000 \text{ yrs.} \end{aligned} \tag{1.74}$$

Photon decoupling and the CMB

The interaction rate of the scattering of photons off electrons is given by

$$\Gamma_\gamma = n_e \sigma_T, \tag{1.75}$$

where $\sigma_T \approx 2 \cdot 10^{-3} \text{ MeV}^{-2}$ is the Thomson cross section. Photons decouple from the plasma roughly at temperature T_{dec} when $\Gamma_\gamma(T_{dec}) \sim H(T_{dec})$. Using the Saha equation, Eq. (1.72), we find

$$T_{dec} \approx 0.27 \text{ eV.} \tag{1.76}$$

Even though the temperature does not decrease substantially between recombination and decoupling, the ionization fraction drops by about one order of magnitude from 10% to 1%. At this stage, the Universe becomes transparent, and the photons start streaming freely, forming what we call today the Cosmic Microwave Background (CMB) radiation. The redshift and time of decoupling are

$$\begin{aligned} z_{dec} &\approx 1100, \\ t_{dec} &\approx 380\,000 \text{ yrs.} \end{aligned} \tag{1.77}$$

We derived these results assuming that the Saha equation holds throughout the recombination and decoupling processes. This is only an approximation, and an exact treatment requires the tracking of the free electron fraction outside of equilibrium through the Boltzmann equation as shown in figure 1.6. In fact, one finds a freeze-out ionization fraction of order $X_e^\infty \sim 5 \cdot 10^{-4}$.

Dark Ages & Reionization

After photon decoupling, space continues expanding and gradually redshifting the background radiation. As these photons are shifted out of the visible spectrum and no other light source exists in the early Universe, apart from the 21-cm line of neutral hydrogen, the Universe becomes dark and this period has been conveniently named the *Dark Ages*. This period lasts for a few hundred million years.

As the first structures started emerging, the ultraviolet radiation from the newly formed stars and quasars produce ionized bubbles which get larger over time, with the growth of the ionizing sources within them. They eventually overlap and the whole Universe is ionized again [208]. This process is referred to as

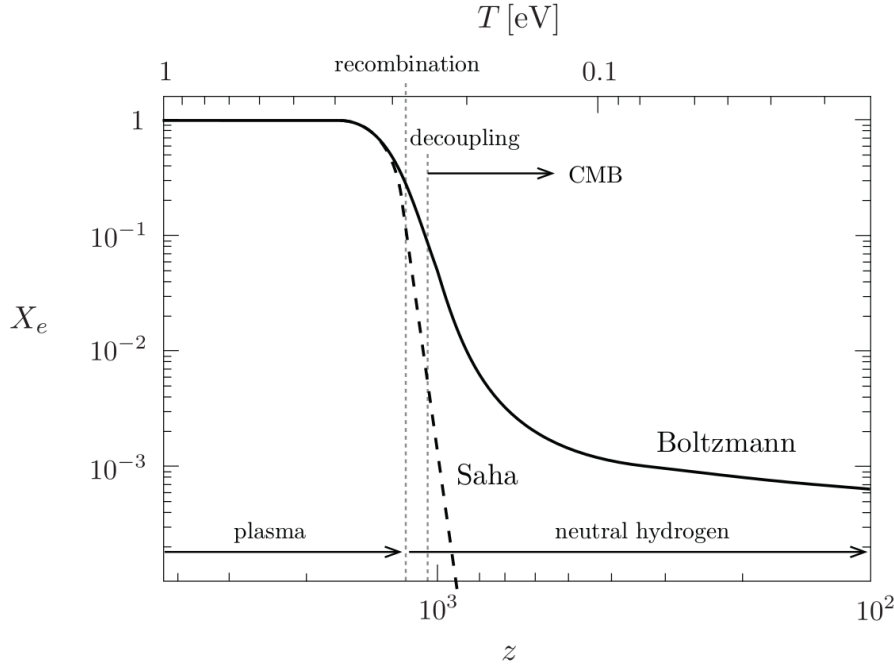


Figure 1.6 – Evolution of the ionization fraction as a function of redshift and temperature. *Credit: D. Baumann [2].*

reionization. One way to put constraints on reionization models is through the observation of the spectra of distant quasars. In particular, the observation of the so-called *Gunn-Peterson trough* [220], detected for the first time from data taken by the Sloan Digital Sky Survey (SDSS), indicates that the Universe reaches the final stage of reionization at around redshift $z \sim 6$ [88]. Analysis of the Lyman- α forest can also provide complementary constraints. CMB measurements are also sensitive to reionization through the Thomson scattering of photons off the free electrons, and can provide powerful constraints on the integrated path of ionized hydrogen to the last scattering surface, we will revisit the effects of reionization on the CMB in greater details in Chapter 3. Finally, a promising new technique that could unveil more insights about this epoch is the measurement of the *spin-flip* 21-cm line of neutral hydrogen [340, 104].

Matter- Λ equality & the present epoch

As we have seen in the second section of this chapter, the energy density of matter scales as a^{-3} , while that of dark energy is constant. With enough time, dark energy becomes inevitably the dominant energy content in the Λ CDM universe. In fact, in our Universe the transition to dark energy domination happened fairly recently at redshift $z \simeq 0.3$. One may ask why we live in an epoch where the matter and dark energy densities are of the same order of magnitude. This is called the *cosmological coincidence problem* and has motivated many models with dynamical dark energy or introducing interactions in the dark sector [458].

Future scenarios

Current projections from the standard model of cosmology, given state-of-art measurements, favor a *heat death* scenario in which the Universe continues

its accelerated expansion indefinitely. Stars will eventually cease to form as the gas supplies will be exhausted over time, and as the last existing stars eventually burn out, matter will be mostly locked in *degenerate states* such as the ones in white dwarfs and neutron stars, which will in turn slowly cool down. Eventually the Universe will get plunged into darkness again. Black holes which will dominate the Universe in the very far future will also evaporate over time, emitting photons, neutrinos and other particles through Hawking radiation [236, 237]. The remaining photons will get redshifted to colossal wavelengths and matter will be extremely diluted with the continued expansion of space towards the infinite future, as the temperatures asymptotically converge towards a minimum value corresponding to the maximum entropy state of the Universe.

Other future scenarios may also be possible. For example, if dark energy is not a cosmological constant and is described by an equation of state parameter $w < -1$ as predicted by the so-called *phantom dark energy* models [122], its energy density will increase over time ripping apart all possible structures in a finite time. This scenario is referred to as the *Big Rip*. Alternatively, the Universe could also experience a re-collapse, the so-called *Big Crunch*, which can be due, for instance, to a negative cosmological constant currently masked by an additional quintessence-like dark energy component driving the current acceleration, but which would become dominated by the negative cosmological constant at later times [119]. Big Crunch scenarios are also a possibility in the framework of supergravity and string theory [271, 428]. Ultimately, our knowledge about the nature of dark energy and the inner workings of our Universe beyond the standard model remain widely uncertain, and any projections of the ultimate fate of the Universe can only be speculative. Hence, the above scenarios should only be considered as pure intellectual curiosities.

1.5 Open questions of the standard model

As we have seen so far the Λ CDM model provides a consistent picture of the physics governing the Universe as a whole, from the earliest times to the present day. It has been corroborated by many independent observations from multiple probes. Its success in this regard allows one to be confident that the highlights of this picture are correct. However, it is likely not the full story. We still lack knowledge about the true nature of the dark sector and its possible interactions. Possible extensions of the standard model of particle physics, such as to explain the matter-antimatter asymmetry for instance, may also have cosmological implications. In addition, the standard model does not provide a natural mechanism that generates the initial perturbations seeding the density fluctuations we see in the CMB, and that later collapse gravitationally to form the structures that we observe. Furthermore, certain features of our Universe seem to require extreme fine tuning of initial conditions, and are thus left unexplained. In the following paragraphs, we will focus on two such fine-tuning problems.

Flatness problem

One can write the Friedmann Equation, Eq. (1.27) in the following form

$$1 - \Omega(a) = \frac{-k}{(aH)^2}, \quad (1.78)$$

with $\Omega(a) \equiv \frac{\rho(a)}{\rho_c(a)}, \quad \rho_c \equiv \frac{3H(a)^2}{8\pi G}.$

We know that for any dominating fluid with $w > -1/3$, the comoving Hubble radius $(aH)^{-1} = \dot{a}^{-1} \propto t^{\frac{1+3w}{3(1+w)}}$ grows with time. Hence, any small deviation from zero spatial curvature should get larger with time. The latest constraints from *Planck* and BAO, are consistent with a spatially flat Universe to a 1σ accuracy of 0.2% [373]. This suggests, that spatial curvature should have been extremely fine tuned at earlier times. For instance, one requires that in the Planck era

$$|\Omega(a_{\text{Planck}}) - 1| \lesssim \mathcal{O}(10^{-61}). \quad (1.79)$$

No natural explanation is provided for this fact in the standard model.

Horizon problem

We can define the *comoving particle horizon*, $\chi_p(t)$, as the maximum distance that light can travel between time 0 and t

$$\chi_p(t) = \tau - \tau_i = \int_0^t \frac{dt'}{a(t')}, \quad (1.80)$$

where τ is the *conformal time*. For a Universe dominated by a fluid with $w > -1/3$, the comoving particle horizon grows with time as

$$\chi_p(t) \propto a^{\frac{1}{2}(1+3w)}. \quad (1.81)$$

This means that scales entering our horizon today in the CMB, have not been in causal contact before. In fact, the CMB counts around $\sim 10^5$ causally disconnected regions [83]. Therefore, the homogeneity of the CMB radiation, which exhibits temperature fluctuations only within one part in 10^5 across the whole sky also seems unnatural in the framework of the standard model. An illustration of the problem is presented in figure 1.7.

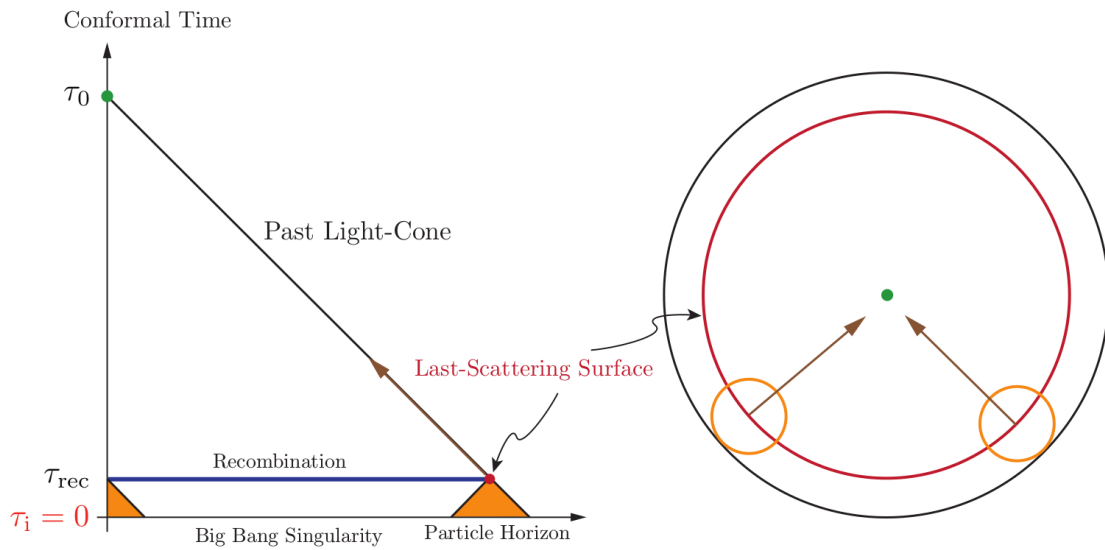


Figure 1.7 – Conformal diagram illustrating the horizon problem in the framework of the standard Hot Big Bang cosmology. The past light cones of two regions on the last scattering surface, separated by 90° , do not intersect. This implies that the two regions are causally disconnected, which is at odds with the observed homogeneity of the CMB. *Credit: D. Baumann [83].*

Inflation

Contents

2.1	Addressing the cosmic puzzles	35
2.2	Inflationary dynamics	36
2.3	Seeding structure in the Universe	39
2.3.1	Metric perturbations	40
2.3.2	Matter perturbations	41
2.3.3	Quantum origin of structures	44
2.4	Observational evidence and constraints	46

In the late 1970s and early 1980s, theoretical physicists Alan Guth, Alexei Starobinsky, Andre Linde and many others developed an elegant solution to the fine tuning problems exposed in the previous chapter. In fact, Guth was working on a solution to the so-called *monopole problem*. The latter stipulates that, *magnetic monopoles*, hypothetical elementary particles with net “magnetic charge” (in analogy with electric charge) predicted by Grand Unified Theories (GUTs), should have been produced abundantly at the high temperatures of the primordial Universe, and be the dominant constituents at the present time [480, 382] contrary to observations. Guth proposed that the early Universe was trapped in a false vacuum state with a high energy density, behaving in a similar fashion to a cosmological constant. He coined the term “inflation”, and it was immediately realized that the resulting exponential expansion of this phase, would not only solve the hypothetical monopole problem, but also the flatness and horizon problems [221]. Soon after, it was demonstrated that inflation also provides an elegant account of the origin of structures through the quantum fluctuations of the inflaton [342, 238, 426, 223, 76]. As of today, many inflationary models have been elaborated, and the inflationary paradigm is given high credence in the scientific community thanks to a strong body of evidence established by decades of cosmological observations. However, controversies remain [262, 263, 222] and a direct signature of this early phase of expansion of the Universe is yet to be observed. The most promising piece of evidence that could amount to a direct

observation of this early phase of the Universe, is the detection of primordial gravitational waves that are believed to be generated by inflation. This would not only strengthen the large body of evidence supporting the theory, but it would also enable a detailed understanding of the physics governing this phase as it would allow us to constrain specific parameters of inflation models, and open up a new window on GUT scale physics, which remain poorly understood at this stage, given that its corresponding energies are orders of magnitude above what is accessible in state-of-the-art particle accelerators. This could, in turn, provide unique insights to guide developments on theoretical physics such as GUTs and theories of quantum gravity.

In this chapter, we will explore how inflation naturally solves the flatness and horizon problems, before we give a detailed account of the physics of a specific class of inflation models called *slow-roll inflation*. We will characterize the statistics of primordial perturbations, and discuss how the state-of-the-art observations support the inflationary paradigm. This chapter is broadly based on lecture notes from [83].

2.1 Addressing the cosmic puzzles

A key underlying feature behind the fine tuning issues that arise in the standard model, is the fact that the comoving Hubble radius $(aH)^{-1}$ grows with time. In fact, as we have seen in Chapter 1, the deviation from flatness is driven by the decreasing $(aH)^{-2}$ factor as we go back in time boosting the amount of fine tuning required to fit the observations. Likewise, the presence of multiple seemingly causally disconnected regions in the last scattering surface is due to the growing comoving Hubble radius; formally one can write the comoving particle horizon in the following form

$$\chi_p(t) = \int_0^a \frac{d \ln a'}{a' H(a')}, \quad (2.1)$$

it becomes then apparent that for an increasing comoving Hubble radius, the contributions to the above integral are dominated by late times. Hence we have $\chi_p(t) \sim (aH)^{-1}$, meaning that regions outside of the comoving Hubble radius have never been in causal contact before.

From the previous observations, it follows that a simple idea to resolve the fine tuning issues is to reverse the behavior and have a shrinking comoving Hubble radius. The flatness problem is then naturally solved as can be derived from Eq. (1.78) for which the solution $\Omega(a) = 1$ becomes an attractor. The horizon problem is also solved by the same mechanism. As the comoving Hubble radius is shrinking, contributions to the integral, Eq. (2.1), from early times become significant, making the particle horizon grow way beyond the comoving Hubble radius. Hence regions of the last scattering surface laying outside the comoving Hubble radius could have been in causal contact in the past, thus explaining the observed homogeneity at superhorizon scales. An illustration of the problem's resolution is given in figure 2.1.

One might ask under what conditions do we have a shrinking comoving Hubble radius? In other terms, we should find what conditions on other quantities of interest are equivalent to $\frac{d}{dt}(aH)^{-1} < 0$. The answer stems directly from the

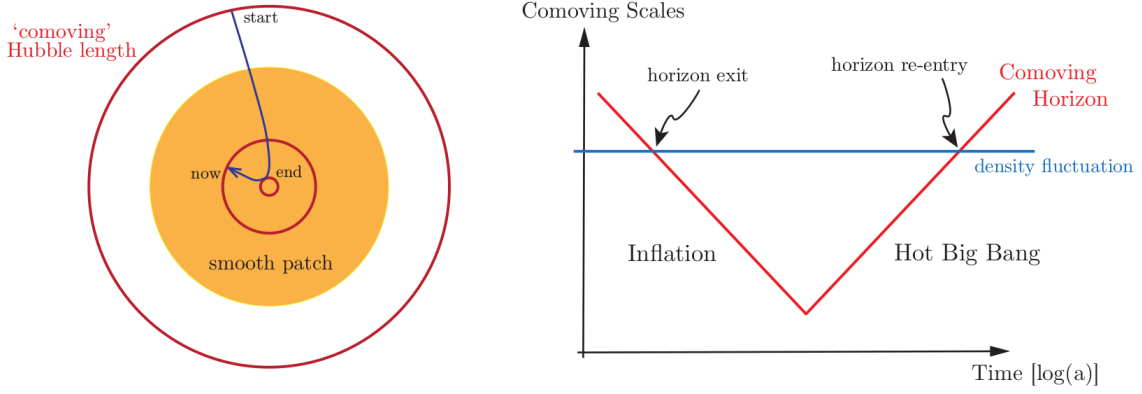


Figure 2.1 – At early times, regions separated by a given cosmological scale of interest, indicated by the blue horizontal line, are causally connected. This scale exits the horizon during the decreasing phase of the comoving Hubble radius.

As the standard Big Bang expansion unfolds, the comoving Hubble radius increases until it reaches again this scale at late times. *Credit: D. Baumann [83].*

following relation

$$\frac{d}{dt}(aH)^{-1} = \frac{-\ddot{a}}{(aH)^2}. \quad (2.2)$$

Hence the condition is equivalent to having an accelerated expansion

$$\frac{d}{dt}(aH)^{-1} < 0 \iff \ddot{a} > 0. \quad (2.3)$$

From the second Friedmann equation, Eq. (1.20), we derive that such conditions can only be generated by a negative pressure fluid satisfying

$$p < -\frac{1}{3}\rho. \quad (2.4)$$

2.2 Inflationary dynamics

In the following, we give a more quantitative account of the physics governing the inflationary phase in the simplest case where the accelerated expansion is governed by a single scalar field, the so-called *inflaton*, rolling a generic potential $V(\phi)$ as shown in figure 2.2.

Although the nature of the inflaton field is a subject of active research, e.g. [96, 343, 202, 211, 267], we will not pursue this question in this chapter. The dynamics of the system, consisting of a scalar field minimally coupled to gravity is given by the action, S , corresponding to the sum of the Einstein-Hilbert action and the action of scalar field with canonical kinetic term

$$S = \int d^4x \sqrt{-g} \left[\frac{1}{2}R + \frac{1}{2}g^{\mu\nu} \partial_\mu \phi \partial_\nu \phi - V(\phi) \right], \quad (2.5)$$

with g the determinant of the metric tensor, and R the Ricci scalar. The energy-momentum tensor for the scalar field can be written as the following

$$T_{\mu\nu}^{(\phi)} \equiv -\frac{2}{\sqrt{-g}} \frac{\delta S_\phi}{\delta g^{\mu\nu}} = \partial_\mu \phi \partial_\nu \phi - g_{\mu\nu} \left(\frac{1}{2} \partial^\sigma \phi \partial_\sigma \phi + V(\phi) \right). \quad (2.6)$$

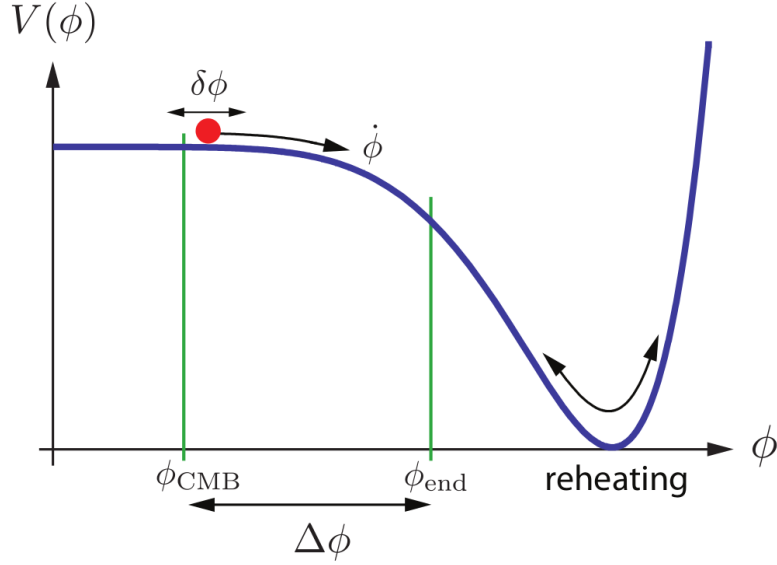


Figure 2.2 – An illustration of the motion of the inflaton field rolling a generic potential $V(\phi)$. The field starts with a dominating potential energy, as the field rolls down the potential, the Universe is inflating, and the kinetic energy grows until it becomes no longer negligible $\frac{1}{2}\dot{\phi}^2 \approx V$. Inflation stops and the inflaton decays converting its energy density into radiation in a process called *reheating*.

Credit: D. Baumann [83].

The equation of motion is

$$\frac{\delta S_\phi}{\delta \phi} = \frac{1}{\sqrt{-g}} \partial_\mu (\sqrt{-g} \partial^\mu \phi) + V_{,\phi} = 0, \quad (2.7)$$

where $V_{,\phi} = \frac{dV}{d\phi}$. In the case of a homogeneous field $\phi(t, \mathbf{x}) \equiv \phi(t)$, and under the assumption of FLRW metric, $T_{\mu\nu}^{(\phi)}$ takes the perfect fluid form, with density ρ_ϕ and pressure p_ϕ given by

$$\begin{aligned} \rho_\phi &= \frac{1}{2} \dot{\phi}^2 + V(\phi), \\ p_\phi &= \frac{1}{2} \dot{\phi}^2 - V(\phi). \end{aligned} \quad (2.8)$$

It follows that the equation of state parameter, w_ϕ , reads

$$w_\phi \equiv \frac{p_\phi}{\rho_\phi} = \frac{\frac{1}{2} \dot{\phi}^2 - V(\phi)}{\frac{1}{2} \dot{\phi}^2 + V(\phi)}. \quad (2.9)$$

When the potential energy $V(\phi)$ dominates over the kinetic energy $\frac{1}{2}\dot{\phi}^2$, the scalar field leads to a negative pressure and accelerated expansion.

Under the above assumptions (homogeneity and FLRW metric), the dynamics of inflation is determined by the following two equations, derived from Eq. (2.7) and the Friedmann equation

$$\begin{aligned} \ddot{\phi} + 3H\dot{\phi} + V_{,\phi} &= 0 \\ H^2 &= \frac{1}{3} \left(\frac{1}{2} \dot{\phi}^2 + V(\phi) \right). \end{aligned} \quad (2.10)$$

We define the slow-roll parameter, ε , as

$$\varepsilon \equiv -\frac{\dot{H}}{H^2} = -\frac{d \ln H}{dN} \quad (2.11)$$

with $dN = H dt$. We will comment on the meaning of the quantity N later, but for the moment, one can relate the acceleration of expansion to this parameter, ε , by observing that by definition we have

$$\frac{\ddot{a}}{a} = H^2(1 - \varepsilon), \quad (2.12)$$

which means that we obtain an accelerated expansion for $\varepsilon < 1$. We can then relate ε to the equation of state parameter, w_ϕ , via the second Friedmann equation in the case of a universe dominated by the inflaton as follows

$$\varepsilon = \frac{3}{2}(w_\phi + 1) = \frac{1}{2} \frac{\dot{\phi}^2}{H^2}. \quad (2.13)$$

In the de Sitter limit, $w_\phi \rightarrow -1$, we get $\varepsilon \rightarrow 0$. In these conditions, we have $\dot{\phi}^2 \ll V(\phi)$. To sustain the accelerated expansion long enough we need to have the second time derivative of ϕ small enough, and thus according to the first line in Eq. (2.10)

$$|\ddot{\phi}| \ll |3H\dot{\phi}| \quad \text{and} \quad |\ddot{\phi}| \ll |V_{,\phi}|. \quad (2.14)$$

This leads to another condition imposed on a second slow-roll parameter, η , defined as

$$\eta = -\frac{\ddot{\phi}}{H\dot{\phi}} = \varepsilon - \frac{1}{2\varepsilon} \frac{d\varepsilon}{dN}, \quad (2.15)$$

where the condition reads $|\eta| \ll 1$. The two conditions $\varepsilon, |\eta| \ll 1$, form the so-called *slow-roll approximation* (SRA). Under these conditions, the background evolution is governed by the following equations of motion

$$\begin{aligned} H^2 &\approx \frac{1}{3} V(\phi) \approx \text{const.}, \\ \dot{\phi} &\approx -\frac{V_{,\phi}}{3H}, \end{aligned} \quad (2.16)$$

and the time evolution of the scale factor approaches the de Sitter solution

$$a(t) \sim e^{Ht}. \quad (2.17)$$

We can also express the SRA as a set of (approximate) conditions on the shape of the potential $V(\phi)$, for this purpose we introduce the dimensionless *potential slow-roll parameters* ϵ_V and η_V defined as follows

$$\begin{aligned} \epsilon_V(\phi) &\equiv \frac{M_{\text{Pl}}^2}{2} \left(\frac{V_{,\phi}}{V} \right)^2, \\ \eta_V(\phi) &\equiv M_{\text{Pl}}^2 \frac{V_{,\phi\phi}}{V}, \end{aligned} \quad (2.18)$$

where M_{Pl} is the Planck mass set to 1 in the following. When the SRA holds, we should have $\epsilon_V, |\eta_V| \ll 1$. In addition, we can also write

$$\varepsilon \approx \epsilon_V, \quad \eta \approx \eta_V - \epsilon_V. \quad (2.19)$$

Note that the slow-roll conditions on the potential parameters are necessary but not sufficient conditions. They essentially ensure that the potential is flat, the first parameter, ϵ_V , being a measure of the slope of $V(\phi)$ and the second parameter, η_V , being a measure of its curvature. However, the SRA may not be satisfied even with a very flat potential if the scalar field has a large kinetic energy.

At the end of inflation we have

$$\varepsilon(\phi_{\text{end}}) = 1, \quad \epsilon_V(\phi_{\text{end}}) \approx 1. \quad (2.20)$$

One important quantity to characterize the inflationary phase, is the number of e-folds before inflation ends, N , which we have introduced earlier. To clarify its meaning we can write it as follows

$$N(\phi) \equiv \ln \frac{a_{\text{end}}}{a}, \quad (2.21)$$

hence it represents a measure of the factor by which the scale factor has expanded between a given instant with field value ϕ , and the end of inflation with field value ϕ_{end} . In other terms, constraining this quantity, is equivalent to setting constraints on the duration of inflation. We can express the number of e-folds in terms of the slow-roll parameters as

$$N(\phi) = \int_{\phi_{\text{end}}}^{\phi} \frac{d\phi}{\sqrt{2\varepsilon}} \approx \int_{\phi_{\text{end}}}^{\phi} \frac{d\phi}{\sqrt{2\epsilon_V}}, \quad (2.22)$$

which allows us to relate it to inflationary models. The number of e-folds separating the time at which the observed CMB fluctuations exit the horizon during inflation and ϕ_{end} , depends on the specific inflation model and the details of reheating. It is scale dependent and can be expressed via the following equation [313, 327, 30]

$$N_*(k) \simeq 67 - \ln \left(\frac{k}{a_0 H_0} \right) + \frac{1}{4} \ln \left(\frac{V_*^2}{M_{\text{Pl}}^4 \rho_{\text{end}}} \right) + \frac{1 - 3w_{\text{int}}}{12(1 + w_{\text{int}})} \ln \left(\frac{\rho_{th}}{\rho_{\text{end}}} \right) - \frac{1}{12} \ln(g_{th}), \quad (2.23)$$

where the subscript $*$ indicates a quantity at horizon exit, the subscript 0 denotes a quantity at the present time, ρ_{end} corresponds to the energy density at the end of inflation, w_{int} is the effective equation of state parameter between the end of inflation and thermalization, ρ_{th} is the thermalization energy scale, and g_{th} is the corresponding number of effective bosonic degrees of freedom. For a fixed physical scale, we can derive a range of possible values for a given inflationary model as shown later in this chapter in figure 2.8.

2.3 Seeding structure in the Universe

Now that we have seen how a homogeneous scalar field can give rise to negative pressure and accelerated expansion, and we have quantitatively studied the dynamics of the inflationary phase. We will tackle the question of the origin of the inhomogeneities that we observe in the CMB, that constitute the seed of all structures in the Universe. We will first start by briefly exposing some elements of linear cosmological perturbation theory, before explaining how quantum fluctuations can generate the primordial inhomogeneities.

2.3.1 Metric perturbations

SVT decomposition

We can write any quantity $Q(t, \mathbf{x})$ as the sum of a time dependent homogeneous background $\bar{Q}(t)$, and a perturbation $\delta Q(t, \mathbf{x})$ that depends on both space and time. In particular, we can write the metric tensor as

$$g_{\mu\nu}(t, \mathbf{x}) = \bar{g}_{\mu\nu}(t) + \delta g_{\mu\nu}(t, \mathbf{x}). \quad (2.24)$$

The perturbed metric tensor is symmetric and contains 10 degrees of freedom describing different gravitational effects. At linear order, it can be decomposed as follows

$$\begin{aligned} ds^2 &= g_{\mu\nu} dx^\mu dx^\nu \\ &= -(1 + 2\Psi) dt^2 + 2aB_i dx^i dt + a^2 [(1 - 2\Phi)\delta_{ij} + E_{ij}] dx^i dx^j. \end{aligned} \quad (2.25)$$

The variables can be decomposed into scalars, vectors and tensors, respectively defined according to how they transform under spatial rotations. More precisely, for a given wavevector \mathbf{k} , when rotating the coordinate system around the wavevector by an angle ψ , a perturbation that transforms as the following

$$X_{\mathbf{k}} \rightarrow e^{im\psi} X_{\mathbf{k}} \quad (2.26)$$

is said to have helicity m . A helicity of 0 correspond to scalar perturbations, ± 1 helicities correspond to vectors, and ± 2 helicities to tensor perturbations. The scalar-vector-tensor (SVT) decomposition first introduced by Bardeen [75] allows to simplify the study of cosmological perturbations, as the different types of perturbations are decoupled at linear order and hence can be treated separately. In particular, B_i can be written as the sum of the gradient of a scalar B , and a divergence-free vector S_i

$$B_i \equiv \partial_i B + S_i, \quad \text{with } \partial^i S_i = 0. \quad (2.27)$$

Similarly E_{ij} can be decomposed into a scalar E , a divergence-free vector F_i , and a divergenceless and traceless symmetric tensor h_{ij} as¹

$$E_{ij} \equiv 2\partial_{(i} E + \partial_{j)} F_i + h_{ij}, \quad \text{with } \partial^i F_i = 0, h_i^i = \partial^i h_{ij} = 0. \quad (2.28)$$

The 10 degrees of freedom of the perturbed metric tensor have been decomposed into:

- 4 *scalars*: Φ , B , Ψ , E corresponding to 4 degrees of freedom
- 2 *vectors*: S_i and F_i corresponding to 4 degrees freedom, given that their divergenceless property constrains 2 out of the initial 6 degrees of freedom.
- 1 *tensor*: h_{ij} corresponding to 2 degrees of freedom, given that the tracelessness and the divergence-free property constrains 1+3 out of the initial 6 degrees of freedom of the symmetric tensor.

In the following we will ignore vector perturbations given that they decay with expansion, and are not generated by inflation [83].

¹ The notation $a_{(i} b_{j)} \equiv a_i b_j + a_j b_i$.

Gauge choice

One important point that we have ignored in the previous development is that when introducing a perturbation δQ , we compare the quantity Q in perturbed spacetime and the same quantity \bar{Q} in the background spacetime. However, this comparison can only be meaningful if we introduce a one-to-one mapping between the two spacetimes. Figure 2.3 shows an illustration of this correspondence. Such mapping is called a *gauge choice*. It is not unique, and no natural choice exists. This gauge freedom can lead to the introduction of fictitious perturbations arising not from physical perturbations but only as an artefact of the choice of coordinates. In order to address this issue, two possible approaches can be adopted. As a first option, one can derive gauge-invariant quantities by combining metric and matter perturbations, these quantities do not represent coordinate effects. The second option is to fix the gauge by introducing some condition to make the coordinate choice unique, after solving the equations, the observable quantities should be independent of the gauge.

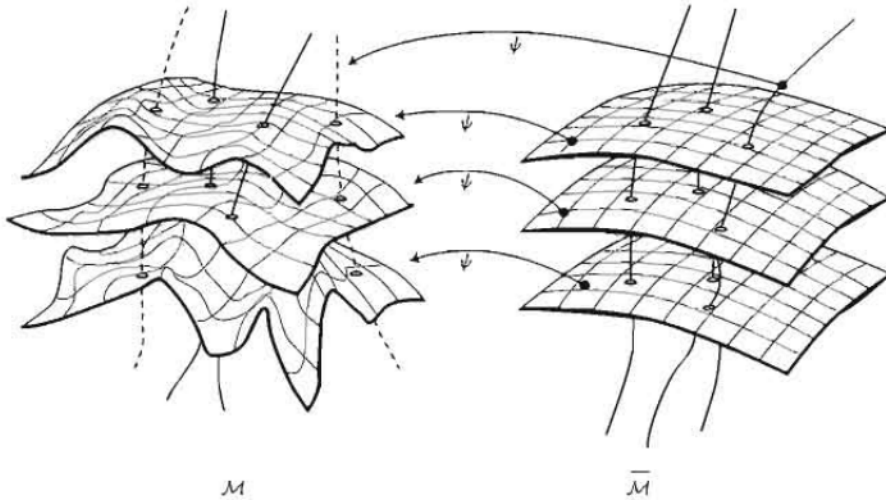


Figure 2.3 – Illustration of a one-to-one mapping, ψ , between the FLRW background spacetime, $\bar{\mathcal{M}}$, and the perturbed spacetime \mathcal{M} . *Credit: Patrick Peter & Jean-Philippe Uzan [369].*

2.3.2 Matter perturbations

The same reasoning can be applied on the stress-energy tensor to decompose it into scalars, vectors and tensors. We can write

$$\begin{aligned} T_0^0 &= -(\bar{\rho} + \delta\rho) \\ T_i^0 &= (\bar{\rho} + \bar{p})v_i \\ T_j^i &= \delta_j^i(\bar{p} + \delta p) + \Pi_j^i. \end{aligned} \tag{2.29}$$

The bulk velocity v_i can be decomposed into a scalar, v , (also called velocity potential) and a divergence-free vector, v_i^V such that $v_i = \partial_i v + v_i^V$, with $\partial^i v_i^V = 0$. Π_j^i is a symmetric traceless 3-tensor describing *anisotropic stress*. Similarly we can

decompose it into scalar, vector and tensor perturbations: $\Pi_{ij} = \Pi_{ij}^S + \Pi_{ij}^V + \Pi_{ij}^T$, with

$$\begin{aligned}\Pi_{ij}^S &= (\partial_i \partial_j - \frac{1}{3} \delta_{ij} \nabla^2) \Pi, \\ \Pi_{ij}^V &= \partial_{(i} \Pi_{j)}, \\ \partial^i \Pi_{ij}^T &= 0.\end{aligned}\tag{2.30}$$

As mentioned earlier, in order to avoid fictitious gauge modes in the study of perturbations, in particular scalar perturbations ($\delta\rho$, δp , v , Π)—given that we ignore vector perturbations and that tensor perturbations are gauge-invariant (at linear order)—one approach is to introduce gauge independent quantities which are constructed out of non trivial combinations of metric and matter perturbations. One such quantity, which is central in the study of the initial density fluctuations induced by inflation is the so-called *comoving curvature perturbation*

$$\mathcal{R} \equiv \Phi - \frac{H}{\bar{\rho} + \bar{p}} \delta q,\tag{2.31}$$

where δq is a scalar defined by $T_i^0 = (\bar{\rho} + \bar{p}) \partial_i v \equiv \partial_i \delta q$. During slow-roll inflation, we have $T_i^0 = -\dot{\bar{\phi}} \partial_i \delta \phi$, the comoving curvature perturbation can hence be written as

$$\mathcal{R} = \Phi + \frac{H}{\dot{\bar{\phi}}} \delta \phi.\tag{2.32}$$

This latter form clarifies the geometrical interpretation of \mathcal{R} as a measure of the spatial curvature of constant- ϕ hypersurfaces.

Statistical characterization

Given the random nature of perturbations, no theory can give a deterministic prediction of the exact inhomogeneities as a function of time and space. Instead, they are treated as stochastic quantities and characterized statistically. As long as a given perturbation is following a linear evolution, the shape of its probability distribution is simply scaled by a factor. In particular, if primordial perturbations follow a Gaussian distribution, they will remain so in the linear regime. In this case, we can characterize the fluctuations of any Gaussian quantity, e.g. $\mathcal{R}(\mathbf{x})$, with the two-point correlation function, which under the assumption statistical homogeneity and isotropy can be written as

$$\xi(\mathbf{x}, \mathbf{x}') \equiv \langle \mathcal{R}(\mathbf{x}) \mathcal{R}(\mathbf{x}') \rangle = \xi(|\mathbf{x} - \mathbf{x}'|)\tag{2.33}$$

Where $\langle \dots \rangle$ denotes the ensemble average. We can represent \mathcal{R} in Fourier space and back in real space via the following relations

$$\begin{aligned}\mathcal{R}(\mathbf{k}) &\equiv \int d^3\mathbf{x} \mathcal{R}(\mathbf{x}) e^{-i\mathbf{k}\cdot\mathbf{x}} \\ \mathcal{R}(\mathbf{x}) &= \frac{1}{(2\pi)^3} \int d^3\mathbf{k} \mathcal{R}(\mathbf{k}) e^{i\mathbf{k}\cdot\mathbf{x}}.\end{aligned}\tag{2.34}$$

In Fourier space, the primordial scalar fluctuations are statistically characterized by the power spectrum of \mathcal{R} which can be written as

$$\langle \mathcal{R}(\mathbf{k}) \mathcal{R}^*(\mathbf{k}') \rangle \equiv (2\pi)^3 \delta(\mathbf{k} - \mathbf{k}') \mathcal{P}_{\mathcal{R}}(k), \quad \Delta_s^2 \equiv \Delta_{\mathcal{R}}^2 = \frac{k^3}{2\pi^2} \mathcal{P}_{\mathcal{R}}(k). \quad (2.35)$$

where we have defined $\mathcal{P}_{\mathcal{R}}(k)$ as the Fourier transform of the two-point correlation function.

The scalar spectral index determines the scale dependence of the power spectrum, and is defined as

$$n_s - 1 \equiv \frac{d \ln \Delta_s^2}{d \ln k}, \quad (2.36)$$

A scale-invariant power spectrum corresponds to $n_s = 1$. The running of the spectral index is defined by the following

$$\alpha_s \equiv \frac{dn_s}{d \ln k}. \quad (2.37)$$

We can then write the power spectrum of scalar fluctuations as a power law of the following form

$$\Delta_s^2(k) = A_s(k_*) \left(\frac{k}{k_*} \right)^{n_s(k_*) - 1 + \frac{1}{2} \alpha_s(k_*) \ln k/k_*}, \quad (2.38)$$

where k_* is an arbitrary pivot scale.

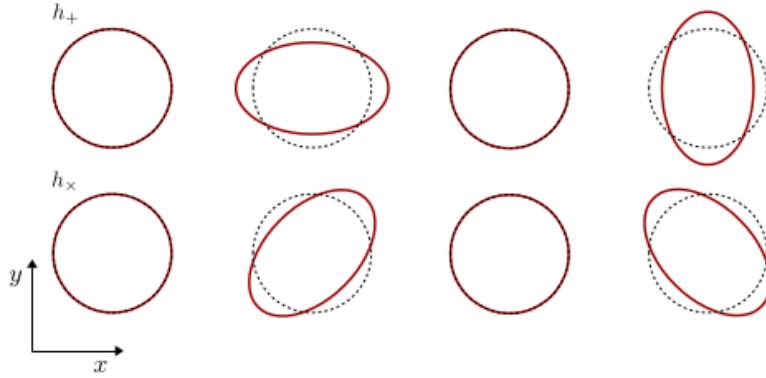


Figure 2.4 – Illustration of the spacetime distortion induced by the two polarization states $+$ (upper panel) and \times (lower panel) of a propagating gravitational wave in the plane perpendicular to its direction. *Credit: Scott Dodelson & Fabian Schmidt [165].*

Tensor modes can be decomposed in two distinct polarization states, $+$ and \times corresponding to the direction of compression or stretching of spacetime perpendicular to the direction of propagation of the gravitational wave (Cf. figure 2.4 for an illustration). The power spectrum of the two polarization modes of h_{ij} , h^+ and h^\times , is given by

$$\langle h^\iota(\mathbf{k}) h^{\iota*}(\mathbf{k}') \rangle = (2\pi)^3 \delta(\mathbf{k} - \mathbf{k}') \mathcal{P}_{h^\iota}(k), \quad \Delta_{h^\iota}^2 = \frac{k^3}{2\pi^2} \mathcal{P}_{h^\iota}(k). \quad (2.39)$$

with $\iota \in \{+, \times\}$. The power spectrum of tensor perturbations is defined as the sum of the power spectra of the two polarization modes

$$\Delta_t^2 \equiv \Delta_{h+}^2 + \Delta_{h\times}^2 \quad (2.40)$$

Similarly, the tilt of the power spectrum is given by

$$n_t \equiv \frac{d \ln \Delta_t^2}{d \ln k}, \quad (2.41)$$

and the power spectrum of tensor fluctuations takes the following form

$$\Delta_t^2(k) = A_t(k_*) \left(\frac{k}{k_*} \right)^{n_t(k_*)}. \quad (2.42)$$

The power spectrum is a sufficient statistic if the random field studied, e.g. \mathcal{R} , is Gaussian. In case of the presence of primordial non-Gaussianities, the power spectra do not encode all the information and higher-order correlation functions are needed (e.g. bispectrum, trispectrum). We will explore this point with greater details in the next chapter.

2.3.3 Quantum origin of structures

To give an account of the density inhomogeneities, one needs to perturb the classical picture we have introduced in the previous section, of a scalar field, $\bar{\phi}(t)$, rolling down a potential. In the presence of quantum fluctuations $\delta\phi(t, \mathbf{x})$, inflation ends at slightly different times in different regions of the Universe. This means that each of these regions experiences a slightly different evolution, ultimately leading to relative fluctuations in density $\delta\rho(t, \mathbf{x})$.

We can write down the action for single-field slow roll models of inflation, and expand it to second order, separately, for both scalar and tensor fluctuations. This allows us to derive an equation of motion for the perturbations, which takes the form of a simple harmonic oscillator. It is called the *Mukhanov equation* and can be written as

$$v_k'' + \left(k^2 - \frac{z''}{z} \right) v_k = 0, \quad (2.43)$$

with

$$v_k = \begin{cases} (a/2)h_k M_{Pl} & \text{tensor} \\ z\mathcal{R}_k M_{Pl} & \text{scalar} \end{cases} \quad (2.44)$$

and

$$z = \begin{cases} a & \text{tensor} \\ a\sqrt{2\varepsilon} & \text{scalar} \end{cases} \quad (2.45)$$

where $(\dots)_k$ denotes the Fourier transform of a quantity depending only on the magnitude, k , of the wavevector, \mathbf{k} . The dependence of z on the background evolution makes the equation difficult to solve. Consequently, general solutions of the Mukhanov equation can only be computed numerically. However, approximate analytical solutions can be obtained in the de Sitter limit under the slow-roll approximation conditions, and can be very useful to build an intuitive understanding.

In this case $\varepsilon \rightarrow 0$ and $z''/z \sim a''/a \sim 2/\tau^2$. The functions $\frac{e^{-ik\tau}}{\sqrt{2k}} \left(1 - \frac{i}{k\tau}\right)$ and $\frac{e^{ik\tau}}{\sqrt{2k}} \left(1 + \frac{i}{k\tau}\right)$ constitute a basis of the solution space of Eq. (2.43). The canonical quantization of v_k , along with the choice of the vacuum state as the Minkowski vacuum of a comoving observer in the far past, allow us to derive two boundary conditions

$$\begin{aligned} \langle v_k, v_k \rangle &\equiv \frac{i}{\hbar} (v_k^* v'_k - v_k'^* v_k) = 1 \\ v_k &\underset{\tau \rightarrow -\infty}{\sim} \frac{e^{-ik\tau}}{\sqrt{2k}} \end{aligned} \quad (2.46)$$

These two conditions completely fix the mode functions on all scales, which are then expressed as

$$v_k = \frac{e^{-ik\tau}}{\sqrt{2k}} \left(1 - \frac{i}{k\tau}\right) \quad (2.47)$$

Using this expression we can derive the primordial power spectrum of scalar and tensor perturbations. First of all, we have

$$\begin{aligned} \frac{1}{a^2} \langle \hat{v}_{\mathbf{k}}(\tau) \hat{v}_{\mathbf{k}'}^*(\tau) \rangle &= (2\pi)^3 \delta(\mathbf{k} - \mathbf{k}') \frac{H^2}{2k^3} (1 + k^2 \tau^2) \\ &\xrightarrow{|k\tau| \ll 1} (2\pi)^3 \delta(\mathbf{k} - \mathbf{k}') \frac{H^2}{2k^3}. \end{aligned} \quad (2.48)$$

The latter result shows that the power spectrum freezes at superhorizon scales. We can now give the scalar and tensor perturbations power spectrum at horizon crossing. Quantities at this instant are indicated by the subscript \star

$$\begin{aligned} \Delta_s^2(k) &\equiv \Delta_{\mathcal{R}}^2(k) = \frac{1}{8\pi^2} \frac{H_\star^2}{M_{Pl}} \frac{1}{\varepsilon_\star}, \\ \Delta_t^2(k) &\equiv 2\Delta_h^2(k) = \frac{2}{\pi^2} \frac{H_\star^2}{M_{Pl}^2}. \end{aligned} \quad (2.49)$$

The *tensor-to-scalar* ratio, r , is defined as

$$r \equiv \frac{\Delta_t^2}{\Delta_s^2} = 16\varepsilon_\star. \quad (2.50)$$

It is a direct measure of the energy scale of inflation

$$V^{1/4} \sim \left(\frac{r}{0.01}\right)^{1/4} 10^{16} \text{GeV}. \quad (2.51)$$

The power spectrum remain frozen until the k modes re-enter the comoving Hubble horizon. As such, it sets the initial conditions for the perturbations that we see in the CMB, and that later grow to form all the structures in the Universe. Notice that H and ε vary slightly during inflation, which means that different scales cross the horizon at different times. As a result, the power spectrum is not exactly but only nearly scale invariant. The spectral indices which characterize the scale dependence of the spectra, satisfy the following relations

$$\begin{aligned} n_s - 1 &= 2\eta_\star - 4\varepsilon_\star, \\ n_t &= -2\varepsilon_\star. \end{aligned} \quad (2.52)$$

In addition the tensor-to-scalar ratio, r , and the tensor tilt, n_t should satisfy the following consistency condition in the case of single field slow-roll models

$$r = -8n_t. \quad (2.53)$$

In the slow-roll approximation, these quantities can be written as a function of the potential slow-roll parameters, ε_V and η_V . In this case, measurement of these quantities contain information about the shape of the inflaton potential $V(\phi)$. H is a measure of the scale of the potential, ε_V is a measure of its first derivative $V'(\phi)$, and η_V is a measure of its second derivative $V''(\phi)$.

2.4 Observational evidence and constraints

Flatness

The total density of the Universe, Ω_{tot} ,

$$\Omega_{tot} = \Omega_b + \Omega_{cdm} + \Omega_\gamma + \Omega_\nu + \Omega_\Lambda, \quad (2.54)$$

fixes the spatial geometry of the Universe. A flat Universe would correspond to $\Omega_{tot} = 1$. As we have seen before, the inflationary mechanism flattens the Universe. Hence, setting constraints on Ω_{tot} is one way of testing the inflationary paradigm. Our measurements of spatial curvature are fundamentally limited by cosmic variance to a standard deviation of order $\sim 10^{-5}$ [464]. Latest results from Planck combined with BAO data [377] are consistent with spatial flatness with a precision of 0.4%. That is still 2 orders of magnitude above the cosmic variance floor, nonetheless it is a remarkable agreement with the theory. One might argue, however, that this is not compelling evidence for inflation given that it is designed to produce this level of spatial flatness, and that some inflationary models also allow for negative curvature [215, 111, 475, 389, 320].

Coherent phases and Superhorizon Fluctuations

Inflation generates curvature perturbations with a nearly scale invariant spectrum, $\Delta_s^2 \propto k^{n_s-1}$, with $n_s \approx 1$. Any given Fourier mode oscillates during inflation as long as it is inside the horizon. Once it exits the horizon, its amplitude remains constant. Upon horizon re-entry, the curvature perturbation sources density fluctuations which then evolve under the influence of gravity and pressure, this causes oscillations in the density field which are in turn coupled to fluctuations in the radiation. In the limit where recombination happens instantaneously, which is a good approximation for sufficiently large scales, each wavelength would be captured at a different phase of its oscillation. These effects will be studied in greater details in the next chapter. However, one important point to be made here is that the peak structure we observe in the power spectra figure 2.5 indicates that for any given scale, the corresponding Fourier modes all oscillate in phase thereby producing the peaks and troughs. If the initial phases of the different modes were arbitrary, as is portrayed in figure 2.6, the structure would be washed out and all scales would have the same amplitude. Inflation

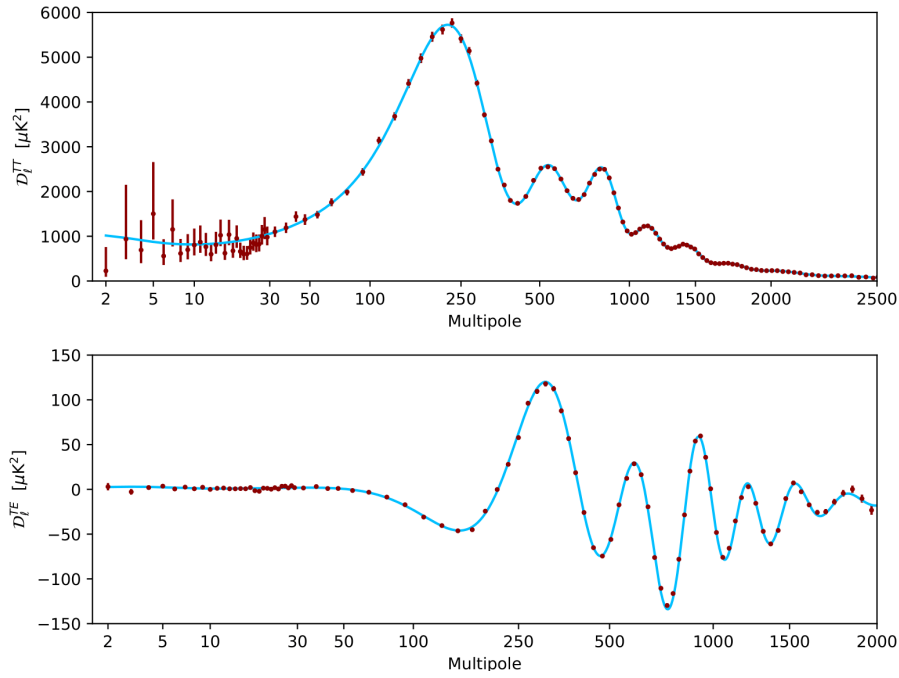


Figure 2.5 – The Planck 2018 angular power spectra of the CMB (TT and TE). The solid blue line is the best-fit Λ CDM model to temperature and polarization data. *Credit:* Planck collaboration.

provides an explanation for the coherence of the initial phases, due to the fluctuations freezing when they exit the horizon during the inflationary period, and the phases being re-synchronized when the modes, corresponding to any given scale, start oscillating at the same time, again at horizon re-entry. The temporal coherence of the Fourier modes then builds up the structure we see in the CMB power spectrum despite having arbitrary initial amplitudes as shown in figure 2.7.

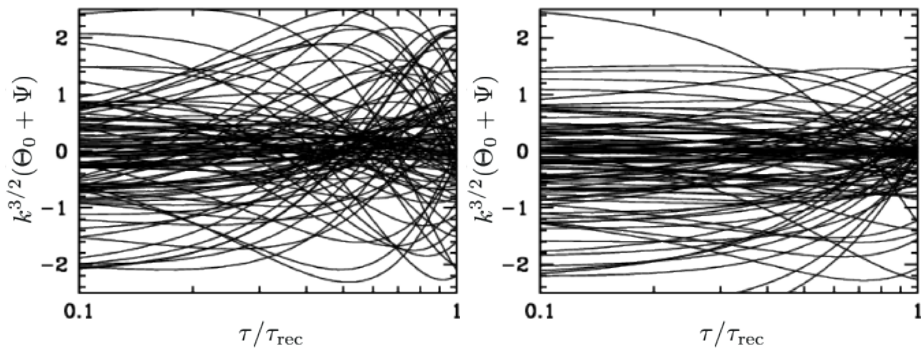


Figure 2.6 – Evolution of monopole and dipole Fourier modes with arbitrary initial phases and different amplitudes, until the time of recombination τ_{rec} . *Credit:* Adapted version from [166] by Daniel Baumann [83].

The peaks and troughs structure seen in the TT angular power spectrum at the top of figure 2.5 resides at $\ell > 200$, which corresponds to angular scales that were within the horizon at recombination. Therefore, one may provide an alternative explanation by conceiving a theory of structure formation which produces these coherent phases through causal mechanisms. However, this cannot be done for

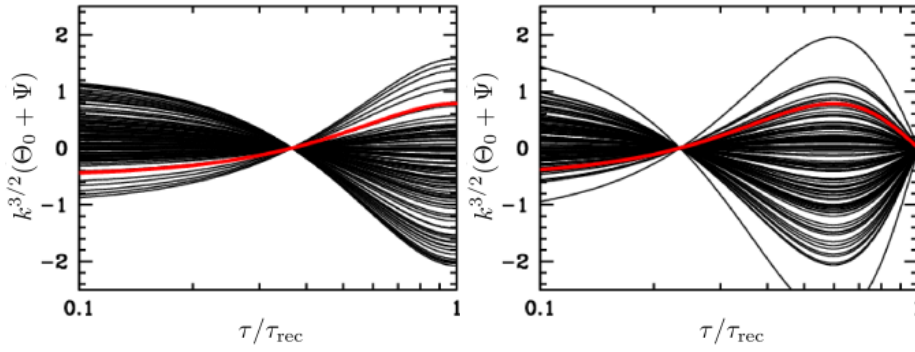


Figure 2.7 – Evolution of monopole and dipole Fourier modes with coherent initial phases and different amplitudes, until the time of recombination τ_{rec} .

Credit: Adapted version from [166] by Daniel Baumann [83].

the TE angular power spectrum which shows a negative peak at $100 < \ell < 200$ which corresponds to scales not within the horizon at recombination. This particular feature is a compelling evidence for an inflation-like mechanism, since no causal physics can be invoked to explain it unless it produces a shrinking comoving horizon.

Scale invariance, Gaussianity, Adiabaticity

The simplest inflationary models predict nearly scale invariant, Gaussian and adiabatic scalar perturbations.

In the latest results from Planck combined with BAO measurements [377], the base Λ CDM cosmology gives a scalar spectral index,

$$n_s = 0.9665 \pm 0.0038 \quad 68\% \text{ CL}, \quad (2.55)$$

hence yielding a nearly scale invariant, red spectrum ($n_s < 1$). This is consistent with theoretical expectations, since the Hubble parameter H evolved in time, making the spectrum of fluctuations scale dependent given that Fourier modes exit the horizon at different times. These results confirm that a scale independent spectrum ($n_s = 1$) is definitely ruled out with a significance of up to 8.4σ (when lensing constraints are included).

Constraints on the scalar spectral index are often given jointly with the tensor-to-scalar ratio in the so-called $n_s - r$ plane, as shown for example in figure 2.8. As the constraints tighten, this allows us to rule out a number of inflationary models or exclude a sizeable region of their parameter space. The tightest constraint to date on the tensor-to-scalar ratio is $r < 0.044$ at 95% CL obtained via the combination of Planck and BICEP2/Keck array data [454].

Another way of constraining inflation is through the study of primordial non-Gaussianities. In the last Planck analysis [375], a number of tests are conducted. Results are consistent with single-field slow roll inflation, and strong constraints on alternative models are derived. However, sensitivities are still not high enough to be able discriminate between different scenarios at this stage.

In single-field inflation, perturbations of the cosmological fluid arise from the same curvature perturbation \mathcal{R} , and hence are characterized by a single degree of freedom, associated with a local shift along the trajectory of the inflaton. This

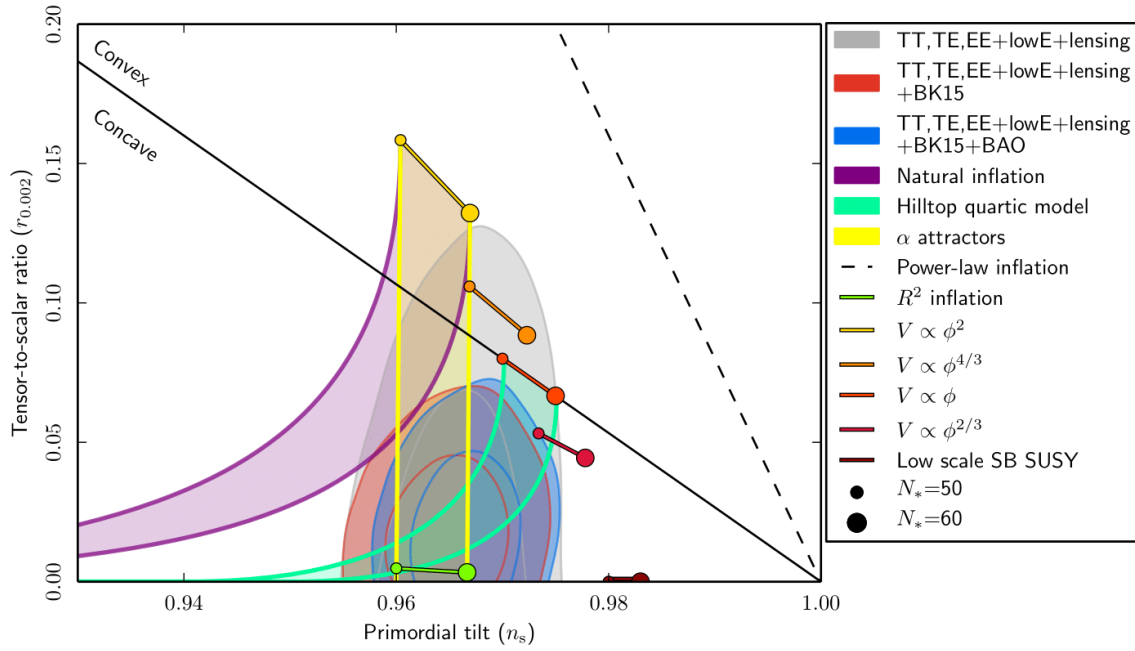


Figure 2.8 – Constraints from *Planck* 2018 data, and their combination with BICEP2/Keck array (BK15) and BAO data, on n_s and r compared to predictions of selected inflationary models (Cf. [377]). The contours represent the 68% and 95% CL regions. *Credit:* Planck collaboration [377]

means that there cannot be spatial variations in the relative densities of different species, perturbations are said to be *adiabatic*. The adiabaticity condition can be written as

$$\delta_m \equiv \frac{\delta \rho_m}{\bar{\rho}_m} = \frac{3}{4} \delta_r \equiv \frac{3}{4} \frac{\delta \rho_r}{\bar{\rho}_r}, \quad (2.56)$$

where δ_X defined above is called *the energy density contrast*, and the m and r subscripts indicate matter (baryons or CDM) and radiation (photons or neutrinos) species respectively. A deviation from adiabaticity would be a signature of physics beyond single-field inflation. In the presence of more than one degree of freedom generating the primordial perturbations, we can have, in addition to the adiabatic mode, multiple non-decaying *isocurvature modes*. These are characterized by an asymptotically vanishing total curvature perturbation on super-horizon scales, and two species with opposite density perturbations compensating each other. These latter perturbations are quantified by the so-called, *entropy perturbation*, where the photon density is taken as a reference. For a given species with equation of state parameter w_i , it reads

$$S_i^d \equiv \frac{1}{1 + w_i} \delta_{i \neq \gamma} - \frac{3}{4} \delta_\gamma. \quad (2.57)$$

There are also isocurvature perturbations of the velocity type, with no initial density perturbations ($\delta_i = 0$), but different initial velocities. These are characterized by

$$S_i^v = \frac{1}{1 - f_\nu} (v_i^0 - v_\gamma^0), \quad (2.58)$$

where f_ν is the ratio of neutrino density in the total radiation density, and v_i^0 the

initial velocity of species i . The only non-zero velocity isocurvature mode is that of neutrinos.

CMB observations can put stringent limits on the amplitudes of isocurvature modes. The latest results from Planck [377], set 95% CL upper bounds on the non-adiabatic contribution, in correlated mixed adiabatic and isocurvature models, to the CMB temperature variance of 1.3% for CDM isocurvature, and 1.7% for neutrino density, and neutrino velocity isocurvature.

As we have seen until now, the CMB provides a wealth of information not only on the content and structure of the Universe, and the different stages of its evolution, but it is also a crucial probe for understanding the physics that sets the initial conditions of the Hot Big Bang, and its observables can be directly linked to the physics of the Early Universe. In the following chapter, we will focus on reviewing the physics of the CMB in more details, what we can learn from it and what is the status of our knowledge to date.

CHAPTER 3

The Cosmic Microwave Background

Contents

3.1 CMB anisotropies on the sphere	52
3.1.1 Temperature anisotropies	52
3.1.2 Polarization anisotropies	53
3.2 Evolution of perturbations	56
3.2.1 Boltzmann equation	57
3.2.2 Line of sight integral	58
3.2.3 Acoustic oscillations	59
3.3 Secondary anisotropies	62
3.3.1 Scattering effects	62
3.3.2 Gravitational effects	64
3.4 Cosmological parameters and observational constraints	69
3.5 Beyond anisotropies	74
3.5.1 Primordial non-Gaussianity	74
3.5.2 Spectral distortions	76
3.5.3 Sunyaev-Zel'dovich Clusters	79
3.5.4 Cosmic birefringence	81

Soon after the discovery of the CMB by Penzias and Wilson in 1965 [365], a consensus on its interpretation as a relic radiation from the Big Bang has formed, and the controversy with the competing steady state model [102, 249, 349] was definitively settled. Further work in the late 1960s and early 1970s by Sachs, Wolfe, Silk, Peebles, Yu, Harrison and Zel'dovich [402, 416, 417, 364, 234, 479] have lead to the prediction that the CMB should feature anisotropies. These results motivated decades of experimental efforts searching for the CMB fluctuations, crowned by their discovery in 1992 by the COBE Differential Microwave

Radiometers (DMR) [423]. These anisotropies constitute a goldmine of information for cosmologists. They have allowed us to make giant leaps in our understanding of the physics governing the Universe, and after decades of work they are still far from revealing all their secrets. In this chapter, we review aspects of the physics that govern the primordial plasma, how it affects the evolution of the primordial perturbations, and their relation to the anisotropies imprinted in the CMB. We also consider other processes that involve the CMB photons leaving observable signatures. Along with the theoretical background we also discuss the state of the art of the measurements and constraints on different cosmological parameters.

3.1 CMB anisotropies on the sphere

In this section we introduce all the mathematical framework and relevant quantities that are used to statistically characterize the observable anisotropies in the CMB radiation, insofar as they are Gaussian.

3.1.1 Temperature anisotropies

In perturbed spacetime, the phase space distribution function of photons depends on (conformal) time, η , position, \mathbf{x} , the photon energy, E , or alternatively we can use the *comoving energy*, $\epsilon \equiv Ea$, and the direction of propagation, \mathbf{n} . We recall that, as long as the background temperature of photons evolves as, $\bar{T}(\eta) \propto a^{-1}$, the background description of photons distribution is given by the Bose-Einstein distribution

$$\bar{f}(\epsilon) = \left[\exp \left(\frac{\epsilon}{a\bar{T}(\eta)} \right) - 1 \right]^{-1} = \left[\exp \left(\frac{\epsilon}{T_0} \right) - 1 \right]^{-1}, \quad (3.1)$$

where T_0 is the CMB temperature at the present day. To study first-order perturbations, we can expand the distribution function into a background and perturbations part, $f = \bar{f}(1 + \delta f)$. We introduce the fractional temperature perturbation, $\Theta(\eta, \mathbf{x}, \mathbf{n}) \equiv \delta T(\eta, \mathbf{x}, \mathbf{n})/\bar{T}(\eta)$. One can show through the Boltzmann equations, which we will cover later in this chapter, that Θ is independent of ϵ , given that the gravitational source terms and the scattering terms are both proportional to $\frac{d \ln \bar{f}}{d \ln \epsilon}$. The distribution function takes the following form

$$f(\eta, \mathbf{x}, \epsilon, \mathbf{n}) = \left[\exp \left(\frac{\epsilon}{T_0(1 + \Theta(\eta, \mathbf{x}, \mathbf{n}))} \right) - 1 \right]^{-1}. \quad (3.2)$$

Assuming that Θ is small ($\ll 1$), we can then write

$$\Theta(\eta, \mathbf{x}, \mathbf{n}) = - \left(\frac{d \ln \bar{f}}{d \ln \epsilon} \right)^{-1} \delta f. \quad (3.3)$$

We assume statistical isotropy and homogeneity of the Universe in the following. To statistically characterize the anisotropies observed with a typical CMB experiment, we first relate Θ to its Fourier transform further expanded in Legendre

polynomials, P_ℓ , as

$$\Theta(\eta, \mathbf{x}, \mathbf{n}) = \int d^3\mathbf{k} e^{-i\mathbf{k}\cdot\mathbf{x}} \sum_{\ell}^{\infty} (-i)^\ell (2\ell + 1) \Theta_\ell(\eta, \mathbf{k}) P_\ell(\mathbf{k} \cdot \mathbf{n}). \quad (3.4)$$

A CMB observer measures the temperature anisotropy of photons coming from a given direction \mathbf{n}_0 , at his position \mathbf{x}_0 , and at the present time η_0 . It can be expressed in a basis of spherical harmonics, $Y_{\ell m}$,

$$\Theta(\eta_0, \mathbf{x}_0, \mathbf{n}_0) = \sum_{\ell m} a_{\ell m}(\eta_0, \mathbf{x}_0) Y_{\ell m}(\mathbf{n}_0). \quad (3.5)$$

Using basic relations between spherical harmonics and Legendre polynomials, we can then suitably express the $a_{\ell m}$ coefficients as a function of Θ_ℓ

$$a_{\ell m}(\eta_0, \mathbf{x}_0) = 4\pi (-i)^\ell \int d^3\mathbf{k} Y_{\ell m}^*(\mathbf{k}) \Theta_\ell(\eta_0, \mathbf{k}). \quad (3.6)$$

Given the orthogonality of the spherical harmonics, the covariance of the $a_{\ell m}$'s is diagonal and given by

$$\langle a_{\ell m} a_{\ell' m'}^* \rangle = C_\ell^{TT} \delta_{\ell\ell'} \delta_{mm'}. \quad (3.7)$$

The covariance in harmonic space is related to the one in real space between two directions \mathbf{n} and \mathbf{n}' via the following

$$C(\mathbf{n}, \mathbf{n}') = \sum_{\ell}^{\infty} \frac{2\ell + 1}{4\pi} C_\ell^{TT} P_\ell(\mathbf{n} \cdot \mathbf{n}'). \quad (3.8)$$

The CMB temperature fluctuation multipoles Θ_ℓ are dominated by the scalar modes sourced by the comoving curvature perturbation, $\mathcal{R}(\mathbf{k})$. In the linear regime, we can relate the two via a transfer function, $\Delta_{T\ell}(\eta, k)$, which encodes the evolution of the initial perturbations, inferred from the Boltzmann equations.

$$\Theta_\ell(\eta, \mathbf{k}) = \mathcal{R}(\mathbf{k}) \Delta_{T\ell}(\eta, k). \quad (3.9)$$

We can finally express the angular power spectrum of CMB temperature anisotropies, C_ℓ^{TT} , as a function of the primordial power spectrum, $\mathcal{P}_\mathcal{R}(k)$

$$C_\ell^{TT} = (4\pi)^2 \int k^2 dk \mathcal{P}_\mathcal{R}(k) \Delta_{T\ell}^2(\eta_0, k). \quad (3.10)$$

3.1.2 Polarization anisotropies

Stokes parameters

The incoming CMB radiation is polarized, and can be fully described in terms of the *Stokes parameters*. Let us define these parameters by first considering a monochromatic electromagnetic plane wave with frequency ν , propagating along the z -direction. At a given point in space, the two components of the electric vector in the $x - y$ plane can be written as

$$\begin{aligned} E_x(t) &= A_x(t) \cos(\nu t - \theta_x(t)), \\ E_y(t) &= A_y(t) \cos(\nu t - \theta_y(t)). \end{aligned} \quad (3.11)$$

In case of the presence of correlations between the two components, the electromagnetic wave is polarized. The four Stokes parameters are defined in terms of the time averages over the amplitudes and phases of the electromagnetic wave

$$\begin{aligned} I &\equiv \langle A_x^2 \rangle + \langle A_y^2 \rangle, \\ Q &\equiv \langle A_x^2 \rangle - \langle A_y^2 \rangle, \\ U &\equiv \langle 2A_x A_y \cos(\theta_x - \theta_y) \rangle, \\ V &\equiv \langle 2A_x A_y \sin(\theta_x - \theta_y) \rangle. \end{aligned} \quad (3.12)$$

I is the total intensity of the radiation, while the remaining three parameters characterize its polarization state. Q and U describe linear polarization, while circular polarization is described by V . The parameters Q and U are coordinate dependent; under rotation of the $x - y$ plane by an angle ψ they transform as spin-2 quantities, with two equivalent representations

$$\begin{pmatrix} Q' \\ U' \end{pmatrix} = \begin{pmatrix} \cos 2\psi & \sin 2\psi \\ -\sin 2\psi & \cos 2\psi \end{pmatrix} \begin{pmatrix} Q \\ U \end{pmatrix} \quad (3.13)$$

$$(Q \pm iU)' = e^{\pm 2i\psi} (Q \pm iU), \quad (3.14)$$

Linear polarization is therefore represented by a spin-2 field, i.e. headless-vectors with magnitude P and orientation α defined as

$$P \equiv \sqrt{Q^2 + U^2} \quad \alpha \equiv \frac{1}{2} \arctan \left(\frac{U}{Q} \right). \quad (3.15)$$

Spin-weighted spherical harmonics

In general, the spherical harmonic transform of a spin- s field $f(\mathbf{n})$ is expressed in a basis of the so-called *spin-weighted spherical harmonic functions*, ${}_{\pm s}Y_{\ell m}$

$$\begin{aligned} {}_{\pm s}a_{\ell m} &= \int d\Omega f(\mathbf{n}) {}_{\pm s}Y_{\ell m}^*, \\ f(\mathbf{n}) &= \sum_{\ell m} {}_{\pm s}a_{\ell m} {}_{\pm s}Y_{\ell m}, \end{aligned} \quad (3.16)$$

with, for $s \geq 0$, we have

$$\begin{aligned} {}_sY_{\ell m} &= \sqrt{\frac{(\ell - s)!}{(\ell + s)!}} \bar{\partial}^s Y_{\ell m}, \\ {}_{-s}Y_{\ell m} &= (-1)^s \sqrt{\frac{(\ell - s)!}{(\ell + s)!}} \partial^s Y_{\ell m}. \end{aligned} \quad (3.17)$$

The operators ∂ and $\bar{\partial}$ are covariant derivatives on the sphere. They are, respectively, spin-raising and spin-lowering operators, i.e acting on a spin- s field $f(\mathbf{n})$, they transform under coordinate rotation by an angle ψ as the following

$$\begin{aligned} (\partial f)' &= e^{-i(s+1)\psi} \partial f, \\ (\bar{\partial} f)' &= e^{-i(s-1)\psi} \bar{\partial} f. \end{aligned} \quad (3.18)$$

Their explicit expression in spherical coordinates is given by [478]

$$\begin{aligned}\bar{\partial}f(\theta, \phi) &= -\sin^s(\theta) \left[\frac{\partial}{\partial\theta} + i \csc(\theta) \frac{\partial}{\partial\phi} \right] \sin^{-s}(\theta), \\ \bar{\partial}f(\theta, \phi) &= -\sin^{-s}(\theta) \left[\frac{\partial}{\partial\theta} - i \csc(\theta) \frac{\partial}{\partial\phi} \right] \sin^s(\theta).\end{aligned}\tag{3.19}$$

E-modes and B-modes decomposition

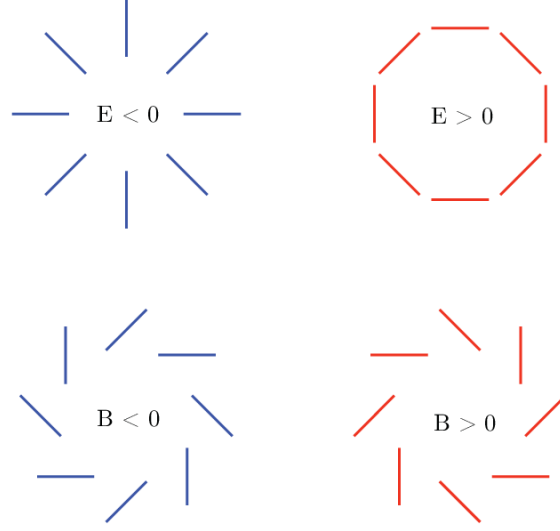


Figure 3.1 – Illustration of E-modes and B-modes polarization patterns.

A very useful decomposition of a spin- s field in harmonic space, is the so-called E-modes and B-modes decomposition. For this purpose we introduce the two following operators

$$\begin{aligned}\mathbf{Y}_{s,\ell m}^E &\equiv \mathbf{D}_s^E Y_{\ell m} = \frac{1}{2} \sqrt{\frac{(\ell-s)!}{(\ell+s)!}} \begin{pmatrix} \bar{\partial}^s + (-1)^s \bar{\partial}^s \\ -i(\bar{\partial}^s - (-1)^s \bar{\partial}^s) \end{pmatrix} Y_{\ell m} \\ \mathbf{Y}_{s,\ell m}^B &\equiv \mathbf{D}_s^B Y_{\ell m} = -\frac{1}{2} \sqrt{\frac{(\ell-s)!}{(\ell+s)!}} \begin{pmatrix} i(\bar{\partial}^s - (-1)^s \bar{\partial}^s) \\ \bar{\partial}^s + (-1)^s \bar{\partial}^s \end{pmatrix} Y_{\ell m}\end{aligned}\tag{3.20}$$

We can then decompose the field into E-modes and B-modes as

$$a_{s,\ell m}^X = \int d\Omega \mathbf{Y}_{s,\ell m}^{X*} f(\mathbf{n}).\tag{3.21}$$

E-modes are invariant under parity transformation, $\mathbf{n} \rightarrow -\mathbf{n}$, while B-modes change sign. For the particular case of the spin-2 linear polarization field, we can write

$$\begin{aligned}a_{\ell m}^E &= -\frac{1}{2}(a_{\ell m} + a_{\ell, -m}) \\ a_{\ell m}^B &= \frac{i}{2}(a_{\ell m} - a_{\ell, -m}).\end{aligned}\tag{3.22}$$

We can also define scalar quantities (rotationally invariant) representing E-modes and B-modes in real space [478]. Figure 3.1 shows an illustration of their patterns in real space.

$$\begin{aligned}\tilde{E}(\mathbf{n}) &= \sum_{\ell m} \sqrt{\frac{(\ell+2)!}{(\ell-2)!}} a_{\ell m}^E Y_{\ell m}(\mathbf{n}), \\ \tilde{B}(\mathbf{n}) &= \sum_{\ell m} \sqrt{\frac{(\ell+2)!}{(\ell-2)!}} a_{\ell m}^B Y_{\ell m}(\mathbf{n}).\end{aligned}\tag{3.23}$$

As in the case of temperature anisotropies, the covariance matrices of the spherical harmonics coefficients are diagonal. In general, for $X, Y \in \{T, E, B\}$ we have

$$\langle a_{\ell m}^{X*} a_{\ell' m'}^Y \rangle = C_\ell^{XY} \delta_{\ell\ell'} \delta_{mm'}.\tag{3.24}$$

Conversely, the angular power spectra can be obtained from the harmonic coefficients via

$$C_\ell^{XY} = \frac{1}{2\ell+1} \sum_m \langle a_{\ell m}^{X*} a_{\ell m}^Y \rangle.\tag{3.25}$$

In the absence of non-Gaussianities, as predicted in the simplest inflation models, these angular power spectra encode all the information of the perturbations. Additionally, since B has opposite parity to T and E their cross-correlations vanish except in the presence of parity violating processes (see for instance the section 3.5.4 on Cosmic Birefringence).

The significant importance of the E/B decomposition of CMB polarization was realized when it was demonstrated in [478, 272] that at linear order, if we ignore vorticity for the reasons outlined in the previous chapter, E-modes are generated by both scalar and tensor perturbations, while B-modes can only be generated by tensor modes, and hence a detection of primordial B-modes would be the smoking gun of the primordial gravitational waves predicted by inflation. We can therefore write the power spectra of E-modes, the TE cross-correlation, and B-modes as a function of the primordial power spectra of scalar and tensor modes and the transfer functions as (scalar modes dominate over tensor modes in EE and TE spectra)

$$\begin{aligned}C_\ell^{EE} &\simeq (4\pi)^2 \int k^2 dk \mathcal{P}_{\mathcal{R}}(k) \Delta_{E\ell}^2(\eta_0, k), \\ C_\ell^{TE} &\simeq (4\pi)^2 \int k^2 dk \mathcal{P}_{\mathcal{R}}(k) \Delta_{T\ell}(\eta_0, k) \Delta_{E\ell}(\eta_0, k), \\ C_\ell^{BB} &= (4\pi)^2 \int k^2 dk \mathcal{P}_h(k) \Delta_{B\ell}^2(\eta_0, k).\end{aligned}\tag{3.26}$$

3.2 Evolution of perturbations

Now that we have covered the mathematical toolbox to describe the CMB anisotropies, we will further explore the link between the CMB observables and the physics of the primordial plasma.

3.2.1 Boltzmann equation

As seen previously, the Boltzmann equation is the master equation describing the evolution of the phase space distribution of the different species in the background as well as the perturbed spacetime, where the distribution function is expanded into a background term, \bar{f} , and a small perturbation δf ,

$$f(\eta, \mathbf{x}, E, \mathbf{n}) = \bar{f}(E) (1 + \delta f(\eta, \mathbf{x}, E, \mathbf{n})). \quad (3.27)$$

In general the Boltzmann equation for the phase space distribution of a given species a takes the following form

$$\frac{df_a}{d\eta} = \mathbf{C}[f_b], \quad (3.28)$$

where the collision operator on the right-hand side is acting, in principle, on the distribution functions of several species interacting with species a . In the following, we will specifically focus on the Boltzmann equation for photons. We can replace the perturbation of the phase space distribution by Θ , and hence derive equations of motion directly for the temperature anisotropies. The electrons and baryons are strongly coupled via Coulomb scattering and are hence treated as a single tightly coupled fluid. The dominant scattering process for photons near recombination is *Thomson scattering* off the free electrons in the plasma. Its differential cross-section is given by [165]

$$\frac{d\sigma}{d\Omega} = \frac{3}{16\pi} \sigma_T [1 + (\mathbf{n}_{\text{in}} \cdot \mathbf{n})^2], \quad (3.29)$$

where \mathbf{n}_{in} and \mathbf{n} are, respectively, the incoming and scattered photons propagation directions, and σ_T the Thomson scattering cross-section. If we take the angular average of Eq. (3.29) over incident radiation, the corresponding collision term can be written as

$$\mathbf{C}[f] = \left(\frac{d\bar{f}}{d \ln \epsilon} \right) \Gamma [\Theta - \Theta_0 - \mathbf{n} \cdot \mathbf{v}_e], \quad (3.30)$$

where $\Gamma \equiv a \bar{n}_e \sigma_T$, \bar{n}_e being the background electron density, \mathbf{v}_e the bulk velocity of electrons and Θ_0 the *monopole* of the temperature anisotropy

$$\Theta_0 \equiv \int \frac{\mathbf{n}_{\text{in}}}{4\pi} \Theta(\mathbf{n}_{\text{in}}). \quad (3.31)$$

The Boltzmann equation for photons given the metric and matter scalar perturbations introduced in the last chapter takes the following form

$$\Theta' + n^i \partial_i \Theta - \Phi' + n^i \partial_i \Psi = -\Gamma [\Theta - \Theta_0 - \mathbf{n} \cdot \mathbf{v}_e], \quad (3.32)$$

where $(\dots)'$ denotes the derivative with respect to conformal time. In Fourier space, it reads

$$\Theta' + ik\mu\Theta - \Phi' + ik\mu\Psi = -\Gamma [\Theta - \Theta_0 - i\mu v_e], \quad (3.33)$$

where $\mathbf{v}_e \equiv iv_e \hat{\mathbf{k}}$ and $\mu \equiv \hat{\mathbf{k}} \cdot \mathbf{n}$, with $\hat{\mathbf{k}}$ the unitary vector oriented in the direction of the wavevector \mathbf{k} . In accordance with statistical isotropy, the equation of

motion does not depend on \mathbf{k} and \mathbf{n} , but only on their relative orientation μ and the wavenumber k . The initial conditions for Θ also depend on μ and not on the two degrees of freedom of \mathbf{n} . This is justified by the fact that at early times, Thomson scattering is efficient ($\Gamma \gg \mathcal{H}$, with $\mathcal{H} \equiv a'/a$ the conformal Hubble parameter) and the electron-baryon fluid and photons are in the tightly-coupled regime. This forces the right-hand side in Eq. (3.32) to vanish, and the photon distribution tends to isotropise as it becomes described only by a monopole and a dipole component, $\Theta \rightarrow \Theta_0 + \mathbf{n} \cdot \mathbf{v}_e$.

3.2.2 Line of sight integral

We can relate the anisotropy observed in a given direction \mathbf{n} to the perturbations on the point of the last scattering surface located in the same direction $\mathbf{x}_* \equiv \mathbf{x}_0 + (\eta_0 - \eta_*)\mathbf{n}$. This is done via the integration of the Boltzmann equation along the line-of-sight. We first introduce two concepts that will be useful for the subsequent derivations.

The first concept is the *optical depth* between times η and η_0 , which characterizes the opacity of the Universe at a given time η , when seen from today (η_0).

$$\tau(\eta) \equiv \int_{\eta}^{\eta_0} \Gamma(\eta') d\eta'. \quad (3.34)$$

The second concept is the *visibility function* which is defined as the probability density that a CMB photon observed today has last scattered at time η . It reads

$$g(\eta) \equiv -\tau'(\eta)e^{-\tau(\eta)}, \quad (3.35)$$

and has negligible values at all times except for two spikes, the first one, very narrow around the time of recombination and the second one, which is smaller and wider, around the time of reionization.

Temperature anisotropies

By introducing the above quantities, we can rewrite the Boltzmann equation (3.32) as the following

$$\frac{d}{d\eta} e^{-\tau} (\Theta + \Psi) = e^{-\tau} (\Phi' + \Psi') + g(\Theta_0 + \Psi - \mathbf{n} \cdot \mathbf{v}_e) \equiv S(\eta, k). \quad (3.36)$$

Integrating the above equation along the line-of-sight and performing Legendre expansion, we can derive the expression of the anisotropy multipoles in Fourier space and write the transfer function as

$$\Delta_{T\ell}(\eta_0, k) = \int_0^{\eta_0} d\eta \left[\underbrace{g(\Theta_0 + \Psi)j_{\ell}(k\chi)}_{\text{Sachs-Wolfe}} - \underbrace{gv_e j'_{\ell}(k\chi)}_{\text{Doppler}} + \underbrace{e^{-\tau}(\Phi' + \Psi')j_{\ell}(k\chi)}_{\text{integrated Sachs-Wolfe}} \right], \quad (3.37)$$

where $\chi = \eta_0 - \eta$, and the derivative in the second term is with respect to the argument of the spherical Bessel function j_{ℓ} and not with respect to conformal time. We can identify three different effects highlighted in the above integral. The Sachs-Wolfe (SW) term contains both the intrinsic temperature perturbation

and the gravitational redshift which are picked up by the visibility function at the last scattering surface. The Doppler term results from the bulk motion of the baryon-photon fluid also picked up at last scattering and projected in the line-of-sight (baryons here refer to both baryons and electrons given their tight coupling). The integrated Sachs-Wolfe (ISW) term represents the integrated redshifting and blueshifting that photons experience along the line-of-sight as they travel through the Universe with dynamical metric fluctuations. Furthermore, the spherical Bessel functions and their derivatives are very peaked around $x \approx \ell$ for large values of ℓ , and hence act as delta-functions in the integral and pick up modes with $k(\eta_0 - \eta) \simeq \ell$.

Polarization anisotropies

Taking into account the polarization dependency in Thomson scattering, we can write the differential cross section as the following

$$\frac{d\sigma}{d\Omega} = \frac{3}{8\pi} \sigma_T |\hat{\mathbf{E}}' \cdot \hat{\mathbf{E}}|^2, \quad (3.38)$$

where $\hat{\mathbf{E}}$ and $\hat{\mathbf{E}}'$ denote the polarization vectors of the incoming and scattered radiation, respectively. Quadrupolar anisotropies in the incoming radiation are the only ones that can generate a net linear polarization in the outgoing radiation [253, 255]. These quadrupolar anisotropies can be sourced either by scalar perturbations, i.e. density fluctuations, or tensor perturbations, i.e. gravitational waves. As previously stated, the E/B decomposition of the linear polarization field provides a useful framework to separate the polarization patterns according to the underlying physical processes that generate them.

Similarly, line-of-sight integral equations can be derived from the Boltzmann equations for the polarization fluctuations [295, 409, 478]. In particular for B-modes, we can write the transfer function as

$$\Delta_{B\ell}(\eta_0, k) = \int_0^{\eta_0} d\eta \left[g \left(\frac{4\psi}{k\chi} + \frac{2\psi'}{k} \right) + 2g' \frac{\psi}{k} \right] j_\ell(k\chi), \quad (3.39)$$

where ψ is defined in [478] as a linear combination of quantities introduced in [379] describing the temperature and polarization perturbations generated by gravitational waves.

3.2.3 Acoustic oscillations

We can derive an oscillator equation for the evolution of the photon temperature monopole Θ_0 in the tight-coupling regime

$$\Theta_0'' + \underbrace{\frac{\mathcal{H}R}{1+R}\Theta_0'}_{\text{damping}} + \underbrace{k^2 c_s^2 \Theta_0}_{\text{pressure}} = \underbrace{-\frac{1}{3}k^2 \Psi}_{\text{gravitational force}} + \underbrace{4\Phi'' + \frac{R'}{1+R}\Phi'}_{\text{dilation effects}}, \quad (3.40)$$

where we defined the sound speed of the photon-baryon fluid

$$c_s^2 \equiv \frac{1}{3(1+R)}, \quad \text{with} \quad R \equiv \frac{3\bar{\rho}_b}{4\bar{\rho}_\gamma} = 0.6 \left(\frac{\Omega_b h^2}{0.02} \right) \left(\frac{a}{10^{-3}} \right). \quad (3.41)$$

We can define a number of relevant length scales of the problem. The comoving *Hubble radius* defined as

$$r_H \equiv (aH)^{-1}, \quad (3.42)$$

which sets the length scale below which modes start to evolve. Modes outside this scale remain frozen. The second length scale we can introduce is the *sound horizon*

$$r_s \equiv c_s (aH)^{-1}. \quad (3.43)$$

In fact, photon fluctuations only start to oscillate after they cross the sound horizon. At early times, R is negligible which yields a sound speed $c_s \approx 1/\sqrt{3}$ and hence the sound horizon is close to the Hubble radius. However as the baryon fraction becomes sizable, the sound horizon becomes significantly smaller than the Hubble radius.

To get an intuitive understanding of the effects at play in these oscillations, let us first consider an approximate solution of Eq. (3.40) where we have ignored time variations of the potentials Φ and Ψ , as well as the baryon-photon momentum density ratio R (this is valid at early times). For adiabatic initial conditions, we can write

$$\Theta_0(\eta, \mathbf{k}) = [\Theta_0(0, \mathbf{k}) + (1 + R)\Psi(\mathbf{k})] \cos(kr_s) - (1 + R)\Psi(\mathbf{k}). \quad (3.44)$$

We can first start by ignoring the effects of the baryons by setting R to zero in the above expression. In this case, the zero-point of the oscillations which corresponds to the equilibrium between pressure and gravity is displaced by $-\Psi > 0$, this means that photons in potential wells are hotter since they get blueshifted through gravitational infall. However, as we have seen previously with the Sachs-Wolfe effect, these photons get redshifted as they have to climb out of the potential wells before they reach the observer.

If we now consider the effects of the baryons, these lead to a decrease in the sound speed and hence a shrinking of the sound horizon, in addition their mass leads to a shift in the equilibrium between pressure and gravity provoking an enhancement of the compression of the fluid in a potential well, as the zero point of the oscillations is now shifted by $-(1 + R)\Psi$. As a consequence, the amplitudes of the peaks from compression become greater than those from rarefaction. Another effect of the baryons, that is not captured by the solution in Eq. (3.44), is that deep inside the sound horizon, the metric perturbations decay and the pressure decreases, but the slowly increasing R leads to the damping term in Eq. (3.40) becoming non-negligible, and hence in this regime the oscillations are decreasing.

At the time of photon decoupling, oscillation modes with different wavenumbers will be in different phases. The discrete set of modes hitting extremas at that time have the following wavenumbers

$$k_n = \frac{n\pi}{r_s(\eta_*)}. \quad (3.45)$$

These modes correspond to peaks in the CMB power spectra. Their characteristic angular scales in the power spectrum are given by

$$\ell_n = \frac{n\pi d_{A*}}{r_s(\eta_*)}, \quad (3.46)$$

where d_{A*} is the angular diameter distance to recombination. The first acoustic peak corresponds roughly to $\ell \approx 200 \approx 2^\circ$. At scales larger than the first peak, if we ignore secondary anisotropies which we'll cover later in this chapter, the temperature power spectrum should be nearly flat, reflecting the near scale invariance of the primordial power spectrum, by virtue of the fact that $\Theta_0 + \Psi = \text{const.}$ in this region. This is called the *Sachs-Wolfe plateau*.

At the smallest scales, the tight-coupling regime is no longer valid, and we need to introduce a third length scale of the problem, which is relevant for the understanding of the damping of CMB fluctuations. If photons are performing a random walk with a step length given by Γ^{-1} before they scatter off an electron, and during a small interval $\delta\eta$ the number of collisions is given by $\Gamma\delta\eta$, the effective distance travelled by photons between times 0 and η , called the *diffusion scale*, will then be given by

$$r_d \equiv \left[\int_0^\eta \Gamma^{-1} d\eta' \right]^{1/2}. \quad (3.47)$$

For small scale modes with wavenumbers $k \gtrsim k_D \equiv \pi/r_d(\eta_*)$ the coupling between photons and electrons is no longer effective, and the diffusion of photons mixes between hot and cold regions, effectively suppressing the fluctuations. This is also known as *Silk damping*, and is modeled by adding an exponential envelope, e^{-k^2/k_D^2} , to the solutions discussed previously. This translates to an angular power spectrum of the following form [250]

$$\frac{\ell(\ell+1)}{2\pi} C_\ell^{TT} \sim \exp(-(\ell/\ell_D)^{1.25}), \quad (3.48)$$

with the damping scale

$$\ell_D = \frac{2\pi d_{A*}}{r_d(\eta_*)}, \quad (3.49)$$

where the diffusion scale at recombination reads

$$r_d \approx 64.5 \left(\frac{\Omega_m h^2}{0.14} \right)^{-0.278} \left(\frac{\Omega_b h^2}{0.024} \right)^{-0.18} \text{ Mpc}. \quad (3.50)$$

Let us now consider the CMB E-mode polarization mainly sourced by scalar perturbations (neglecting tensor modes). Since these are generated by the quadrupolar anisotropies of the photon distribution, their characteristics are related to the radiative viscosity π_γ which is given by [252]

$$\pi_\gamma = \left(\frac{k}{k_D} \right)^{-1} (k_D \eta_*)^{-1} v_\gamma, \quad (3.51)$$

where v_γ is the bulk velocity of the photons. This leads to a scaling of the polarization power spectrum as ℓ/ℓ_D with a peak at the damping scale corresponding to 10% of the temperature fluctuations in amplitude, and then falling at higher scales due to the sourcing anisotropies being suppressed with diffusion damping. The peaks of the polarization power spectrum are out of phase with the temperature peaks, given that the velocity oscillates as

$$v_\gamma \propto \sin(kr_s), \quad (3.52)$$

The cross-correlation between temperature and E-mode polarization oscillates at twice the acoustic frequency, given that

$$(\Theta + \Psi)v_\gamma \propto \cos(kr_s) \sin(kr_s) \propto \sin(2kr_s). \quad (3.53)$$

3.3 Secondary anisotropies

In addition to the primary anisotropies imprinted on the CMB at recombination, additional processes affect the CMB photons along their path from the last scattering surface to our detectors, and imprint specific signatures on the CMB radiation. We refer to these as secondary anisotropies. In this section, we focus on these secondary effects and divide them in two categories: scattering and gravitational effects.

3.3.1 Scattering effects

Reionization

As we have outlined in the first chapter, the release of free electrons after the reionization of the Universe imprints a specific signal in the CMB due to the Thomson scattering of the photons off the free electrons. In particular, this effect integrated over the line-of-sight is characterized by the optical depth to reionization which is defined as [373]

$$\tau_{\text{reio}} \equiv n_H \sigma_T \int_0^{z_{\text{max}}} dz X_e(z) \frac{(1+z)^2}{H(z)}, \quad (3.54)$$

where n_H is the number density of hydrogen nuclei today, X_e is the free electron fraction, and z_{max} is the maximum redshift where reionization is still expected to contribute in a non-negligible way to the scatter of photons. The re-scattering of photons randomizes the photons propagation directions thereby washing out the anisotropies at sub-horizon scales. More specifically, at the power spectrum level, this leads to a modulation of the power spectrum amplitude by a scale independent suppression factor, $A_s e^{-2\tau_{\text{reio}}}$, at $\ell \gtrsim 20$ [391]. Conversely, at large scales ($\ell \lesssim 20$), the re-scattering of the local quadrupoles of the radiation field generates an excess of power in polarization shaped as a *bump* with amplitude scaling as τ_{reio}^2 [391].

Figure 3.2 demonstrates these features. On the left panel, showing the impact of reionization on temperature, as we go from the black line to the light blue line, the optical depth is increased thereby emphasizing the power suppression of sub-horizon scales at reionization. The dashed red line shows the effect of a variation of the amplitude of primordial scalar perturbations, A_s , demonstrating its near perfect degeneracy with the optical depth (the dashed red line coincides in nearly all scales with the dark blue line). On the right panel, showing the impact on the E-mode polarization, the power is also suppressed on scales smaller than the horizon at reionization as in temperature. However, on scales larger than the horizon we see a reionization bump. The measurement of this feature is very important for the determination of τ_{reio} as it allows to break the degeneracy with A_s as can be seen from the comparison between the dashed red line and the dark

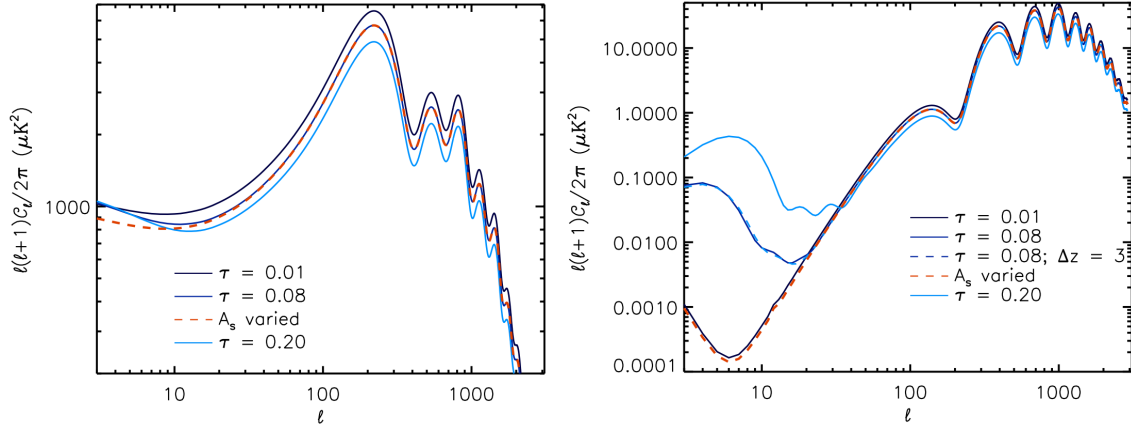


Figure 3.2 – *Left*: Impact of the optical depth to reionization on the temperature power spectrum. *Left*: Impact of the optical depth to reionization on the E-mode polarization power spectrum. *Credit*: [391]

blue one. Finally, we also notice that the E-mode signal is largely insensitive to the duration of reionization as long as the optical depth remains the same as we see no difference between the solid dark blue line and the dashed one.

Sunyaev-Zel’dovich effects

The Sunyaev-Zel’dovich effect [433] is the distortion of the CMB spectrum in a given direction on the sky, resulting from the inverse Compton scattering of photons off the free electrons present in a hot gas along the line-of-sight, typically within galaxy clusters or superclusters. We can distinguish between two processes: the first one is the thermal SZ (tSZ) effect, where the upscattering of photons is caused by the random motion of thermal electrons, and the second one is the kinetic SZ (kSZ) effect, which is due to the bulk motion of the electrons with respect to the CMB photons.

The thermal SZ effect produces a localized spectral distortion, called *Compton-y distortion*, resulting in a temperature anisotropy which can be written as

$$\frac{\Theta_{\text{tSZ}}}{T} = y \left[\frac{2\pi\nu}{T} \coth\left(\frac{\pi\nu}{T}\right) - 4 \right], \quad (3.55)$$

where T is the CMB monopole temperature, and y is the Compton- y parameter defined in a given direction \mathbf{n} as

$$y(\mathbf{n}) \equiv \int n_e \frac{T_e}{m_e} \sigma_T ds, \quad (3.56)$$

where m_e is the electron mass, n_e the electron number density, T_e the electron temperature, and ds the distance along the line-of-sight \mathbf{n} . Additionally, the thermal SZ effect can also give rise to a polarized signal (pSZ) due to the Thomson scattering of quadrupoles in the radiation field [405]. However, this signal has yet to be detected.

The kinetic SZ effect also produces a temperature anisotropy given by

$$\frac{\Theta_{\text{kSZ}}}{T} = -v_r \tau_{\text{cluster}}, \quad (3.57)$$

where τ_{cluster} is the optical depth of the medium within the cluster, and v_r is the radial bulk velocity of the electron gas.

The SZ effects are also important probes for cosmology as we will see in section 3.5.3.

3.3.2 Gravitational effects

Gravitational potential wells

The falling and climbing out of potential wells that photons experience along their path from the last scattering surface to the observer, leaves a net imprint on the CMB power spectra due to the imperfect cancellation of the induced blueshifting or redshifting. We can distinguish between two effects that operate in this way. The first one is the integrated Sachs-Wolfe effect [402] which we have briefly mentioned previously. In a static gravitational potential well, photons are first blueshifted as they fall in it then redshifted as they climb out of it, and the two effects are perfectly compensated. However, if the potential evolves with time then the symmetry between the gravitational infall and the climbing is broken, the blueshifting and redshifting are no longer compensated and photons get a net shift in their energies. In fact, one can show with Einstein equations that after the change in the equation of state of the Universe at the time of radiation and matter equality, sub-sound-horizon metric fluctuations do not stabilize quickly and keep evolving beyond photon decoupling, hence leading to a specific imprint on the CMB called the Early Integrated Sachs-Wolfe (EISW) which is strongest for scales crossing the sound horizon near photon decoupling. In terms of angular scales, the EISW mainly enhances the first peak of the temperature power spectrum near $\ell \approx 200$. Additionally, one can also show that the transition to the dark energy dominated regime also triggers variations of the metric fluctuations on all scales thereby imprinting a signal on the CMB called the Late Integrated Sachs-Wolfe (LISW) which is strongest at the largest observable scales. At the power spectrum level, the LISW induces a tilt in the Sachs-Wolfe plateau. Apart from the integrated Sachs-Wolfe effect, second order effects triggered by the emergence of non-linear structures leading to rapidly evolving potentials can also produce secondary anisotropies through the same process. This is referred to as the Rees-Sciama effect [390].

Gravitational lensing effect

The clustering of large-scale structures have a gravitational effect on CMB photons, as it remaps the radiation field thereby modifying the temperature and polarization anisotropies. This effect is called *weak-gravitational lensing*. The term *weak* refers to the fact that the effect is very subtle and cannot be detected on an individual galaxy, but is rather quantified by averaging over a large statistical sample.

Formally, we can model the effect by a two-dimensional vector field α on the sky, which gives the deflection angle in every direction

$$\mathbf{n} \longrightarrow \mathbf{n}' \equiv \mathbf{n} + \alpha(\mathbf{n}). \quad (3.58)$$

We can decompose the deflection field α into a longitudinal term given by the lensing potential ϕ and a transversal term given by the lensing curl-potential Ω

$$\begin{aligned} \mathbf{n}' &= \mathbf{n} + \alpha \\ &\equiv \mathbf{n} + \nabla\phi + \nabla \times \Omega \\ &= \begin{pmatrix} \theta \\ \phi \end{pmatrix} + \begin{pmatrix} \partial_\theta \phi \\ \partial_\varphi \phi \end{pmatrix} + \begin{pmatrix} \partial_\varphi \Omega \\ -\partial_\theta \Omega \end{pmatrix} \\ &\equiv \mathbf{M}(\mathbf{n}) \cdot \mathbf{n}. \end{aligned} \quad (3.59)$$

The matrix \mathbf{M} is called the *magnification matrix*. In the weak lensing regime it can be written as

$$\begin{aligned} \mathbf{M}(\mathbf{n}) &= \begin{pmatrix} \frac{\partial \mathbf{n}'_\theta}{\partial \theta} & \frac{\partial \mathbf{n}'_\theta}{\partial \varphi} \\ \frac{\partial \mathbf{n}'_\varphi}{\partial \theta} & \frac{\partial \mathbf{n}'_\varphi}{\partial \varphi} \end{pmatrix} = \begin{pmatrix} 1 + \partial_\theta^2 \phi + \partial_\theta \partial_\varphi \Omega & \partial_\theta \partial_\varphi \phi + \partial_\varphi^2 \Omega \\ \partial_\theta \partial_\varphi \phi - \partial_\varphi^2 \Omega & 1 + \partial_\varphi^2 \phi - \partial_\theta \partial_\varphi \Omega \end{pmatrix} \\ &\equiv \begin{pmatrix} 1 - \kappa + \gamma_Q & \gamma_U - \omega \\ \gamma_U + \omega & 1 - \kappa - \gamma_Q \end{pmatrix}. \end{aligned} \quad (3.60)$$

All of the parameters introduced above depend on the sky direction \mathbf{n} . The convergence κ characterizes the magnification of the features with lensing, i.e. given a feature of angular size θ at the last scattering surface, its apparent angular size would be $(1 + \kappa)\theta$, ω gives its clock-wise rotation in radians. The shear component γ_Q characterizes the stretching of features by $(1 + \gamma_Q)$ along the horizontal axis and their compression along the vertical axis by $(1 - \gamma_Q)$. While γ_U exerts similar effects along axes rotated clockwise by 45° .

Let us now relate the lensing effect to metric perturbations. At first order, we can neglect the curl potential, and hence our task reduces to relating the lensing potential ϕ to the gravitational potential. More specifically we introduce the so-called *Weyl potential* defined as $\Psi_W \equiv (\Phi + \Psi)/2$.

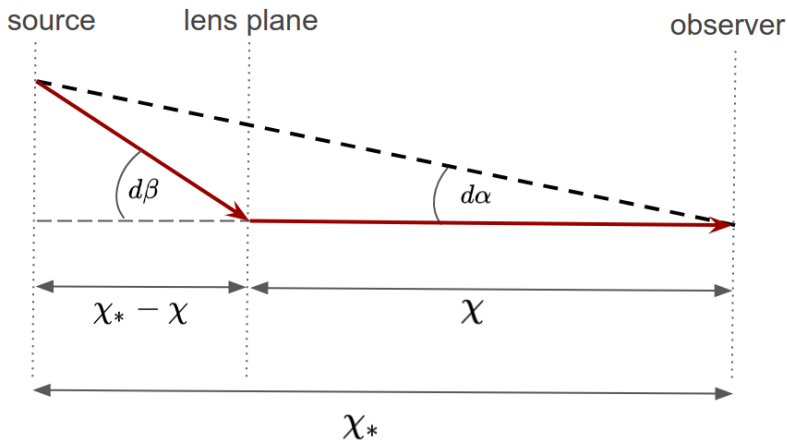


Figure 3.3 – Geometrical representation of a single deflection induced by gravitational lensing.

One can show that the infinitesimal deflection angle $d\beta$ along the trajectory of photons during conformal time $d\eta$ is given by

$$d\beta(\eta, \mathbf{x}) = -\frac{2}{S_k(\chi)} \nabla \Psi_W(\eta, \mathbf{x}) d\eta, \quad (3.61)$$

where χ is the comoving distance to the source. If we consider a light ray emitted at time $\eta = 0$, and reaching an observer located at $\mathbf{x} = 0$ at time η_* from a direction \mathbf{n} , and experiencing a single deflection event at time η , then given that the deflection angle of the source as seen by the observer, $d\alpha$ subtends the same comoving size as the deflection angle at the lensing plane, $d\beta$ (see figure 3.3) we can write

$$S_k(\chi_*) d\alpha(\eta_*, 0) = S_k(\chi_* - \chi) d\beta(\eta, \mathbf{n}(\chi_* - \chi)). \quad (3.62)$$

We obtain the deflection angle in a sky direction \mathbf{n} , by integrating the above equation along comoving distance χ (i.e. to the lensing plane)

$$\alpha(\mathbf{n}) = -2 \int_0^{\chi_*} d\chi \frac{S_k(\chi_* - \chi)}{S_k(\chi_*) S_k(\chi)} \nabla \Psi_W(\eta, \mathbf{s}(\eta, \mathbf{n})), \quad (3.63)$$

where $\mathbf{s}(\eta, \mathbf{n})$ parametrizes the perturbed photon geodesic reaching the observer from direction \mathbf{n} . At first order, we can assume that photon geodesics follow a straight unperturbed line, $\chi\mathbf{n}$. This is referred to as the *Born approximation*. Under this assumption we can finally write the lensing potential of CMB photons as a simple weighted line-of-sight integral of the gravitational potential

$$\phi(\mathbf{n}) = -2 \int_0^{\chi_*} d\chi \frac{S_k(\chi_* - \chi)}{S_k(\chi_*) S_k(\chi)} \Psi_W(\eta, \chi\mathbf{n}). \quad (3.64)$$

CMB lensing potential power spectrum

Similarly to the temperature and polarization anisotropies, we can compute the power spectrum of the lensing field in order to extract the physical information encoded in its statistical properties. We can expand the lensing potential in spherical harmonic space as

$$\begin{aligned} \phi_{LM} &= \int d\Omega Y_{LM}^* \phi(\mathbf{n}) \\ \langle \phi_{LM} \phi_{L'M'}^* \rangle &= \delta_{LL'} \delta_{MM'} C_L^{\phi\phi}. \end{aligned} \quad (3.65)$$

We can define the power spectrum of the gravitational potential (assuming statistical homogeneity) as the following

$$\langle \Psi_W(\eta, \mathbf{k}) \Psi_W^*(\eta', \mathbf{k}') \rangle \equiv \frac{2\pi^2}{k^3} \mathcal{P}_\Psi(\eta, \eta', k) \delta(\mathbf{k} - \mathbf{k}'), \quad (3.66)$$

from which we can infer [311]

$$C_L^{\phi\phi} = 16\pi \int \frac{dk}{k} \int_0^{\chi_*} d\chi \int_0^{\chi_*} d\chi' \mathcal{P}_\Psi(\eta_0 - \chi, \eta_0 - \chi', k) j_L(k\chi) j_L(k\chi') W(\chi) W(\chi'), \quad (3.67)$$

where we have defined W as

$$W(\chi) \equiv \frac{S_k(\chi_* - \chi)}{S_k(\chi_*) S_k(\chi)}. \quad (3.68)$$

In the linear regime, we can write the gravitational potential as the product of the primordial spectrum and a transfer function

$$\Psi_W(\eta, \mathbf{k}) = \Delta_\Psi(\eta, k) \mathcal{R}(\mathbf{k}), \quad (3.69)$$

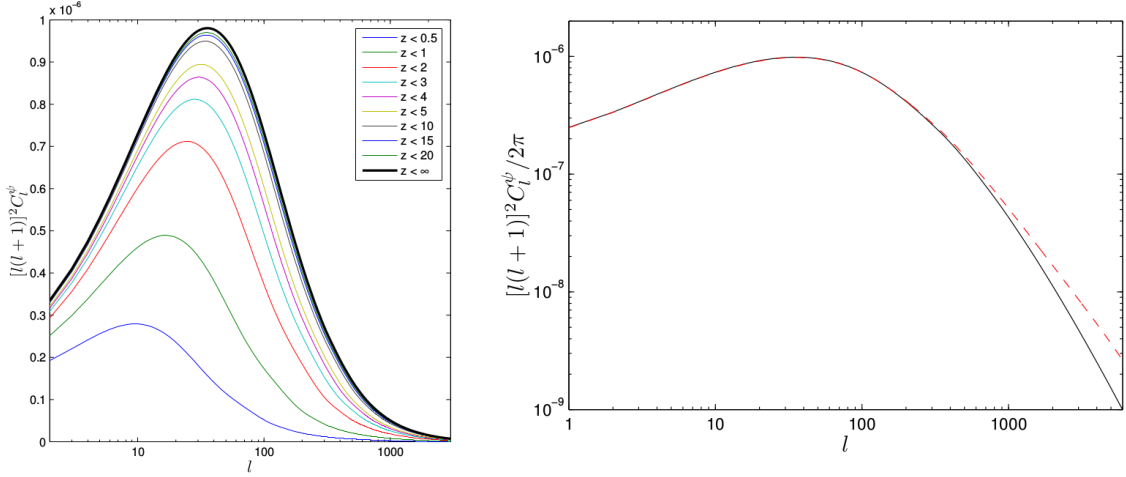


Figure 3.4 – *Left*: The cumulative contributions to the angular lensing power spectrum as a function of maximum redshifts. The black curve indicates the total lensing power spectrum. *Right*: A comparison between the angular lensing power spectrum in the linear regime and non-linear corrections from HALOFIT [422]. *Credit*: [311]

in this case the angular power spectrum can be rewritten as

$$C_L^{\phi\phi} = 16\pi \int \frac{dk}{k} \mathcal{P}_{\mathcal{R}}(k) \left[\int_0^{\chi_*} d\chi \Delta_{\Psi}(\eta_0 - \chi, k) j_{\ell}(k\chi) W(\chi) \right]^2. \quad (3.70)$$

In the above expression, we can see that the lensing power spectrum is affected by the primordial power spectrum of scalar perturbations $\mathcal{P}_{\mathcal{R}}$, how it evolves through the transfer function and the geometrical term given by W and spherical Bessel functions. In particular, the lensing potential measured from the CMB receive significant contributions from high redshifts, as can be seen in the left panel of figure 3.4, thereby circumventing the limits in typical redshifts of galaxy surveys ($z \lesssim 1$). In addition, the non-linear evolution of galaxy clustering at late-times can also affect lensing power spectrum as shown in the right panel figure 3.4. The non-linear effects are typically computed via N-body simulations [153, 201, 124, 118], although analytic formulae [422] can yield percent level accuracy up to $\ell \sim 2000$.

Effects on the CMB observables

In this subsection, we write down the effects of the weak gravitational lensing on the CMB observables. Let us consider the temperature T , and the spin-2 polarization field $\pm P = Q \pm iU$. We indicate lensed quantities by a tilde. Neglecting the curl potential, we can write the lensed CMB fields as a function of the unlensed ones as

$$\begin{aligned} \tilde{T}(\mathbf{n}) &= T(\mathbf{n} + \nabla\phi(\mathbf{n})), \\ \pm \tilde{P}(\mathbf{n}) &= \pm P(\mathbf{n} + \nabla\phi(\mathbf{n})). \end{aligned} \quad (3.71)$$

Performing a second order Taylor expansion in ϕ around the unperturbed line-of-sight \mathbf{n} , one can show that the spherical harmonic coefficients for both the lensed

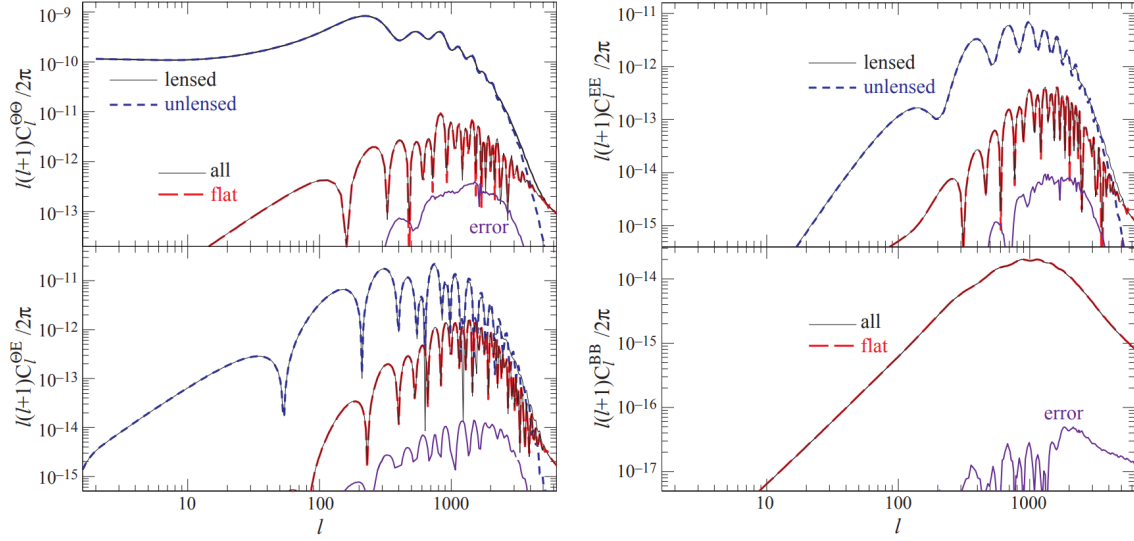


Figure 3.5 – *Left*: TT and TE lensed and unlensed spectra and their differences in all-sky and flat-sky approximation as well as the error induced by the latter. *Right*: EE and BB lensed and unlensed spectra with their differences. Here we assume no tensor modes, and hence the unlensed BB vanishes. *Credit*: [251]

temperature and polarization fields can be written as [251]

$$\begin{aligned}\tilde{a}_{\ell m}^T &= a_{\ell m}^T + \sum_{\ell' m'} \sum_{LM} \phi_{LM} a_{\ell' m'}^T \left(F_{\ell m LM \ell' m'} + \frac{1}{2} \phi_{L' M'}^* G_{\ell m LM \ell' m' L' M'} \right), \\ \pm 2 \tilde{a}_{\ell m} &= \pm 2 a_{\ell m} + \sum_{\ell' m'} \sum_{LM} \phi_{LM} \pm 2 a_{\ell' m'} \left(\pm 2 F_{\ell m LM \ell' m'} + \frac{1}{2} \phi_{L' M'}^* \pm 2 G_{\ell m LM \ell' m' L' M'} \right).\end{aligned}\tag{3.72}$$

where the precise definition of the convolution kernels F and G can be found in [251]. Finally, we can write down the expressions of the lensed CMB power spectra as a function of the unlensed ones and to the lowest order in $C_\ell^{\phi\phi}$ [251]

$$\begin{aligned}\tilde{C}_\ell^{TT} &= (1 - \ell(\ell+1)R_\phi) C_\ell^{TT} + \sum_{L\ell'} \frac{F_{\ell L \ell'}^2}{2\ell+1} C_L^{\phi\phi} C_{\ell'}^{TT}, \\ \tilde{C}_\ell^{TE} &= (1 - (\ell^2 + \ell - 2)R_\phi) C_\ell^{TE} + \sum_{L\ell'} \frac{F_{\ell L \ell'} F_{\ell L \ell'}}{2\ell+1} C_L^{\phi\phi} C_{\ell'}^{TE}, \\ \tilde{C}_\ell^{EE} &= (1 - (\ell^2 + \ell - 4)R_\phi) C_\ell^{EE} + \frac{1}{2} \sum_{L\ell'} \frac{2F_{\ell L \ell'}^2}{2\ell+1} C_L^{\phi\phi} (\epsilon_{\ell L \ell'} C_{\ell'}^{EE} + \beta_{\ell L \ell'} C_{\ell'}^{BB}), \\ \tilde{C}_\ell^{TB} &= (1 - (\ell^2 + \ell - 2)R_\phi) C_\ell^{TB} + \sum_{L\ell'} \frac{F_{\ell L \ell'} F_{\ell L \ell'}}{2\ell+1} C_L^{\phi\phi} C_{\ell'}^{TB}, \\ \tilde{C}_\ell^{EB} &= (1 - (\ell^2 + \ell - 4)R_\phi) C_\ell^{EB} + \frac{1}{2} \sum_{L\ell'} \frac{2F_{\ell L \ell'}^2}{2\ell+1} C_L^{\phi\phi} (\epsilon_{\ell L \ell'} C_{\ell'}^{EB} - \beta_{\ell L \ell'} C_{\ell'}^{EB}), \\ \tilde{C}_\ell^{BB} &= (1 - (\ell^2 + \ell - 4)R_\phi) C_\ell^{BB} + \frac{1}{2} \sum_{L\ell'} \frac{2F_{\ell L \ell'}^2}{2\ell+1} C_L^{\phi\phi} (\beta_{\ell L \ell'} C_{\ell'}^{EE} + \epsilon_{\ell L \ell'} C_{\ell'}^{BB}),\end{aligned}\tag{3.73}$$

where we have defined the tensors ϵ and β as

$$\begin{aligned}\epsilon_{\ell L \ell'} &\equiv \frac{1 + (-1)^{\ell+L+\ell'}}{2}, \\ \beta_{\ell L \ell'} &\equiv \frac{1 - (-1)^{\ell+L+\ell'}}{2},\end{aligned}\tag{3.74}$$

and introduced

$$R_\phi \equiv \frac{1}{2} \sum_L L(L+1) \frac{2L+1}{4\pi} C_L^{\phi\phi}.\tag{3.75}$$

The above expressions obtained from a Taylor expansion are only correct up to a 10% accuracy level. More accurate results can be obtained by numerical codes such as CAMB for example [134].

An important effect one can readily infer from the above expressions is that weak lensing induces a mixing between the two parity modes E and B. In particular, even if primordial scalar perturbations do not source any B-modes, we get a significant leakage from E-modes through its convolution with the lensing potential. Given the relative amplitudes of primordial E-modes and B-modes, this becomes particularly problematic for the detection of the latter. This effect can be (partially) corrected for in a specific procedure called *delensing* (Cf. Chapter 5 Section 5.2.3).

Figure 3.5 shows a comparison between lensed and unlensed temperature and polarization auto and cross spectra, as well as the errors induced by calculations that assume a flat-sky approximation. A major effect of lensing is the smoothing of the sharp acoustic peaks, that is revealed by the fact that the difference between lensed and unlensed spectra is oscillating. In addition, the difference between the two signals also peaks in power around $\ell \sim 1000$, which indicates another effect of lensing which is the redistribution of power to smaller scales.

3.4 Cosmological parameters and observational constraints

We can summarize all the effects that we have previously covered, including both the initial anisotropies characterized by the primordial spectra and how they evolve up to recombination, and the secondary effects that impact CMB photons along their path between recombination and us, into a handful of parameters. In the standard flat Λ CDM model, there are six parameters. We consider an extension of this model where there are three additional parameters, N_{eff} , M_ν and r . The model is then characterized by the following 9 parameters,

A_s , the amplitude of primordial power spectrum of scalar perturbations,

n_s , the tilt of the primordial power spectrum of scalar perturbations,

Ω_b , the energy density of baryons today,

Ω_{cdm} , the energy density of cold dark matter today,

Ω_Λ , the energy density of dark energy today,

- τ_{reio} , the optical depth to reionization,
- N_{eff} , the effective number of relativistic species,
- M_ν , the total mass of neutrinos,
- r , the tensor-to-scalar ratio.

Table 3.1 summarizes the best-fit values of these parameters or their upper bounds from the latest Planck results (combined with BAO) [373, 454], the upper bounds on r are derived by combining Planck data with BICEP2/Keck 2015 [39]. Figure 3.6 shows the state of the art of the measurements of the temperature and polarization auto and cross spectra as well as the lensing power spectrum. As of today, the temperature anisotropies have been characterized with high precision up to fundamental limits in a wide range of multipoles ($\ell \simeq 3000$), there is still room for improvement in the measurements of the E-mode polarization, and we are still in the beginnings of the scientific exploitation of the B-mode signal. In fact, the first detection of lensing B-modes through cross-correlations with the Cosmic Infrared Background (CIB) [235] have been conducted by the South Pole Telescope in 2013 [230], followed the next year by a direct measurement of the lensing B-mode signal by the POLARBEAR experiment [42]. The primordial B-mode signal, however, has yet to be detected.

base Λ CDM parameters	68% CL
$\ln(10^{10} A_s)$	3.047 ± 0.014
n_s	0.9665 ± 0.0038
$\Omega_b h^2$	0.02242 ± 0.00014
$\Omega_{\text{cdm}} h^2$	0.11933 ± 0.00091
Ω_Λ	0.6889 ± 0.0056
τ_{reio}	0.051 ± 0.006
Extensions	95% CL
M_ν [eV]	< 0.12
N_{eff}	$2.99^{+0.34}_{-0.33}$
r	< 0.044

Table 3.1 – Summary of constraints on cosmological parameters from results of [373, 454, 39].

Let us examine in detail how the three added parameters in our extension of the flat Λ CDM model affect the CMB observables.

Effects from neutrinos

Neutrinos can have observable signatures on the CMB characterized by two different parameters, N_{eff} and M_ν . Let us discuss the effects of each one of them separately.

Effective number of relativistic species, N_{eff}

Neutrinos contribute to N_{eff} when they are in the relativistic regime. As noted before in the first chapter, N_{eff} can also receive contributions from other light relics

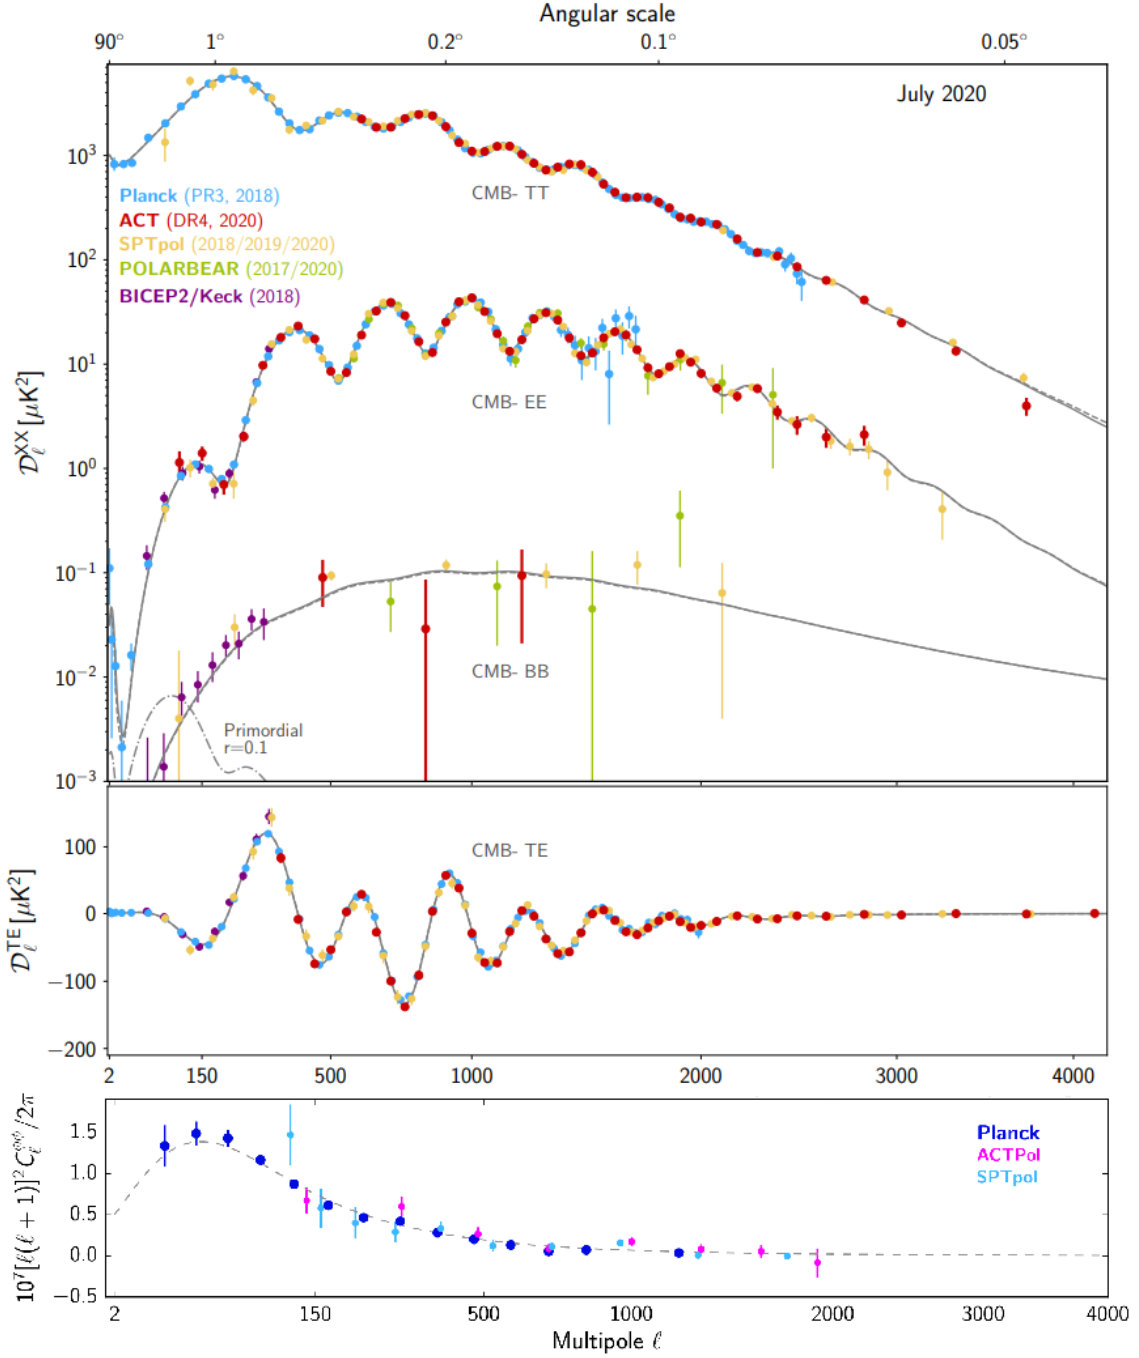


Figure 3.6 – A compilation of state of the art measurements of CMB angular power spectra and the lensing power spectrum from multiple ground and satellite experiments. *Credit:* Planck & ACT collaborations [143].

in the early Universe and hence its measurement is also important for the search of these particles. N_{eff} can have both background and perturbative effects on the CMB observables. The background effect can be considered in two ways. First, if we keep the densities of the other species fixed and increase N_{eff} , we decrease the redshift of the radiation-to-matter equality resulting in the enhancement of the overall acoustic peaks. This effect can be produced with the variation of other parameters, Ω_{cdm} for instance, and hence is not a true signature of N_{eff} . A better way to emphasize the background effect of N_{eff} , is to increase all densities propor-

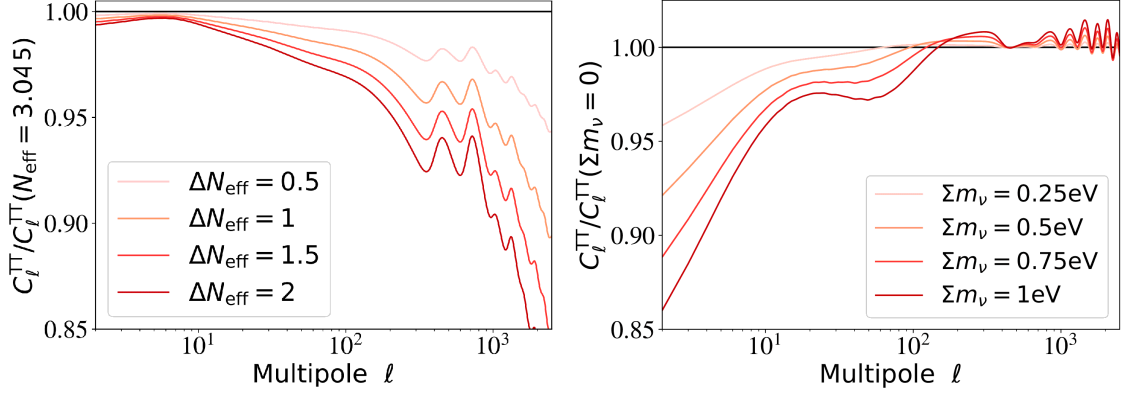


Figure 3.7 – *Left*: Ratio of C_ℓ^{TT} for different values of $\Delta N_{\text{eff}} = N_{\text{eff}} - 3.045$ over a reference model with $\Delta N_{\text{eff}} = 0$. *Right*: Ratio of C_ℓ^{TT} for different values of the total mass of neutrinos over a reference model with massless neutrinos. *Credit*: Particle Data Group [219].

tionally to keep the redshift of the radiation-to-matter equality fixed. In this case, as can be seen in the left panel of figure 3.7 the effect is an increase in the diffusion scale at decoupling, causing power suppression at higher multipoles. The perturbative effect of N_{eff} is a phase shift of the acoustic peaks towards larger scales and a decrease in their amplitudes. This is due to the fact that the photon-baryon fluid feels the gravitational influence of free-streaming neutrinos. The smoothing of acoustic peaks results from the smoother distribution of neutrinos, while their phase shift, which is a distinctive feature of N_{eff} not degenerate with other parameters, results from the fact that neutrino perturbations are propagating at velocities exceeding the sound speed of the photon-baryon fluid [81, 219]. Polarization measurements are particularly important for N_{eff} since they allow a more precise determination of the phase shift because of the sharper acoustic peaks, and also bring complementary information to break degeneracies affecting the damping tail [84].

Total mass of neutrinos, M_ν

The second effect by which neutrinos can impact CMB observables is their total mass, M_ν . Given current constraints on neutrino mass, they become non-relativistic after photon decoupling, hence increasing M_ν while keeping the density of other matter components fixed would not affect the early cosmological evolution preceding the non-relativistic transition. However it can still affect the CMB temperature and polarization power spectra through four main effects. The first one stems from the fact that increasing the total matter density (through increasing M_ν) affects the late background evolution and increases the angular diameter distance to recombination d_{A*} , and it also affects the late ISW effect through a modification of the redshift of matter-to- Λ equality. By allowing h and Ω_Λ to vary we can manage to keep one of the above two quantities fixed. In the right panel of figure 3.7, the Hubble parameter is allowed to vary to keep a fixed d_{A*} . The result is a decrease of the LISW effect which manifests itself with the power suppression at $\ell \leq 20$. Given that we can play on parameters M_ν and h to fix d_{A*} , the two parameters are degenerate. The second effect comes from the fact

that the non-relativistic transition of neutrinos causes a change in the equation of state of the Universe thereby sourcing small metric fluctuations, and can show in the CMB spectra through the early ISW effect if the transition happens shortly after photon decoupling. In figure 3.7 (right panel), this effect explains the dip seen for multipoles between 20 and 200. The third effect is that a higher M_ν reduces the effect of weak lensing due to a suppression of the matter power spectrum induced by neutrino free streaming. This effect is responsible for the oscillations seen in figure 3.7 for $\ell \geq 200$. It can also be probed by the measurement of the lensing power spectrum itself. The fourth effect is due to the neutrinos with the smallest momenta which transition to the non-relativistic regime before the average ones, affecting photon perturbations gravitationally and leading to an enhancement in C_ℓ^{TT} for $\ell \geq 500$. This effect is not very noticeable in figure 3.7 however, due to its compensation with the lensing effect. Current CMB constraints on neutrino mass are driven by the first and third effects. Combining these data with BAO brings additional information on the angular diameter distance at small redshift and provides very strong bounds on M_ν [373, 219]. Constraints from the lensing power spectrum require a measurement of A_s from CMB spectra and hence are limited by τ_{reio} due to the degeneracy between the two parameters at $\ell \geq 20$. In the future, experiments such as CLASS [232] and LiteBIRD [330] will target a precise measurement of the reionization bump and would thereby significantly improve current limits. In addition, small scale measurements from stage-3 CMB experiments would also allow to improve the uncertainty on τ_{reio} through the measurement of the kSZ signal [117].

Effects from primordial gravitational waves

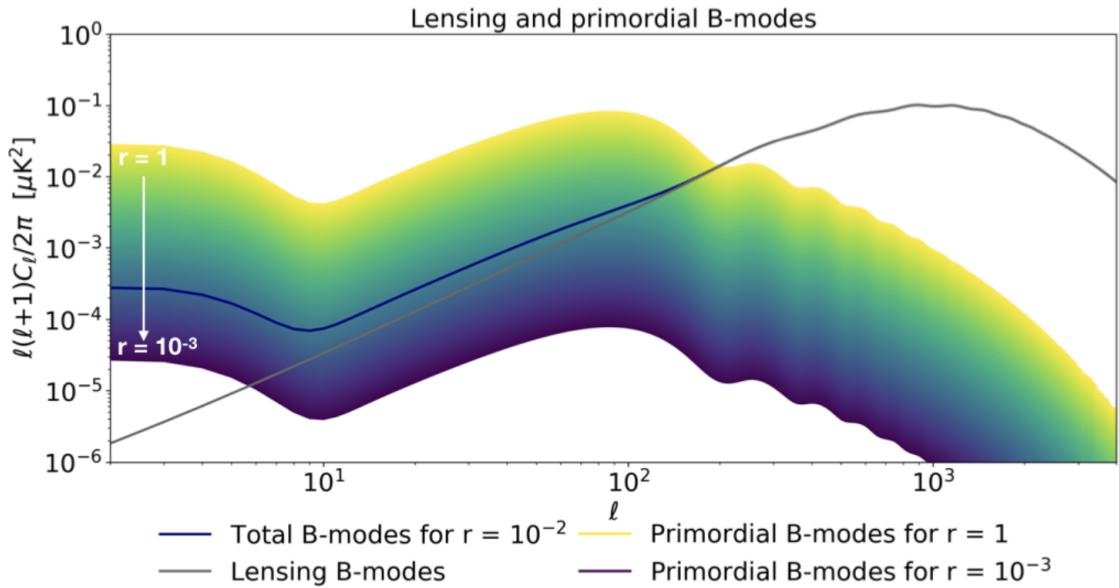


Figure 3.8 – Lensing and primordial B-modes for different values of r . *Credit:* Dominic Beck.

Tensor modes in the temperature and E-mode polarization CMB anisotropies are several orders of magnitude fainter than scalar modes. The way to detect

primordial gravitational waves hence goes through the measurement of the large-scale B-mode power spectrum. The lensing signal gets weaker at large scales, and the primordial signal picks up around degree scales. The amplitude of the primordial B-mode spectrum is proportional to the tensor-to-scalar ratio r

$$C_\ell^{BB} = r C_\ell^{BB}|_{r=1}. \quad (3.76)$$

Figure 3.8 shows the primordial BB spectrum for different values of r compared with the lensing B-modes spectrum. Current CMB experiments target a measurement of the recombination bump around $\ell \sim 80$, and the tightest upper limit to date is $r < 0.044$ at 95% CL [454].

3.5 Beyond anisotropies

In this last section, we give a short overview of other scientific targets that can be pursued with the study of the CMB signal beyond what is accessible with the measurement of the angular power spectra.

3.5.1 Primordial non-Gaussianity

In our previous developments, we assumed that the CMB temperature and polarization fields were Gaussian. Under this assumption, all the physically relevant information is encoded in the power spectra. However, if we relax this assumption and consider some level of non-Gaussianity as expected by a number of inflation models as well as alternatives, then computing higher-order statistics gives us access to additional information thereby allowing us to derive strong constraints that should discriminate between the predictions of different models. One should note that galactic foregrounds as well as secondary anisotropies such as SZ, ISW and Rees-Sciama effects, in addition to lensing and extragalactic point sources can also be a source of non-Gaussianity, and therefore need to be properly accounted for in the search for primordial non-Gaussianities.

Higher-order statistics

In the following, we will mainly focus on the so-called *bispectrum* which is the Fourier transform of the three-point correlation function

$$\langle \mathcal{R}(\mathbf{k}_1) \mathcal{R}(\mathbf{k}_2) \mathcal{R}(\mathbf{k}_3) \rangle = (2\pi)^3 \delta(\mathbf{k}_1 + \mathbf{k}_2 + \mathbf{k}_3) B_{\mathcal{R}}(k_1, k_2, k_3), \quad (3.77)$$

where the delta function ensures that the three wavevectors form a triangle, $\mathbf{k}_1 + \mathbf{k}_2 + \mathbf{k}_3 = 0$, as a consequence of the translational invariance of the background.

A non vanishing bispectrum can be probed via the CMB through the measurement of the angular bispectrum defined as

$$B_{m_1 m_2 m_3}^{\ell_1 \ell_2 \ell_3} \equiv \langle a_{\ell_1 m_1} a_{\ell_2 m_2} a_{\ell_3 m_3} \rangle. \quad (3.78)$$

Injecting in the above the expression of the harmonic coefficients as a function of \mathcal{R} and the transfer function, we can relate the CMB bispectrum to the primordial one. We can write [291, 193]

$$B_{m_1 m_2 m_3}^{\ell_1 \ell_2 \ell_3} = \mathcal{G}_{m_1 m_2 m_3}^{\ell_1 \ell_2 \ell_3} b_{\ell_1 \ell_2 \ell_3}, \quad (3.79)$$

where $\mathcal{G}_{m_1 m_2 m_3}^{\ell_1 \ell_2 \ell_3}$ is the *Gaunt integral* defined as

$$\mathcal{G}_{m_1 m_2 m_3}^{\ell_1 \ell_2 \ell_3} \equiv \int d\Omega Y_{\ell_1 m_1}(\mathbf{n}) Y_{\ell_2 m_2}(\mathbf{n}) Y_{\ell_3 m_3}(\mathbf{n}), \quad (3.80)$$

and $b_{\ell_1 \ell_2 \ell_3}$ is the *reduced bispectrum* which contains all the physical information in $B_{m_1 m_2 m_3}^{\ell_1 \ell_2 \ell_3}$, and is given by

$$b_{\ell_1 \ell_2 \ell_3} = \left(\frac{2}{\pi}\right)^3 \int dk_1 dk_2 dk_3 (k_1 k_2 k_3)^2 B_{\mathcal{R}}(k_1, k_2, k_3) \Delta_{\ell}(k_1) \Delta_{\ell}(k_2) \Delta_{\ell}(k_3) \times \int_0^\infty r^2 dr j_{\ell}(k_1 r) j_{\ell}(k_2 r) j_{\ell}(k_3 r). \quad (3.81)$$

Non-Gaussianity shapes and constraints

Primordial non-Gaussianity comes in different shapes, each corresponding to a class of physically motivated models. Typically models can be associated with a bispectrum shape that peaks at a specific triangle configuration. For a detailed review of inflation models and their predicted non-Gaussianity, one may refer to [79]. We will focus here on three specific shapes which were constrained by Planck [375]. One first way of modeling primordial non-Gaussianity is the so-called *local non-Gaussianity* where we introduce a non-linear correction term to the Gaussian perturbation \mathcal{R}_g

$$\mathcal{R}(\mathbf{x}) = \mathcal{R}_g(\mathbf{x}) + f_{\text{NL}}^{\text{local}} [\mathcal{R}_g(\mathbf{x})^2 - \langle \mathcal{R}_g(\mathbf{x})^2 \rangle]. \quad (3.82)$$

This yields a bispectrum of the following form

$$B_{\mathcal{R}}^{\text{local}}(k_1, k_2, k_3) = 2f_{\text{NL}}^{\text{local}} [P_{\mathcal{R}}(k_1)P_{\mathcal{R}}(k_2) + P_{\mathcal{R}}(k_2)P_{\mathcal{R}}(k_3) + P_{\mathcal{R}}(k_3)P_{\mathcal{R}}(k_1)] \\ = 2A_s^2 f_{\text{NL}}^{\text{local}} \left[\frac{1}{(k_1 k_2)^{4-n_s}} + \text{cycl.} \right]. \quad (3.83)$$

For local non-Gaussianity the bispectrum peaks when one of the wavenumbers is very small compared with the other two, e.g. $k_3 \ll k_1 \approx k_2$. This is called the *squeezed triangle* configuration referring to the shape of the triangle formed by the three wavevectors. This mode typically corresponds to multi-field inflation models [80, 92, 460], the curvaton model [314, 321], inhomogenous reheating [171, 287] and bouncing cosmology models [147, 299, 309].

Other models of non-Gaussianity correspond to a peak of the signal in the *equilateral triangle* configuration ($k_1 = k_2 = k_3$). The corresponding expression of the bispectrum is given by [411]

$$B_{\mathcal{R}}^{\text{equil}}(k_1, k_2, k_3) = 6A_s^2 f_{\text{NL}}^{\text{equil}} \times \left[-\frac{1}{(k_1 k_2)^{4-n_s}} - \frac{1}{(k_2 k_3)^{4-n_s}} - \frac{1}{(k_1 k_3)^{4-n_s}} - \frac{2}{(k_1 k_2 k_3)^{2(4-n_s)/3}} \right. \\ \left. + \left(\frac{1}{k_1^{(4-n_s)/3} k_2^{2(4-n_s)/3} k_3^{4-n_s}} + 5 \text{ perms.} \right) \right]. \quad (3.84)$$

We can also define the orthogonal to the equilateral mode, which is defined by

$$\begin{aligned}
 B_{\mathcal{R}}^{\text{ortho}}(k_1, k_2, k_3) = & 6A_s^2 f_{\text{NL}}^{\text{ortho}} \\
 & \times \left[-\frac{3}{(k_1 k_2)^{4-n_s}} - \frac{3}{(k_2 k_3)^{4-n_s}} - \frac{3}{(k_1 k_3)^{4-n_s}} - \frac{8}{(k_1 k_2 k_3)^{2(4-n_s)/3}} \right. \\
 & \left. + \left(\frac{3}{k_1^{(4-n_s)/3} k_2^{2(4-n_s)/3} k_3^{4-n_s}} + 5 \text{ perms.} \right) \right].
 \end{aligned} \tag{3.85}$$

These models correspond typically to single-field inflation models with higher-derivative interactions [138, 64, 418].

In the Planck 2018 data release [375], the following constraints were derived on the above three shapes of non-Gaussianity

$$\begin{aligned}
 f_{\text{NL}}^{\text{local}} &= -0.9 \pm 5.1 \\
 f_{\text{NL}}^{\text{equil}} &= -26 \pm 47 \\
 f_{\text{NL}}^{\text{ortho}} &= -38 \pm 24
 \end{aligned} \tag{3.86}$$

These results, corresponding to the 68% CL (statistical), are consistent with the base Λ CDM model ($f_{\text{NL}} = 0$) as well as predictions from standard single-field slow roll inflation models where non-Gaussianities should be unobservably small. However, as mentioned in the previous chapter, the sensitivities still need to be improved before we start putting some tension on the competing models.

3.5.2 Spectral distortions

As we have mentioned before, the CMB spectrum was measured by the FIRAS instrument onboard the COBE satellite [200] and was shown to be indistinguishable from a blackbody spectrum at temperature $T_0 = 2.72548 \pm 0.00057$ K [198] up to the 10^{-5} level. However several processes are believed to generate departures from a pure blackbody spectrum, encoding key information about the thermal history of the Universe. Characterizing these *spectral distortions* can be a powerful way to constrain cosmological models.

Main types of spectral distortions

In the early Universe, the relevant interactions between photons and electrons are the Compton scattering ($e^- + \gamma \rightleftharpoons e^- + \gamma$) where the number of photons is conserved, and the *double Compton scattering* ($e^- + \gamma \rightleftharpoons e^- + \gamma + \gamma'$) as well as the *Bremsstrahlung radiation* ($e^- \rightleftharpoons e^- + \gamma$)—also called *free-free*—where photons can be produced (or destroyed). At high redshifts ($z > \text{few} \times 10^6$) all three interactions are very efficient, photons thermalize efficiently and thus follow the Bose-Einstein equilibrium distribution, with a null chemical potential given the non conservation of photon number due to the double Compton and Bremsstrahlung processes. In these conditions, energy injection mechanisms simply heat up the primordial plasma increasing the background temperature (Cf. figure 3.9). At lower redshifts however, these processes are no longer efficient and energy releasing mechanisms create spectral distortions. We can

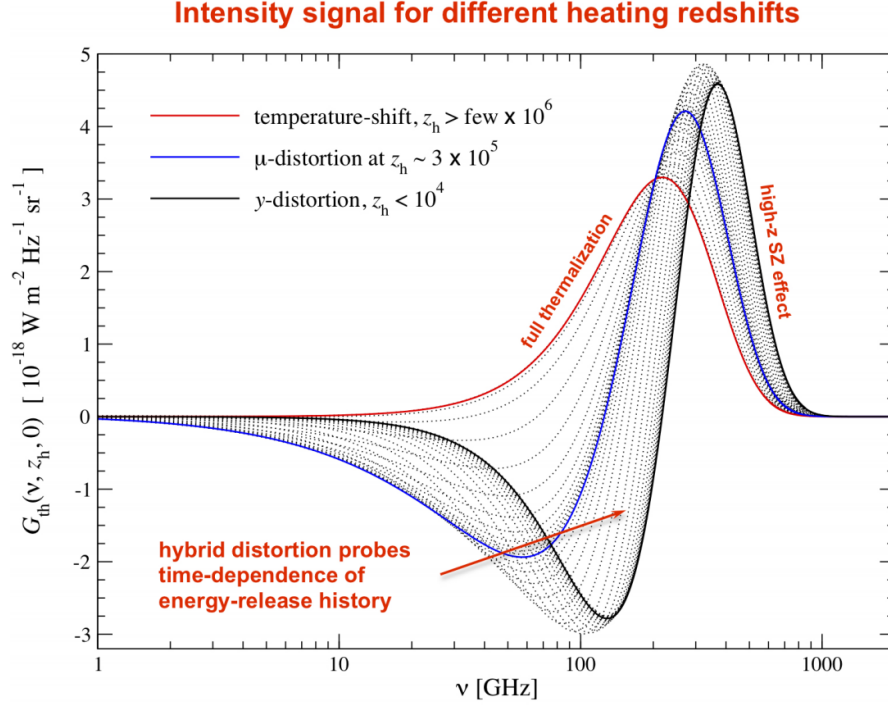


Figure 3.9 – Spectral distortions of the CMB after energy injection at different redshifts z_h . *Credit: Jens Chluba [4].*

distinguish two main types of distortions, the so-called μ -type and y -type distortions. We briefly explain each of these two effects below (see figure 3.9 for an illustration). Note that additional distortions not mentioned here can provide extra information, for instance the r -type distortion which contains the time-dependent information from the gradual transition between the μ -type and y -type distortions. For more detailed reviews of the topic, one may refer to [142, 318, 443].

μ -type distortion: For $10^5 < z \lesssim \text{few} \times 10^6$, the double Compton and free-free processes become inefficient while the Compton scattering process is still in equilibrium. This means that photons still thermalize but no process can ensure a null chemical potential. In particular, when energy injection happens the temperature of the photon bath keeps track of the temperature of the electrons, however given that the number of photons has to be conserved they cannot relax to a pure black-body spectrum and thus acquire a non-zero chemical potential. This distortion is usually described in terms of the dimensionless parameter

$$\tilde{\mu} \equiv \frac{\mu}{T_\gamma}. \quad (3.87)$$

y -type distortion: For $10^3 \lesssim z < 10^4$, Compton scattering ceases to be efficient and photons can no longer be thermalized. The result is a spectral distortion similar to what happens in the Sunyaev-Zel'dovich effect, with the exception that this effect is not localized and can be observed in all directions in the sky. This distortion is characterized by the Compton- y parameter which reads

$$y = \int \sigma_T n_e \frac{T_e - T_\gamma}{m_e} ds, \quad (3.88)$$

where the integral is over the photon path. The sign of y determines the flow of energy. For $y > 0$ energy is transferred from electrons to photons (*Comptonization*), while energy flows in the opposite way for $y < 0$ (*Compton cooling*). Note that in the definition of the Compton parameter in the Sunyaev Zel'dovich section we did not introduce T_γ : this is because typically the hot gas of electrons in clusters has $T_e \sim \text{few} \times 10^7$ K, and therefore we have $T_e \gg T_\gamma$.

Constraining fundamental physics

CMB spectral distortions open a unique window on the Universe, allowing us to probe physical processes happening in a wide range of redshifts. In this way we can pursue a broad set of science targets. Below we give a very short overview. More details can be found in [141].

Dark matter: Spectral distortions can offer complementary information to tighten the constraints on dark matter annihilation and decay. Spectral distortions can provide a direct measurement of the lifetimes, t_X , of decaying particles, for $t_X \simeq 10^6 - 10^{12}$ s. Moreover, we can also constrain annihilating particles using μ -type distortions which are sensitive to particles with $m \lesssim 100$ keV.

Axion-like particles: Axion-like particles (ALPs) couple with photons in the presence of an external magnetic field. These interactions distort the photon energy spectrum and can thus constitute a robust way of detecting these particles given that the CMB spectrum is well characterized, and that CMB photons would couple to these particles as they go through the galactic and extra-galactic magnetic fields. This effect would enable the exploration of a parameter space of ALP masses and coupling strengths which is not currently accessible to particle physics experiments.

Primordial black holes: Spectral distortions can put tight limits on the abundance of primordial black holes (PBHs) by exploiting mainly two effects. The first one stems from the dissipation of large density perturbations that collapse into PBHs. The second one is due to the Hawking radiation emitted by evaporating PBHs.

Inflation: Spectral distortions are sensitive to the properties of the primordial density perturbations, and can therefore provide key constraints on inflation or other competing models. The dissipation of these perturbations through Silk damping produce observable distortions in the CMB spectrum. In particular, the measurement of the μ -type distortion can provide constraints on the amplitude and scale dependence of the power spectrum around $k \simeq 10^3 \text{ Mpc}^{-1}$, complementing measurements from CMB anisotropies and large-scale structure which are only limited to the largest scales, $k \simeq 10^{-4} - 0.1 \text{ Mpc}^{-1}$. The wide range of scales can then be suitably exploited to derive bounds not only on the tilt of the spectrum, but also on the running of the tilt and the running of the running. This would be very useful to constrain slow-roll models of inflation. In addition, spectral distortions also depend on the type of perturbations, i.e. whether adiabatic or iso-curvature, and can be sourced by tensor perturbations and primordial non-Gaussianity, thereby providing numerous ways by which one can

test the different inflation scenarios.

Observational constraints

Much of the work is still left to be done in this area. The state-of-the-art upper limits were derived roughly some 25 years ago by the COBE/FIRAS instrument [199]

$$\begin{aligned} |y| &< 1.9 \times 10^{-5}, \\ |\tilde{\mu}| &< 9 \times 10^{-5}. \end{aligned} \quad (3.89)$$

The experimental efforts that followed have been focusing solely on the measurement of the temperature and polarization anisotropies. However, a number of experimental concepts aiming at the measurement of spectral distortions have recently been proposed such as PIXIE [290], PRISM [61], and PRISTINE [8], in addition spectral distortions are also considered among the priorities for a future large-scale mission of ESA in the context of its long term science program, Voyage 2050 [10].

3.5.3 Sunyaev-Zel'dovich Clusters

As we have mentioned earlier in this chapter, the Sunyaev-Zel'dovich effects result from the reverse Compton scattering of CMB photons off a distribution of high energy electrons in the intra-cluster medium (ICM) of galaxy clusters. The thermal SZ effect induces an increased intensity of the CMB at frequencies higher than ~ 220 GHz and a decrease at lower frequencies as can be seen in figure 3.10. The first detection of the thermal SZ effect was performed by the Owens Valley Radio Observatory (OVRO) in 1984 [98]. It took several years later, until 2012, for the kinematic SZ effect to be detected [228, 341]. Today, state of the art CMB experiments have provided catalogs with thousands of SZ clusters [32, 100, 244] which constitute a valuable resource to conduct cosmological analyses as well as for the study of the formation and evolution of structures.

Traditionally, our understanding of the ICM has relied mostly on X-ray observations. The X-ray surface brightness depends on both the gas temperature and density

$$S_X = \frac{1}{4\pi(1+z)^3} \int n_e^2(l) \Lambda_{ee}(T_e(l), Z) dl, \quad (3.90)$$

where $n_e(l)$ and $T_e(l)$ are respectively the electron density and temperature along the line-of-sight l , and Λ_{ee} is the X-ray emissivity, z the cluster's redshift, and Z is the metallicity. The surface brightness above is expressed in $\text{cts arcmin}^{-2} \text{ s}^{-1}$. When converted to units of intensity an additional factor $(1+z)^{-1}$ needs to be introduced. As such, the X-ray intensity suffers from the so-called cosmological dimming effect with a scaling $\propto (1+z)^{-4}$. The observations of SZ effects can complement X-ray observations given that it does not suffer from cosmological dimming, the tSZ effect has a different line-of-sight dependence and thus allow one to infer complementary information on the thermodynamic properties of the ICM, and through the kSZ effect we can obtain information on the kinematics of the ICM. With high-angular resolution imaging experiments, such as NIKA [29], MUSTANG [294] and ALMA [284], we are able to spatially resolve SZ clusters,

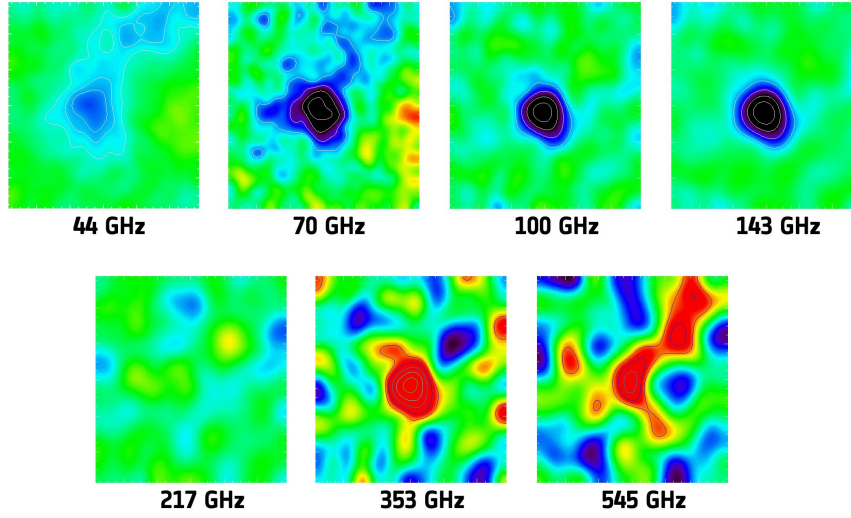


Figure 3.10 – Maps of the Abell 2319 cluster in seven frequencies obtained by the Planck satellite. We see the local spectral distortion of the CMB caused by the SZ effect which results in a reduced intensity at low frequencies and an increased intensity at high frequencies. The transition happens at 217 GHz where the corresponding map shows no signal. *Credit:* Planck Collaboration.

and use complementary probes, such as shown in figure 3.11 to study the thermodynamic structure of the ICM and the distribution of matter and dark matter. Conducting these measurements in a wide range of redshifts helps improve our understanding of the physical processes that govern galaxy and cluster formation.

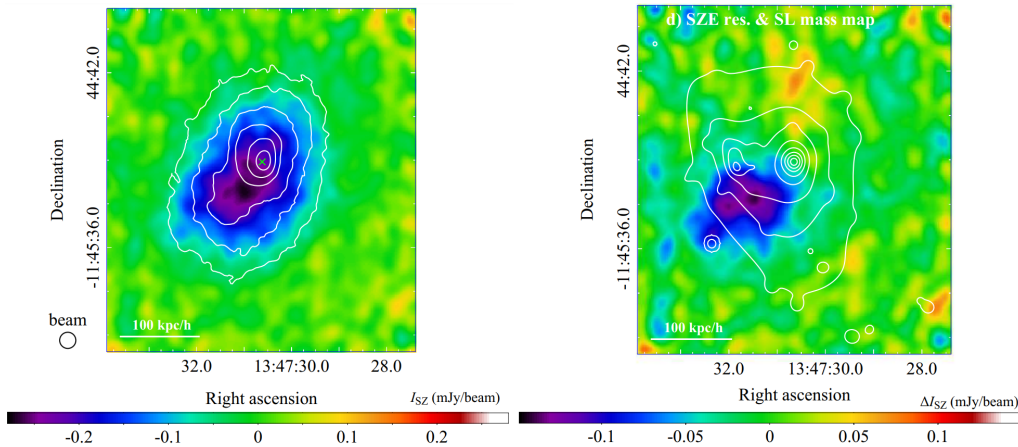


Figure 3.11 – *Left:* SZE surface brightness of RXJ1347.5-1145 by ALMA overlaid with X-ray contours from *Chandra*. *Right:* Residual tSZ signal revealed after subtracting the mean profile. Lensing contours from HST are overlaid, revealing the location of dark matter. *Credit:* S. Ueda et al. [457].

In cosmology, measuring cluster abundance as a function of redshift allows to set constraints on matter density Ω_m , the growth of structure $\sigma_8(z)$, as well as neutrino mass and dark energy equation of state. In addition, statistical analyses of the tSZ and kSZ signals are also emerging as powerful cosmological tools and

would provide further constraints to test cosmological models in the future. Upcoming CMB experiments such as the Simons Observatory [36] or future ones like CMB-S4 [16] have dedicated programs to fully exploit the potential of SZ effects.

3.5.4 Cosmic birefringence

Cosmic birefringence refers to the rotation of the CMB linear polarization induced by parity violating processes, converting E into B-modes and vice-versa, and producing non-vanishing TB and EB correlations in the CMB maps. We can trace back this effect to two main mechanisms, the *Chern-Simons mechanism* and *Faraday rotation*.

The Chern-Simons mechanism

This effect [127] is the result of the coupling of CMB photons with pseudoscalar, axion-like fields $\phi(t)$ associated with models of dark matter and dark energy [194, 326]. The rotation of the CMB polarization angle in a given direction \mathbf{n} can be written as [381]

$$\alpha(\mathbf{n}) = \frac{1}{2} g_{\phi\gamma} \Delta\phi(\mathbf{n}), \quad (3.91)$$

where $g_{\phi\gamma}$ is the Chern-Simons coupling coefficient and $\Delta\phi$ is the change of $\phi(t)$ along the photons trajectory between the time of decoupling and the time of observation. In the presence of anisotropies, the power spectrum of α , for a massless scalar field and in the large-scale limit ($\ell \lesssim 100$), is given by [121]

$$\sqrt{\frac{\ell(\ell+1)C_\ell^{\alpha\alpha}}{2\pi}} = \frac{H_I g_{\phi\gamma}}{4\pi}, \quad (3.92)$$

where H_I is the Hubble parameter during inflation.

Faraday rotation

In the presence of a magnetic field, this effect [296, 231] induces a polarization rotation that is proportional to the projection of the magnetic field parallel to the direction of propagation. Primordial Magnetic Fields (PMFs), introduced to explain the puzzling origins of magnetic fields in galaxies and galaxy clusters, could thus leave a signature in the CMB polarization through this effect. We can write the power spectrum of the rotation angle, in the large scale limit ($\ell \lesssim 100$) and for nearly scale-invariant PMFs, as the following [38]

$$\sqrt{\frac{\ell(\ell+1)C_\ell^{\alpha\alpha}}{2\pi}} = 1.9 \times 10^{-4} \left(\frac{\nu}{150 \text{ GHz}} \right)^{-2} \left(\frac{B_{1 \text{ Mpc}}}{1 \text{ nG}} \right), \quad (3.93)$$

where ν is the observing frequency and $B_{1 \text{ Mpc}}$ is the magnetic field averaged over 1 Mpc.

Observational constraints

The state of the art upper limits on the primordial magnetic fields are set by the POLARBEAR experiment using measurements of the B-mode power spectrum alone [37]

$$B_{1 \text{ Mpc}} < 3.9 \text{ nG}, \quad (3.94)$$

this result is compatible with Planck 2015 limit [31] using both temperature and polarization data, $B_{1\text{ Mpc}} < 4.4\text{ nG}$, both results are given at the 95% confidence level. The Atacama Cosmology Telescope (ACT) has also given a 95% CL upper limit on the Chern-Simons coupling [348]

$$g_{\phi\gamma} \leq \frac{4.0 \times 10^{-2}}{H_I}. \quad (3.95)$$

It is worth mentioning that such measurements are limited by the systematic uncertainty in the calibration of the polarization angles of the instrument. A recently proposed method [335] where the polarization miscalibration and the birefringence angles are estimated simultaneously, allows to mitigate these systematic errors. Its application on Planck 2018 polarization data yields a measurement of a cosmic birefringence angle [334]

$$\alpha = 0.35 \pm 0.14 \text{ deg} \quad (68\% \text{ CL}), \quad (3.96)$$

thereby excluding $\alpha = 0$ with a statistical significance of 2.4σ .

After reviewing many aspects of the richness of the CMB as a cosmological probe, in addition to all the complementary information it can provide on other areas of physics such as neutrino physics and extensions of the standard model of particle physics. It is worth to note that to consolidate our understanding of the physics, combining the CMB results with other cosmological probes can be a powerful way, not only of verifying the consistency of the results, but also to set complementary constraints and break degeneracies enabling high precision measurements or very tight bounds on the cosmological parameters of interest. In the presence of inconsistencies between the data, it can also reveal the presence of unknown systematics or errors in the analyses and thereby ultimately lead to the improvement of the robustness of our results, or alternatively signal new physics. In the next chapter, we will review a broad range of observables in cosmology and what we can learn from them.

Cosmology from multiple probes

Contents

4.1	Galaxy surveys	84
4.1.1	Baryon Acoustic Oscillations	85
4.1.2	Redshift-space distortions	86
4.1.3	Weak lensing and cosmic shear	88
4.2	Lyman- α forest	89
4.3	21-cm cosmology	90
4.4	Type-Ia supernovae	92
4.4.1	Hubble diagram	92
4.4.2	Cosmic distance ladder and the H_0 tension	94
4.4.3	Future cosmological targets	94
4.5	Gravitational waves	95

Significant progress has been made in our understanding of the physics governing our Universe. Observations of the Cosmic Microwave Background constitute one of the pillars that helped achieve this level of understanding as emphasized in the last chapter. However it is not the only tool available to cosmologists. The success of modern cosmology relies on a number of equally important observables which are sensitive in different ways to the properties of the Universe and therefore offer unique insights helping us build a comprehensive global picture of the physics at play at different stages of the Universe’s evolution. In addition, the combination of independent information from multiple probes allow us to break degeneracies between parameters enabling highly precise measurements or constraints on cosmological parameters.

The main aim of this chapter is to offer a sense of how the synergies between multiple probes in cosmology are exploited, and how they help us to build a robust understanding of the physics of our Universe. The list of examples treated is not intended to be exhaustive. We will avoid technical details to focus on the global picture.

4.1 Galaxy surveys

One of the most prominent ways of constraining cosmological models and testing their consistency is to probe the large scale distribution of matter in the Universe. Since matter exists predominantly in the form of dark matter, this distribution cannot be observed directly. We resort instead to mapping the distribution of galaxies, which are biased tracers of the underlying matter density field [269, 74].

The 3D distribution of galaxies can be mapped with photometric and spectroscopic sky surveys. Spectroscopic data enables the measurement of redshift, while photometric data allows to determine the angular positions of galaxies on the sky. Although measurements of redshifts can also be performed through photometry, the corresponding errors remain significantly higher than those of spectroscopic measurements. Figure 4.1 shows an example of a slice of the 3D map of galaxies constructed by the Sloan Digital Sky Survey (SDSS) [99], the distribution is represented in the so-called *wedge diagram* where the Milky Way is located at the center, the radial axis represents redshift, the polar angle specifies the right ascension, and colors encode the local density (denser regions in red).

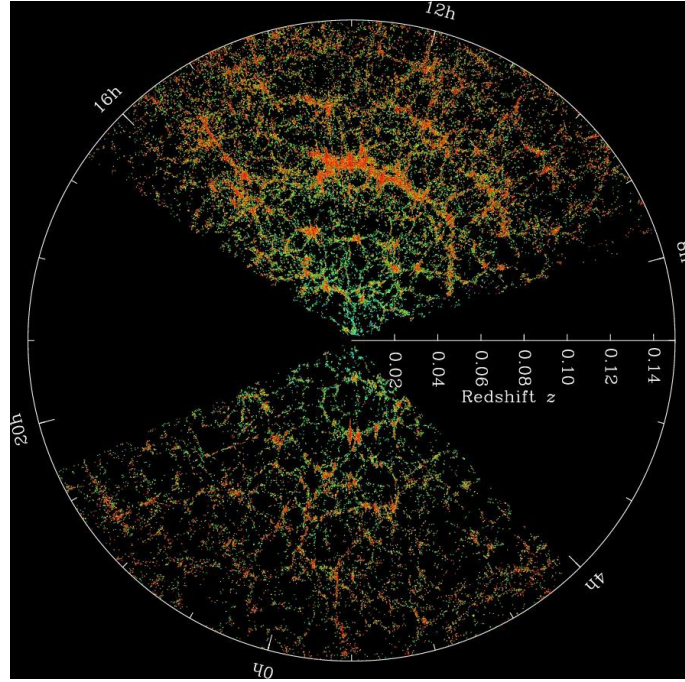


Figure 4.1 – 2D slice in the 3D distribution map of galaxies measured by the Sloan Digital Sky Survey (SDSS). Each point represents a galaxy. *Credit:* M. Blanton and the Sloan Digital Sky Survey.

On large scales, a simple ansatz for relating the galaxy distribution to the matter distribution reads

$$\delta_g(z, \mathbf{n}) \equiv \frac{n_g(z, \mathbf{n})}{\bar{n}_g(z)} - 1 = b(z)\delta(z, \mathbf{n}), \quad (4.1)$$

where δ_g and δ are respectively the galaxy and matter density contrast, $n_g(z, \mathbf{n})$ and \bar{n}_g are the local and mean number density of galaxies, and $b(z)$ is the bias

parameter assumed to be scale independent at large scales and is a function of redshift z [163]. This parameter encodes the complicated physics of galaxy formation. The galaxy power spectrum, P_{δ_g} , can then be written as a function of the matter power spectrum, P_δ , as

$$P_{\delta_g}(z, k) = b^2(z) P_\delta(z, k). \quad (4.2)$$

4.1.1 Baryon Acoustic Oscillations

As we have discussed in the previous chapter, the primordial plasma contains acoustic oscillations which manifest themselves in the peaks we see in the CMB power spectra. At the drag epoch, i.e., when baryons decouple from the photons, the baryon pressure drops and the sound waves are frozen. The characteristic scale of these acoustic waves is then given by the sound horizon at the drag epoch. Its corresponding comoving scale is roughly, $r_s \simeq 150$ Mpc, computed from the matter and baryonic matter densities, $\Omega_m h^2$ and $\Omega_b h^2$, as well the radiation content of the Universe. These parameters can be precisely measured by the CMB. Given that these sound waves should also leave their imprint on the clustering of matter, using galaxies as a tracer, we expect to see a peak in the correlation function of galaxies, located at the sound horizon scale r_s (also called BAO scale). This peak has been first detected by SDSS in 2005 [179], figure 4.2 shows one of the latest measurements of the BAO feature using the Luminous Red Galaxies (LRGs) sample of eBOSS [86] which is a subprogram of SDSS.

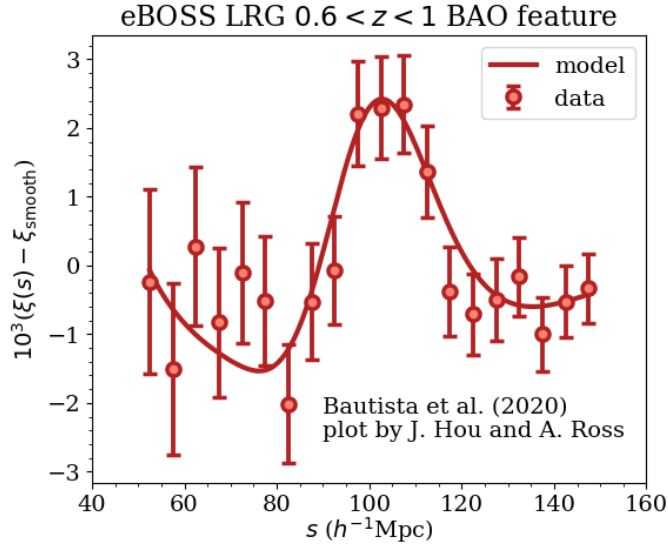


Figure 4.2 – The BAO peak in the galaxy correlation function measured from the LRG sample of eBOSS for redshifts $0.6 < z < 1$, representing an effective volume of 2.7 Gpc^3 . Credit: Bautista et al. [86].

This measurement opens up the possibility to use the sound horizon as a standard ruler. Since its absolute calibration can be performed using CMB observations, we can use its apparent size at different redshifts, in order to infer measures of distance as a function of redshift. We can use the BAO scale both in the direction of the line-of-sight and in the transverse direction. In the line-of-sight,

measuring the redshift subtended by the BAO feature, Δz , allows a direct measurement of the Hubble parameter, $H(z) = \Delta z/r_s$. In the transverse direction, determining angular size of the BAO feature, $\Delta\theta$, allows to estimate the (comoving) angular diameter distance given by $d_A(z) = r_s/\Delta\theta$. These measurements can then be used to derive constraints on dark energy, curvature, H_0 , as well as the matter density, and the total mass of neutrinos.

Combining BAO with the CMB is a powerful way to constrain the parameters that control measures of distance. In fact, by measuring distance at different redshifts we are measuring different combinations between these parameters allowing us to break degeneracies between them. One such example is given in the left panel of figure 4.3 for curvature, where CMB data alone favors a closed Universe but their combination with BAO favors a flat Universe with $\Omega_k = 0.0001 \pm 0.0018$ (68% CL) [53]. The right panel also shows the complementarity between the two data sets in constraining a flat w CDM model, the BAO data alone do not exhibit strong degeneracy between matter density and the dark energy equation of state parameter, w , and its combination with CMB data is able to break the degeneracies between the two parameters in the CMB data set yielding an equation of state for dark energy, $w = -1.034^{+0.061}_{-0.053}$ (68% CL), consistent with a cosmological constant [53]. The figure also shows constraints from type Ia supernovae, which is another powerful cosmological probe discussed later in this chapter.

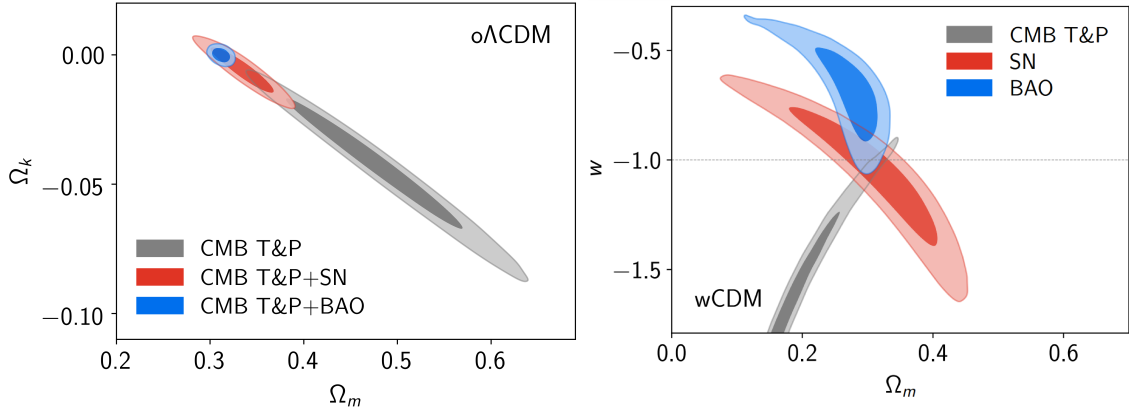


Figure 4.3 – *Left*: Constraints in the $\Omega_m - \Omega_k$ plane, on a model with cosmological constant and free curvature ($o\Lambda$ CDM), from the combination of CMB temperature and polarization data alone (gray), CMB + SNe Ia data (red), and CMB + BAO (blue). *Right*: Constraints in the $\Omega_m - w$ plane, on a flat w CDM model, from CMB data alone (gray), SNe Ia data alone (red), and BAO data only (blue). *Credit*: SDSS [53].

4.1.2 Redshift-space distortions

One important point that we have not considered in our previous discussion, is that the observed redshifts of galaxies are not only the result of the Hubble flow, but also of their peculiar motion. The latter induces a distortion in the positions of galaxies in the wedge diagrams along the radial direction which corresponds to redshift. As a result, the galaxy correlation function is not isotropic in redshift space [270]. These so-called *redshift-space distortions* (RSD) manifest themselves

in two ways. On small scales, the high velocity dispersion of galaxies within a cluster for example, produces an elongated structure pointing towards the observer in redshift space, this particular pattern is referred to as “Fingers of God”. On larger scales, as galaxies fall towards the center of an overdensity, we see galaxies closer to us than the center of the overdensity as moving faster than the Hubble flow and hence with a larger apparent redshift, while galaxies behind the center are moving slower than the Hubble flow and thus their apparent redshift is smaller. As a consequence, in redshift space the distribution appears to be squashed along the line-of-sight. Figure 4.4 shows an example of the anisotropic two-point correlation function as a function of transverse and radial separation, inferred from 2dF Galaxy Redshift Survey (2dFGRS) data [239]. We can clearly see the squashing of the distribution for large separations ($\gtrsim 10h^{-1}$ Mpc), while the Fingers of God are visible on smaller separations.

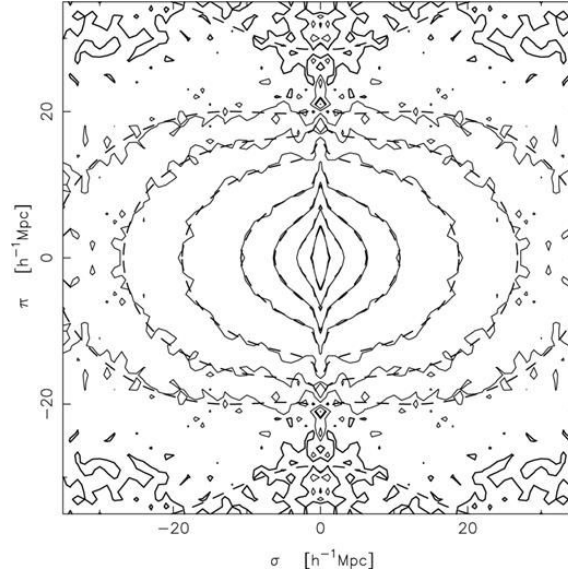


Figure 4.4 – The two-point correlation function ξ_g as a function of transverse separation $\sigma = cz\Delta\theta/H_0$ and radial separation $\pi = c\Delta z/H_0$. Solid lines represent contours of constant ξ_g . The dashed curves represent the best-fitting model.

Credit: Hawkins et al. [239].

The RSD effect is sensitive to the rate at which structures grow. We can define the growth rate $f(z)$ as

$$f = \frac{\partial \ln \sigma_8}{\partial \ln a}, \quad (4.3)$$

where $\sigma_8(z)$ characterizes the amplitude of linear matter perturbations on a co-moving scale of $8h^{-1}$ Mpc. With RSD measurements we can constrain $f\sigma_8$ improving our understanding of structure formation and dynamics. The left panel of figure 4.5 shows constraints on $\Omega_m - \sigma_8$ assuming Λ CDM cosmology, the combination of RSD measurements from SDSS with other probes such as CMB lensing from Planck, or weak lensing measurements (covered in the next subsection) from DES, allows to break degeneracies and yields tighter constraints which are overlapping with the contours derived from CMB temperature and polarization data alone, thus confirming the overall consistency of the results [53]. RSD measurements can also be used to test theories of gravity [23]. The right panel of figure

4.5 shows constraints on two parameters μ_0 and Σ_0 , characterizing deviations of the metric perturbations, Ψ and Φ , from predictions of general relativity (GR). Given that RSD and lensing measurements have complementary sensitivities to these parameters, their combination can provide stringent tests of gravity. As the figure shows the current results are consistent with GR prediction $\mu_0 = \Sigma_0 = 0$ [53].

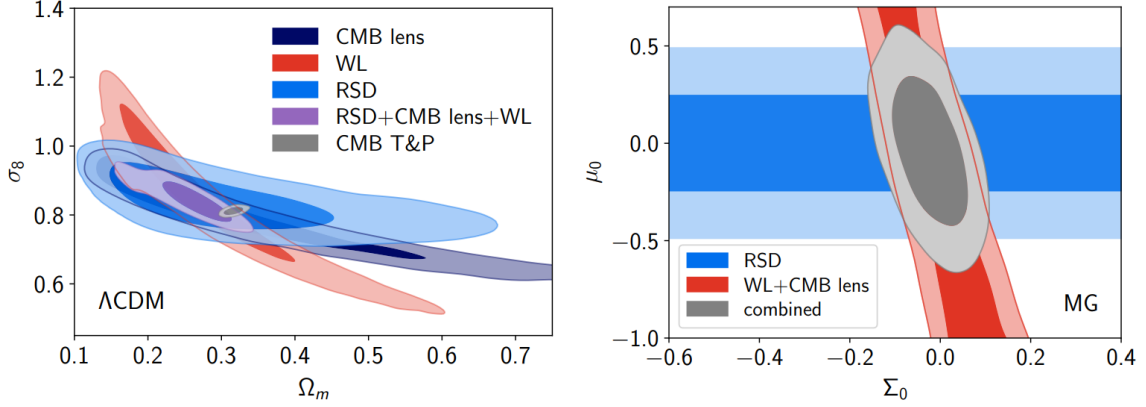


Figure 4.5 – *Left:* $\Omega_m - \sigma_8$ 68% and 95% constraints. Blue contours represent constraints from RSD data alone, red contours correspond to weak lensing from DES, dark blue to CMB weak lensing from Planck, light purple to the combination of RSD, and lensing from both DES and Planck, while gray contours correspond to Planck temperature and polarization data. *Right:* Constraints (68% and 95%) on $\Sigma_0 - \mu_0$ from RSD data alone (blue), weak lensing from DES and CMB lensing from Planck (red), and their combination (gray).
Credit: SDSS [53].

4.1.3 Weak lensing and cosmic shear

As explained in detail in the previous chapter, the presence of massive structures in the path of photons distorts the shapes and sizes of images of the sources. These effects are of course not exclusive to CMB photons, but also apply to background galaxies and clusters in galaxy surveys. Typically, these effects happen in the weak lensing regime, where distortions are of the order of 1%. Given that the intrinsic shapes and sizes of galaxies are not known, weak lensing can only be measured statistically through the correlations of shapes and sizes between pairs of galaxies. Photons coming from neighboring galaxies should propagate through the same gravitational field, the resulting distortions are thus expected to be correlated. Figure 4.6 shows an illustration of the correlation between the ellipticities of distant sources and their relation with the projected matter distribution. The shape distortions are referred to as *cosmic shear*, they are much more used in practice because we can reasonably assume that the statistical average shape of galaxies should be spherical, given that there should be no preferred direction for the elongation axis of galaxies, and that there is less scatter in galaxy shapes than in their sizes [467]. Weak lensing measurements are a powerful probe to constrain the distribution of matter in a way that is independent of galaxy bias, although there are many other biases (e.g., intrinsic alignment [456], multiplica-

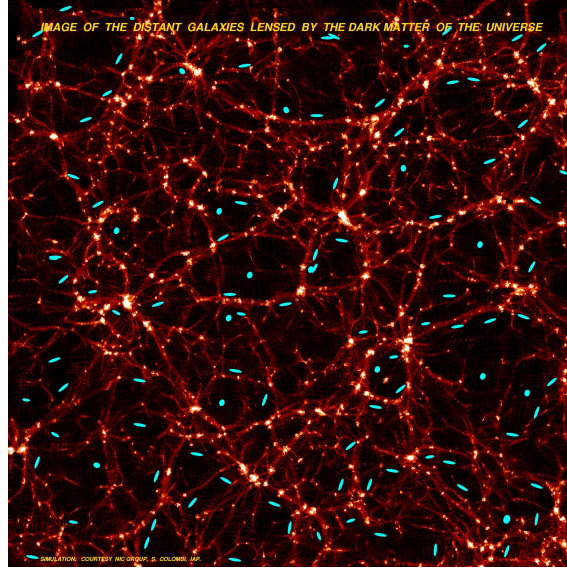


Figure 4.6 – Simulated images of distant galaxies, shown in light blue, lensed by the dark matter shown in the red and white filaments. Weak gravitational lensing stretches galaxy images along directions which are parallel to the dark matter filaments. *Credit:* S. Colombi, IAP / Canada-France-Hawaii Telescope (CFHT) Corporation.

tive bias [316], etc). These constraints can be inferred from the computation of the shear power spectrum, related to line-of-sight integrals over the matter power spectrum. In the linear regime at low redshift, its amplitude scales as $\propto \Omega_m^2 \sigma_8^2$. These scalings are altered in the non-linear regime and the measurements become sensitive to different combinations of Ω_m and σ_8 . As we have seen in figure 4.5, weak lensing measurements are complementary to other cosmological probes and can be used to break degeneracies between cosmological parameters and yield strong constraints or verify consistency testing cosmological models or theories of gravity.

4.2 Lyman- α forest

Another useful tool in cosmology comes from the spectral analysis of the light emitted by quasi-stellar objects (QSOs) or *quasars*. Quasars are some of the most luminous objects in the entire Universe, with luminosities orders of magnitude higher than that of their own host galaxies. These are active galactic nuclei (AGNs) whose radiation is powered by matter accretion onto a central supermassive black hole. They are characterized by a broad energy distribution, emitting continuum radiation from γ -rays and X-rays to the far infrared wavelengths. In the context of cosmology, the very high luminosity of these objects can be exploited to use them as background light tracing the distribution of matter at very high redshifts, by studying absorption lines in their measured spectra.

Let us consider the effect of the presence of neutral hydrogen in the intergalactic medium (IGM) as the background light from quasars propagate through it. At discrete wavelengths, the hydrogen atom absorbs a photon and gets excited, with principal quantum number $n \geq 2$, before it relaxes back to the ground state, $n = 1$.

The corresponding electronic transitions from excited states to the ground state are known as the *Lyman series*. The first transition from $n = 2$ to $n = 1$ is known as the Lyman- α line, and corresponds to an absorption or emission wavelength, $\lambda_{Ly\alpha} = 1215.67 \text{ \AA}$ in the absorber or emitter rest frame. When excited hydrogen atoms relax to the ground state they do emit photons in all directions, the incident light is thus scattered away and therefore shows up as missing energy at the corresponding wavelength when examining the quasar spectrum. However as the background quasar signal gets redshifted while propagating through the IGM, the neutral hydrogen clouds should leave a series of absorption lines at wavelengths shorter than the Lyman- α emission peak. These absorption features are known as the Lyman- α forest. Figure 4.7 shows an illustration of this process as light from a distant quasar travels through an expanding universe, and the neutral hydrogen clouds imprint absorption lines in its spectrum.

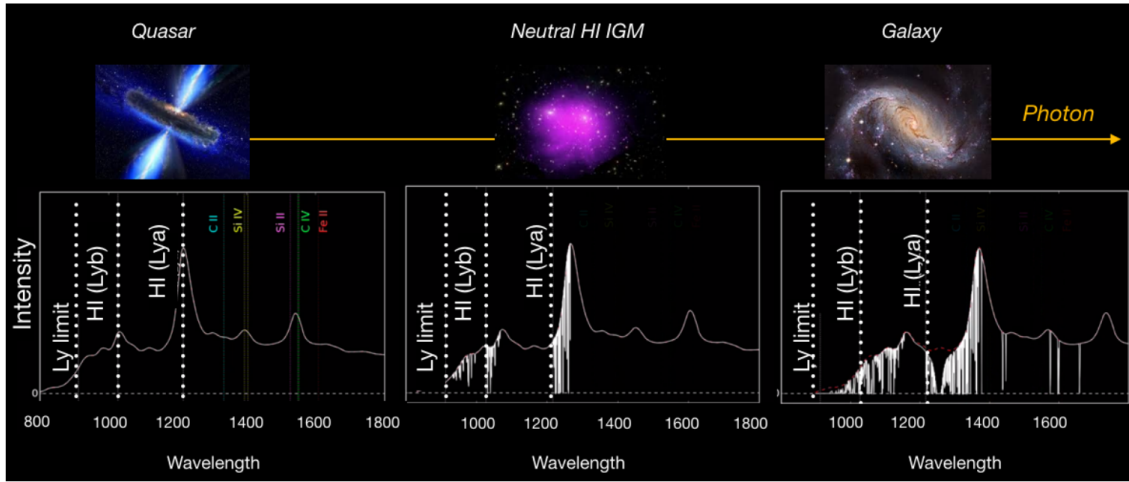


Figure 4.7 – Illustration of the evolution of the spectrum of light emitted from a distant quasar (*left*) as it travels through the IGM (*middle*) and near a galaxy (*left*) towards our telescopes. When travelling near a galaxy, saturated absorption lines are produced which are distinct from the low-density absorptions in the IGM, in addition metal absorption lines also appear at larger wavelengths than the Ly α emission line. *Credit:* Solène Chabanier [132] / Andrew Pontzen [6].

The first detection of the Lyman- α forest was in 1971 [319]. As explained above it is sensitive to the column density distribution of neutral hydrogen in the IGM. Given that, at high redshifts, the absorbing gas density in the IGM is relatively low ($\delta \lesssim 10$), the pressure forces are small and its dynamics on large scales are driven by the gravitational influence of dark matter perturbations [97]. The two-point correlation function of Lyman- α lines is thus a tracer of the matter distribution and can then be used to constrain cosmologies [466, 169].

4.3 21-cm cosmology

One of the most promising cosmological probes for the future is the 21 cm intensity mapping [385]. The 21 cm emission line of neutral hydrogen arises from the hyperfine splitting of its ground state induced by the interaction between the magnetic moments of the proton and the electron. The transition between the two

hyperfine states corresponds to a wavelength of 21.1 cm or a frequency of 1420 MHz. Since its first detection in 1951 [189], it has been used as an astrophysical probe to map the velocity distribution of gas in the local Universe and is therefore a useful tracer of galactic dynamics. However, modern radio telescopes are aiming at using the 21 cm line in cosmological contexts to probe the gas along the line-of-sight up to very high redshifts inaccessible by other probes. In particular, it will be a powerful way to peer through the epoch of reionization and perform a precise mapping of the ionized structures, as is targeted by future experiments such as SKA [9]. In addition, the 21 cm intensity mapping surveys aim to measure the integrated signal from the large scale distribution of galaxies without resolving the light from individual galaxies, this will be used as a tracer of the matter distribution, measuring the associated BAO and RSD signals, which will enable to constrain the dark energy equation of state, and modified gravity models with unprecedented statistical power given the large comoving volume that is possible to probe with this technique. Furthermore, it will enable constraining non-Gaussianity, statistical tests of homogeneity and isotropy, and its combination with other probes will also allow to improve the uncertainties on many other cosmological parameters [322].

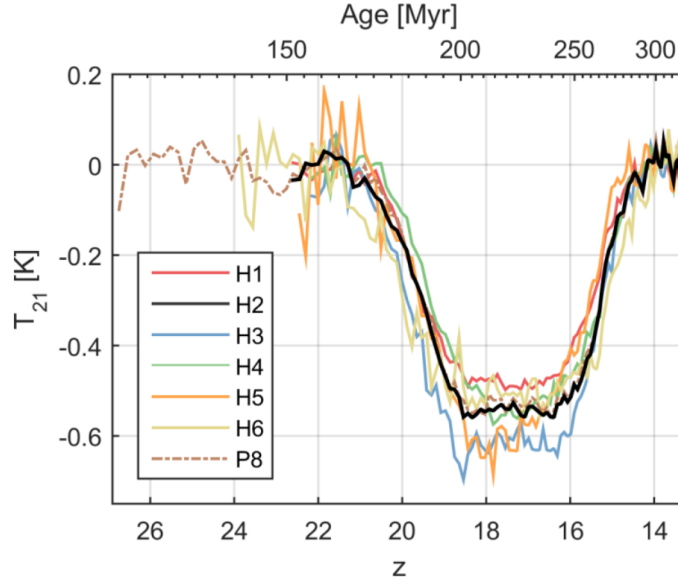


Figure 4.8 – Best-fitting absorption profiles detected by EDGES for different hardware configurations and data analysis pipelines. *Credit:* J. D. Bowman et al. [104].

One of the recent results that is worth mentioning, is the one reported by the Experiment to Detect Global Signature (EDGES) [338]. The experiment performs low angular resolution measurements of the sky averaged global 21 cm signal, from a radio quiet area located at the Murchison Radio-astronomy Observatory (MRO) in western Australia. Figure 4.8 shows an absorption signal detected in 2018, at 78 MHz between redshifts $15 \leq z \leq 20$ [104]. The observable 21 cm signal is given as a differential brightness temperature relative to the background

radiation

$$T_{21}(z) \approx 0.023 x_{\text{HI}}(z) \sqrt{\left(\frac{0.15}{\Omega_m}\right) \left(\frac{1+z}{10}\right) \left(\frac{\Omega_b h}{0.02}\right)} \left[1 - \frac{T_R(z)}{T_S(z)}\right] \text{ K}, \quad (4.4)$$

where $x_{\text{HI}}(z)$ is the neutral hydrogen fraction at redshift z , Ω_m and Ω_b are the matter and baryon density parameters, T_R is the background radiation temperature, and T_S is the 21cm spin temperature which defines the relative population of the hyperfine states. The absorption feature is theoretically expected. At early times neutral hydrogen and CMB photons should be in thermal equilibrium ($T_S = T_R$), and no 21 cm signal should be detectable. As the first sources of Lyman- α photons appear, a coupling between the spin temperature and the cool gas kinetic temperature appears through the so-called *Wouthuysen-Field effect* [474, 196, 197], causing an absorption signal as the equilibrium between neutral hydrogen and the CMB is broken ($T_S < T_R$). As the gas is heated first by stellar radiation and later by X-ray sources, typically accretion disks around stellar remnants, the absorption signal is expected to halt when the gas temperature reaches the background radiation temperature. However, the intriguing fact about this result is that the amplitude of the absorption profile is two times greater than the largest predictions, indicating either a lower gas temperature or a higher background radiation temperature. One of the exciting proposed explanations is that the cooling of the gas could result from interactions between dark matter and baryons [77]. If confirmed this would be the first detected non-gravitational signature of dark matter.

4.4 Type-Ia supernovae

In this section we will focus on a specific type of stellar explosions or supernovae, the so-called *type Ia supernovae* (SNe Ia) which is an important probe in cosmology. Type Ia supernovae occur when a white dwarf in a binary system accretes matter from its companion star until it crosses the limit of carbon nuclear fusion, runaway fusion reactions are then triggered releasing enough energy to rip the star apart in a violent explosion. The two interesting properties of these explosions in cosmological contexts is that they are very bright, and therefore can be seen at high redshifts, and their maximum luminosity is fairly consistent, therefore they constitute, in principle, good standard candles for measures of luminosity distance which are particularly useful for constraining cosmological models as we have seen before.

4.4.1 Hubble diagram

One of the challenges we are faced with when attempting to construct a Hubble diagram with type Ia supernovae, is that they are not truly perfect standard candles. Their peak luminosity in the B band show around 40% dispersion. However, several methods were devised to correct for this variability by exploiting specific correlations in the properties of SNe Ia. One of them is the peak luminosity-decline rate relationship discovered by Mark M. Phillips [370], namely the brighter the supernovae the slower its luminosity tends to decline

after reaching the maximum. This effect can be characterized either by the decline in magnitude 15 days after maximum in the B band, $\Delta m_{15}(B)$, introduced by Phillips [370], or by the linear stretch factor s introduced by Perlmutter's group [367] which simply stretches or contracts the supernovae light curves along the time axis. There exists another correlation between the brightness and the color index of the supernovae, namely that bluer supernovae tend to be brighter [397, 453]. This effect can be either intrinsic or the result of dust extinction. By correcting for these effects we are able to reduce the dispersion of the luminosity peak, and thus SNe Ia are “standardisable” candles. We should also note that there are other correlations such as with the host galaxy properties [432] or spectral properties [461] that need to be taken into account and corrected for. Summarizing all of the above, the standardized distance modulus $\mu \equiv 5 \log (d_L/10\text{pc})$ is written as

$$\mu = m_B^* - M_B + \alpha(s - 1) - \beta C, \quad (4.5)$$

where m_B^* is the measured peak magnitude in rest-frame B band, and C is the color index at maximum brightness. The three parameters (m_B^* , s , C) are fitted to the supernovae light curves using specific models [224, 146]. M_B is the absolute peak magnitude in the B band for “standard” supernovae ($s = C = 0$), along with α and β they constitute nuisance parameters that are fitted to the data simultaneously with the cosmological parameters. Figure 4.9 shows the Hubble diagram obtained using data from SDSS-II and SNLS [94]. These results can also be used in combination with other probes to set constraints on dark energy.

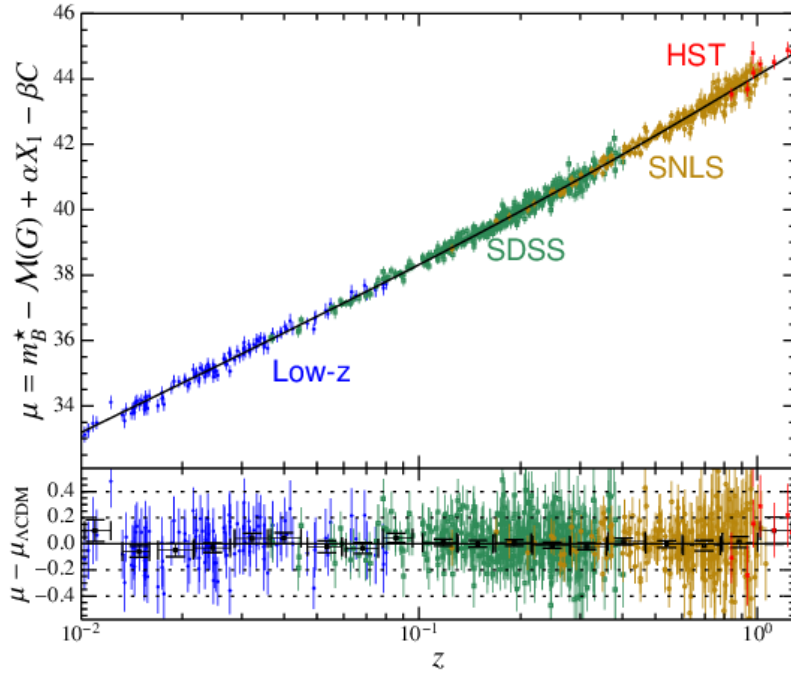


Figure 4.9 – Hubble diagram constructed with the “JLA” sample comprising 740 SNe Ia, the black line is the best-fit Λ CDM model for a fixed $H_0 = 70$ km/s/Mpc.

Credit: M. Betoule et al. [94].

4.4.2 Cosmic distance ladder and the H_0 tension

We can use SNe Ia to measure the local expansion rate H_0 . Such measurement, however, relies on building the so-called *cosmic distance ladder*, given that direct distance estimates of distant galaxies are not possible. The first rung of the distance ladder relies on the use of local geometric distance anchors to determine distances to hosts of Cepheids. These include, for instance, masers in NGC 4258 [260] or detached eclipsing binaries (DEB) in the Large Magellanic Cloud (LMC) [372]. Trigonometric parallaxes can also be used for Cepheids in the Milky Way [89, 394]. Cepheids are variable stars with a well known period-luminosity relation first determined by Henrietta Leavitt [307, 308]. This relation is calibrated using the distance measurements of the first rung of the ladder. The second rung of the ladder is then built using the previous relation to find distances of host galaxies which contain both Cepheids and SNe Ia. This is used to calibrate their intrinsic peak magnitude. As a consequence we can determine distances to SNe Ia at very large distances in the Hubble flow, where peculiar velocities are negligible, which constitutes the third rung in the distance ladder. The resulting distance-redshift relation is then used to infer a model independent value of H_0 . The recent value reported by SH0ES [395], $H_0 = 74.03 \pm 1.42 \text{ km s}^{-1} \text{ Mpc}^{-1}$, is in a 4.4σ tension with the one inferred from the best-fit Λ CDM from Planck [373]. With other independent measurements of H_0 such as the one with strongly lensed quasar time delays (H0LiCOW) [473] or from BAO data [53], the tension on H_0 is not only confirmed but becomes more significant. Figure 4.10 shows a compilation of the state-of-the-art emphasizing the discrepancies between early-time and late-time measurements of H_0 . This tension could point towards a lack of understanding of systematics in either types of measurements, or more interestingly, it could be the sign of new physics [392, 282, 344, 301, 357].

4.4.3 Future cosmological targets

With powerful future surveys such as the Vera C. Rubin Observatory [317], the sample size of SNe Ia will increase by hundreds of times. In addition to the improvement of the constraining power on the dark energy equation of state, this will enable performing statistical tests of isotropy and homogeneity by searching for angular dependence in the redshift-distance relation, thereby determining whether the dark energy equation of state is directionally dependent. Moreover, SNe Ia will also be used as tracers of the large scale structures to measure the BAO signal. Another target of interest is the measurement of the peculiar velocities of galaxies, constraining the growth rate of structure and testing deviations from GR [248]. Furthermore, it will be possible to measure the weak lensing effect on SNe Ia which causes either magnification or demagnification. Using SNe Ia brightness distribution cross-correlated with the large scale structure traced by foreground galaxies, the weak lensing signal can be measured and employed to constrain cosmological parameters as well as for consistency checks with other probes. Finally, time-delay cosmography with strong lens systems as is used for quasars can also be applied to SNe Ia and would be of interest particularly for measurements of H_0 .

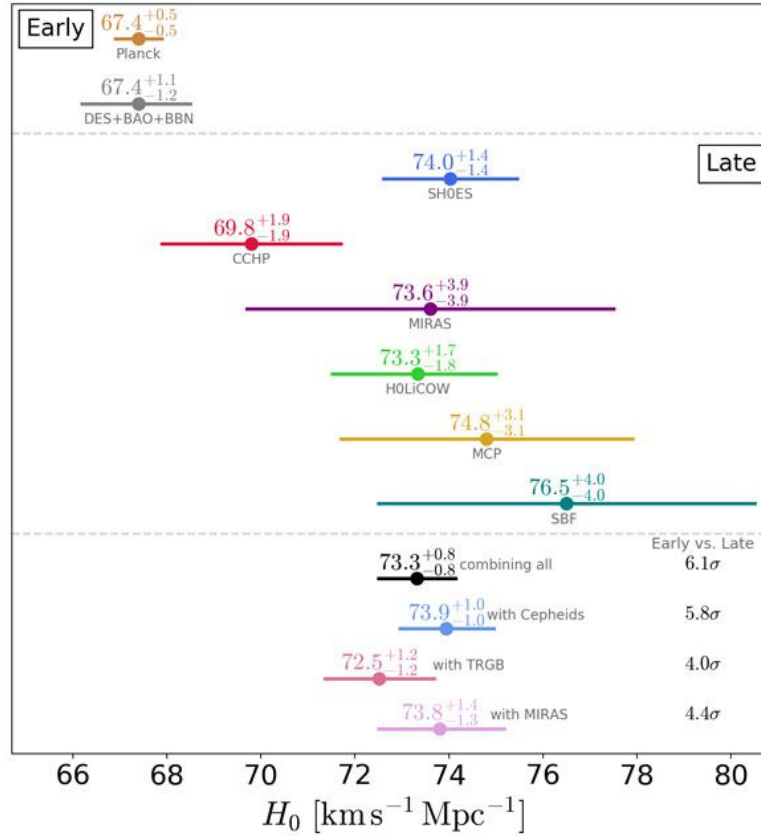


Figure 4.10 – Compilation of the different recent measurements of H_0 giving rise to a 4σ to 6σ tension between early-time and late-time measurements. *Credit:* Vivien Bonvin.

4.5 Gravitational waves

Gravitational waves are a promising cosmological probe for the future. They are already adding valuable contributions to the field, enabling stringent tests on general relativity and modified gravity to be conducted for instance [22, 19, 20]. First predicted by Einstein in 1916 [176, 178], the experimental evidence of their existence was not found until decades later. In 1974, the first binary pulsar system was discovered. By tracking the radio emissions of one star in the binary system over the years, it was later established that their orbital period was decreasing in a way that is fully consistent with the GR prediction that the system should be emitting gravitational waves [470, 444]. This effect, providing indirect evidence for gravitational waves, was later confirmed through analyses of other binary neutron star systems [424]. After decades of dedicated experimental efforts to build gravitational waves detectors, the first direct detection was achieved by LIGO’s interferometers in 2015 [18] with the merger of two black holes of around 30 and 35 solar masses.

Gravitational waves can be generated by a variety of sources, and their frequencies are inherently related to the properties of the sources, which in turn determines the methods used for their detection. Figure 4.11 summarizes the spectrum of gravitational waves in terms of their frequencies and their associated detection techniques. As emphasized in the previous chapter, the method of choice

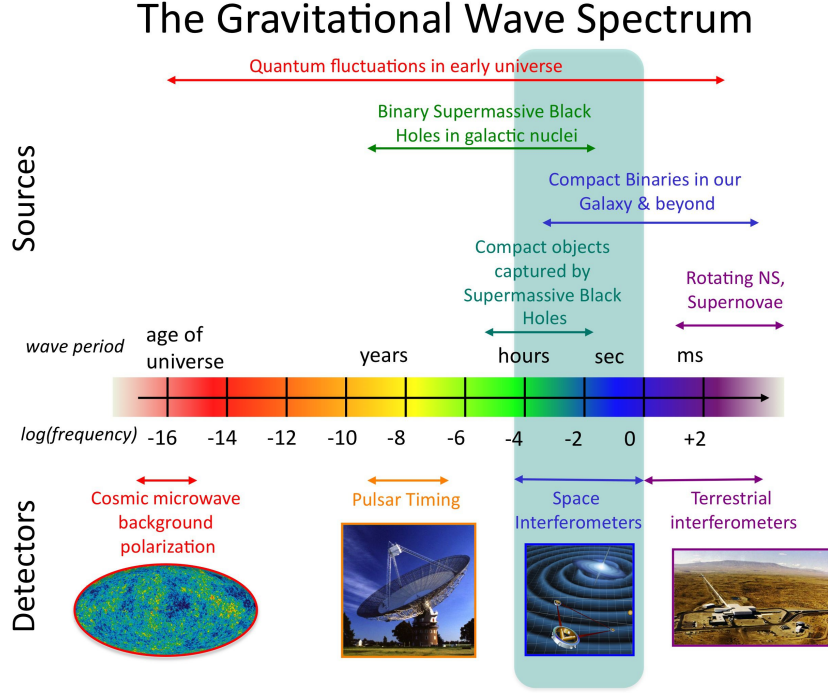


Figure 4.11 – The spectrum of possible frequencies for gravitational waves sources and their corresponding detectors. *Credit:* NASA Goddard Space Flight Center.

for detecting primordial gravitational waves, whose wavelengths stretch to cosmological scales and therefore have extremely low frequencies—in the leftmost end of the spectrum in figure 4.11—is an indirect one through the B-mode signal in the anisotropies of the CMB. Pulsar timing arrays techniques, rely on the use of the stable spin periods of millisecond pulsars to look for correlations in the deviations of the pulse arrival times caused by low frequency gravitational waves. These techniques are mostly sensitive to nanohertz frequency gravitational waves sourced by supermassive black holes, and potentially cosmic strings, phase transitions or inflation [451, 114, 286, 306]. The NANOGrav experiment have recently reported possible “first hints” of a low-frequency ($\sim 1 - 100$ nHz) gravitational wave background [67], however the evidence is not yet conclusive and more data is needed to confirm the nature of the observed signal. Interferometry techniques from the ground have been detecting many compact binary mergers since the first detection in 2015. These use Michelson interferometers, with arms lengths of about $L = 3$ km. With a typical gravitational wave strain of $h = \delta L / L = 10^{-21}$, the detector should be sensitive to a difference of about $\delta L \sim 10^{-18}$ m in arms lengths. The required sensitivity, in addition to the need for a robust control of various noise sources (e.g. seismic, thermal, quantum) makes the measurement very difficult. Despite these challenges detections are now made routinely by the LIGO and Virgo detectors [21], and a sizable statistical sample is being progressively collected paving the way for exciting science analyses. In addition to regular upgrades of the LIGO and Virgo detectors, new experiments are joining or in the process of joining the detectors network such as KAGRA in Japan [268]

and LIGO India [404]. The prospects for space interferometers are also exciting with LISA [59], which is expected to be launched in the mid 2030s. Its 2.5 million kilometers arms will enable detecting gravitational waves with much lower frequencies than is achievable from the ground, and will therefore be sensitive to sources inaccessible from the ground. Examples of sources that LISA will be able to probe include the stochastic gravitational waves background, supermassive black hole mergers, or extreme mass-ratio inspirals (EMRIs) [59].

Standard sirens

A measurement of interest for cosmology enabled by gravitational waves is the use of the so-called *standard sirens* [407, 246]. These are gravitational waves with an electromagnetic counterpart, which allows to identify the host galaxy of the source (or a set of possible hosts). The gravitational wave signal encodes information about the luminosity distance to the source, and we can get the redshift from the electromagnetic counterpart. This enables the construction of a redshift-distance relation and hence allows a measurement of H_0 in a way that is completely independent of the distance ladder measurements mentioned in the previous section, or the early-time model dependent measurements from the CMB or BAO. Figure 4.12 shows a recent standard-siren measurement of H_0 with the GW170817 event combined with statistical information from binary black hole mergers from the first and second observing runs of the Advanced LIGO and Virgo detectors [17]. As it stands, the error bars on H_0 from gravitational waves measurements are still too high to weigh in on the current H_0 tension. However the measurement accuracy is expected to improve in the coming years, reaching about 1% with $\mathcal{O}(100)$ detections [137].

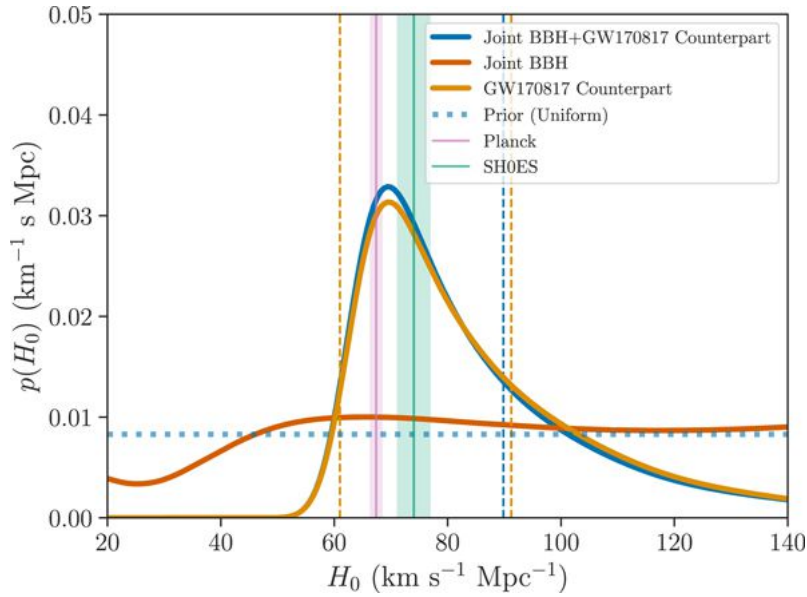


Figure 4.12 – Standard-siren measurement of H_0 compared with the other measurements from Planck[373] and SH0ES [395]. *Credit:* LIGO and Virgo collaborations [17].

Extracting cosmology with current and future CMB experiments

Contents

5.1	The experimental landscape	99
5.1.1	Overview of CMB experiments	99
5.1.2	Selected Chilean experiments	102
5.1.3	Future experiments	110
5.2	Challenges of CMB observations	111
5.2.1	Systematics	112
5.2.2	Astrophysical foregrounds	115
5.2.3	Delensing	118
5.3	Data analysis pipelines for modern CMB experiments	119

After reviewing the overall status of our understanding of the physics governing the Universe as a whole, and the main scientific questions that motivate the numerous cosmological observations of the past, present and future, we will now more thoroughly focus on CMB observations. In this chapter, we provide a detailed review of the process by which the cosmological information is processed and extracted in these experiments. The discussion will be led in three main steps: we first start by giving an overview of the landscape of CMB observations, introducing some of the major CMB experiments of the last decades along with some of their main achievements. As a member of the POLARBEAR/Simons Array and Simons Observatory experiments, I focus more on these experiments before finishing with a summary of some of the future experiments which are expected to start their operations by the end of this decade. The second section of this chapter will concern the main challenges encountered in the analysis of the data collected by CMB experiments, and the last section will cover the different stages of the data analysis process from the collection of the raw detector data to the estimation of cosmological parameters.

5.1 The experimental landscape

CMB observatories constitute a broad and rich landscape, both in terms of their scientific targets and their employed technologies. In this section we discuss aspects of CMB experiments with a focus on the science drivers.

5.1.1 Overview of CMB experiments

Ever since the first measurement of the CMB by Penzias and Wilson [365] in the mid 1960s, there have been numerous experimental efforts, spanning ground, balloon, as well as satellite experiments, to characterize this radiation ever more precisely. In the 1990s, a major milestone was attained after the measurement of the blackbody spectrum of the CMB by the FIRAS spectrometer and the measurement of the large scale CMB temperature anisotropies by the DMR instrument, both onboard the COBE satellite [423, 328]. In the present day, the temperature anisotropies have been extensively studied and the focus has shifted more towards the measurement of the polarization anisotropies. These efforts are expected to continue in the future with experiments that will soon start operating or others which are being prepared for the end of this decade. Figure 5.1 represents a summary of the history of observations post COBE/DMR featuring some of the major CMB experiments with some of their important scientific achievements.

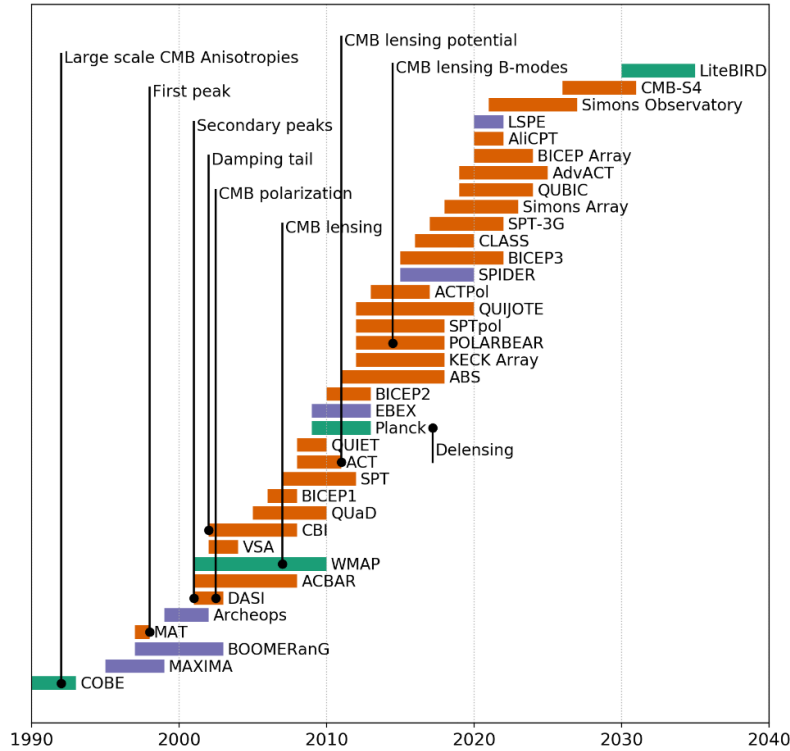


Figure 5.1 – History of some of the major CMB experiments along with their important milestones. Satellite experiments are shown in green, balloon-borne in purple, and ground-based in orange. *Credit: Dominic Beck [87].*

The measurement of the first acoustic peak was conducted by the Mobile Anisotropy Telescope on Cerro Toco (MAT/Toco) [333] as well as multiple

balloon-borne experiments, the Balloon Observations Of Millimetric Extragalactic Radiation and Geophysics (BOOMERanG) [157], The Millimeter Anisotropy eXperiment IMaging Array (MAXIMA) [227] and Archeops [91]. Subsequent measurements from the Degree Angular Scale Interferometer (DASI) [225], BOOMERanG [156], and the Very Small Array (VSA) [408] unveiled the secondary peaks in the CMB temperature power spectrum. High-resolution measurements of the CMB temperature up to $\ell = 3500$ performed by the Cosmic Background Imager [361], followed by similar measurements from the Arcminute Cosmology Bolometer Array Receiver (ACBAR) [302], revealed the CMB damping tail. In 2002, the DASI experiment reported the first detection of E-mode polarization with a high confidence of 4.9σ [297].

The two satellite successors of COBE have helped achieve significant progress in the characterization of the CMB anisotropies. The first satellite, the Wilkinson Microwave Anisotropy Probe (WMAP), was launched in 2001 towards the Sun-Earth L2 Lagrangian point from which it observed the sky for nine years, with sub-degree angular resolution, in five frequency bands from 23 to 94 GHz. WMAP delivered high resolution, full-sky maps in both temperature and polarization, enabling measurements of the CMB temperature and E-mode power spectra with unprecedented precision thereby drastically improving uncertainties on cosmological parameters and converting cosmology to a precision science. The first evidence for weak gravitational lensing of the CMB with a 3.4σ detection, was derived by cross-correlating the reconstructed lensing potential, from WMAP sky maps, and radio galaxy counts from the NRAO VLA Sky Survey (NVSS) [421]. The second satellite, Planck, launched in 2009, observed the sky for more than four years in nine frequency bands ranging from 30 to 857 GHz down to 5 arcminute resolution. It delivered the most precise full-sky maps of CMB temperature and polarization to date, enabling sub-percent precision measurements of the flat Λ CDM parameters. Figure 5.2 shows an illustration of the resolution of the maps delivered by the three generation of satellite experiments. Planck also delivered full-sky lensing maps and allowed the demonstration of delensing techniques (Cf. the second section of this chapter) for the first time [305].

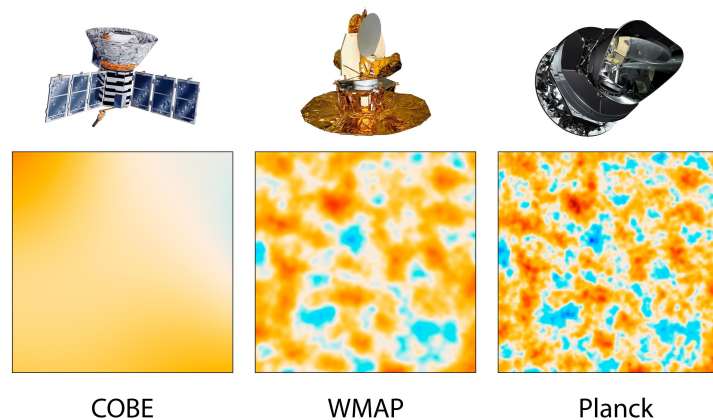


Figure 5.2 – 10 deg^2 sky patches from the full-sky maps delivered by the three generation of satellites, COBE, WMAP, and Planck (left to right). *Credit:* NASA/JPL-Caltech/ESA.



Figure 5.3 – Illustration of some of the main operating ground-based CMB experiments at the South Pole and at the Atacama desert in Chile. *Top left:* SPT. *Top right:* BICEP. *Bottom left:* POLARBEAR. *Bottom right:* ACT. *Credit:* SPT / BICEP / Keck / POLARBEAR / ACT collaborations.

Currently, many ground experiments are making high precision measurements of the CMB anisotropies and are regularly upgrading their instruments to improve their sensitivities. Two of the major experiments are located on Cerro Toco in the Atacama desert in Northern Chile, at an altitude of 5190 m. These are the Atacama Cosmology Telescope (ACT) and the POLARization of the Background Radiation (POLARBEAR). Two other major experiments are operating from the Amundsen-Scott South Pole Station. These are the Background Imaging of Cosmic Extragalactic Polarization (BICEP) and the South Pole Telescope (SPT). In the following we give a brief overview of these experiments. We leave the discussion of the POLARBEAR experiment to the next section, where we also cover the upcoming Simons Observatory which results from the fusion of the POLARBEAR and ACT collaborations.

ACT is a 6m telescope designed for arcminute resolution measurements. It started its first observations in 2007, with the Millimeter Bolometric Array Camera (MBAC), in three frequency bands, 148, 218, and 277 GHz, until 2010. The second generation receiver, ACTPol, conducted observations between 2013 and 2016, with multi-chroic polarization sensitive detectors at 98 and 150 GHz bands [450]. The third generation receiver, AdvACT, expands on ACTPol by adding high frequency (150/230 GHz) and low frequency (27/39 GHz) multichroic arrays, it has been surveying almost half of the sky since 2016 [293]. In 2011, ACT reported the first detection of the power spectrum of the CMB lensing potential with a significance of 4σ [154]. It also achieved the first statistical detection of the kinematic Sunyaev Zel'dovich effect [228].

The BICEP/Keck program are a series of small-aperture refracting telescopes targeting the degree-scale CMB B-mode signal, currently producing the deepest large-scale polarization maps. BICEP1 conducted operations from 2006 to 2008 at two frequency bands, 100 and 150 GHz [275]. BICEP2 improved on the previous instrument by greatly increasing the number of detectors. It performed

observations from 2010 to 2012, at the 150 GHz frequency band [353]. The Keck array consists of five receivers similar to BICEP2, it started observations in 2010. In addition to observations at the 150 GHz frequency band, it also deployed two receivers operating at 95 GHz and two others operating at 220 GHz [425]. BICEP3 started science observations in 2015 with 2560 detectors operating at 95 GHz [48]. It constitutes a prototype of the four BICEP3-like receivers of the BICEP Array, which should feature more than 30,000 detectors observing in six frequency bands from 30 to 270 GHz. The 30/40 GHz receiver was deployed during the last summer of 2020. The BICEP Array program is expected to reach a sensitivity to the tensor-to-scalar ratio, $\sigma(r) \sim 0.003$ [259]. Combined with Planck, the BICEP2/Keck array data set the tightest constraints on the tensor-to-scalar ratio, $r < 0.044$ at the 95% confidence level [454].

SPT is a 10 m telescope conducting high resolution (arcminute) surveys of the CMB. The SPT-SZ survey which was operating from 2008 to 2011, observed 2500 deg^2 of the sky at three frequency bands, 95, 150, 220 GHz [135]. Equipped with polarization-sensitive detectors, the SPTpol receiver conducted a 500 deg^2 survey at 95 and 150 GHz between 2012 and 2016 [70]. Since 2017, the third generation receiver, SPT-3G, started observations with 16,140 tri-chroic detectors in three frequency bands, 95, 150 and 220 GHz, targeting 1500 deg^2 with a map depth of $2 - 3 \mu\text{K-arcmin}$ in temperature at 95 and 150 GHz [130].

5.1.2 Selected Chilean experiments

POLARBEAR

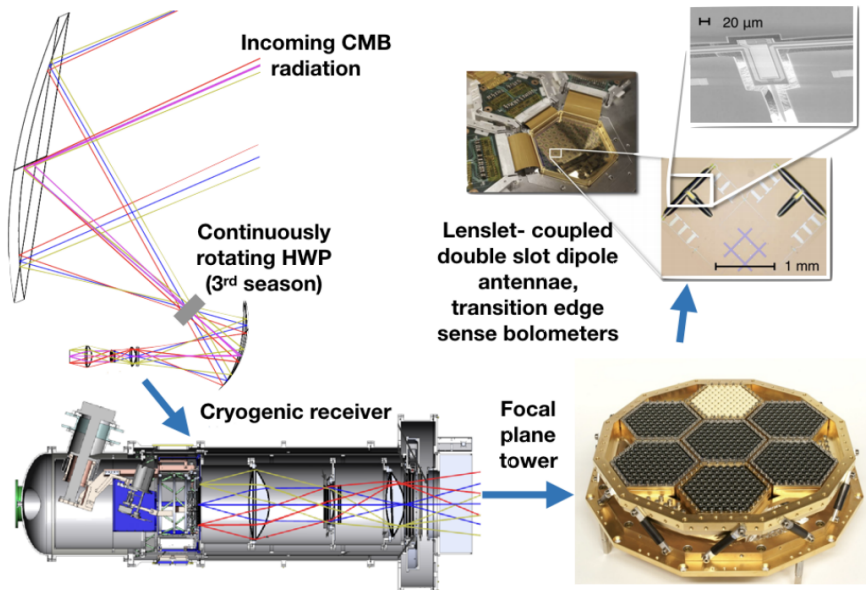


Figure 5.4 – Diagram illustrating the path of incoming light towards the polarization sensitive detectors of the POLARBEAR receiver. *Credit: Neil Goeckner-Wald.*

The POLARBEAR experiment [280] is installed on the Huan Tran Telescope (HTT) with an off-axis Gregorian-Dragone design which satisfies the

Mizugutchi-Dragone (MD) condition minimizing cross polarization and astigmatism [452]. The primary mirror with a diameter of 2.5 m enables relatively high resolution measurements (down to $\ell \approx 3000$) with a 3.5' full width at half maximum (FWHM) beam at 150 GHz. The incoming light is focused to a cryogenic receiver, where it goes through several optical components focusing the sky signal to the detector array and minimizing thermal loading or spurious signals from other optical elements. The focal plane consists of seven wafers populated by 637 polarization sensitive pixels cooled at 250 mK, with an angular footprint of 2.4° on the sky. Each pixel is a lenslet-coupled double slot antenna. The signals picked up by the antennae are then dissipated in transition-edge sensor (TES) bolometers. Measurements of the POLARBEAR experiment from 2012 to 2017 were sensitive to a single frequency band, 150 GHz. Figure 5.4 shows a schematic diagram of the optical path through different elements of the telescope from the primary mirror all the way down to the polarization sensitive pixels. In February 2014, a continuously rotating half-wave plate for polarization modulation was installed at the focus of the primary mirror in the HTT. Details about its operation and performance can be found in [440].

Scan strategy In the first two seasons, POLARBEAR was targeting three $3^\circ \times 3^\circ$ sky patches, so-called *small patches*, located in regions with low foreground contamination, for sub-degree scale B-mode measurements. Since July 2014, the experiment was observing a larger sky patch of 670 deg^2 area, so-called *large patch*, allowing to also constrain inflation with degree-scale B-mode measurements. Figure 5.5 shows the location of the observed regions on the sky.

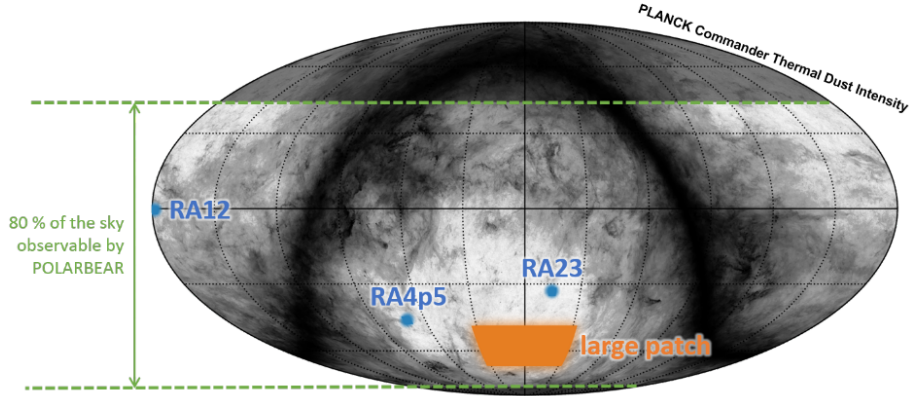


Figure 5.5 – The four sky patches observed by the POLARBEAR experiment.

The blue regions correspond to the small patches named RA4.5, RA12, and RA23 after their right ascension coordinates. The orange region corresponds to the large patch. *Credit: Dominic Beck [87]*

Typically, ground-based CMB experiments, including POLARBEAR, conduct their sky surveys in the form of a series of constant elevation scans (CES) to avoid atmospheric loading. For any given CES, the telescope observes the sky with left and right scans in azimuth. Except during the *turnarounds* where the telescope has to stop and change the scan direction at the edges of the scanned sky area, the scan velocity is kept constant. For small patch observations, the scan speed is set to $0.75^\circ \text{ s}^{-1}$, while it is set to 0.4° s^{-1} for large patch observations. The duration of

each CES depends on the scanned sky patch, it is set to 15 min for small patch observations, with roughly 200 subscans (left to right or right to left scan) per CES, and 1 hour for large patch observations, yielding about 70 subscans per CES. The telescope is repointed after every 15 min CES in small patch observations, while it is only repointed every 4 hours in large patch observations. The detectors are retuned every hour however. Science observations are also interrupted every 24 or 48 sidereal hours for the recycling of the cryogenic refrigeration system. After the fridge cycle, calibration measurements are performed. One important property that is ensured by the fact that the same sky patch is observed at different elevations in both rising and setting configurations, is the so-called *cross-linking*, which refers to the fact that the same sky is observed with multiple attack angles. An illustration is provided in figure 5.6. This is particularly crucial for the map-making step, since it allows to mitigate noise correlations effects which can be strong for ground observations given the presence of atmosphere. The presence of a HWP also helps to significantly diversify the pixel attack angles thereby improving the reconstruction of the polarization signal.

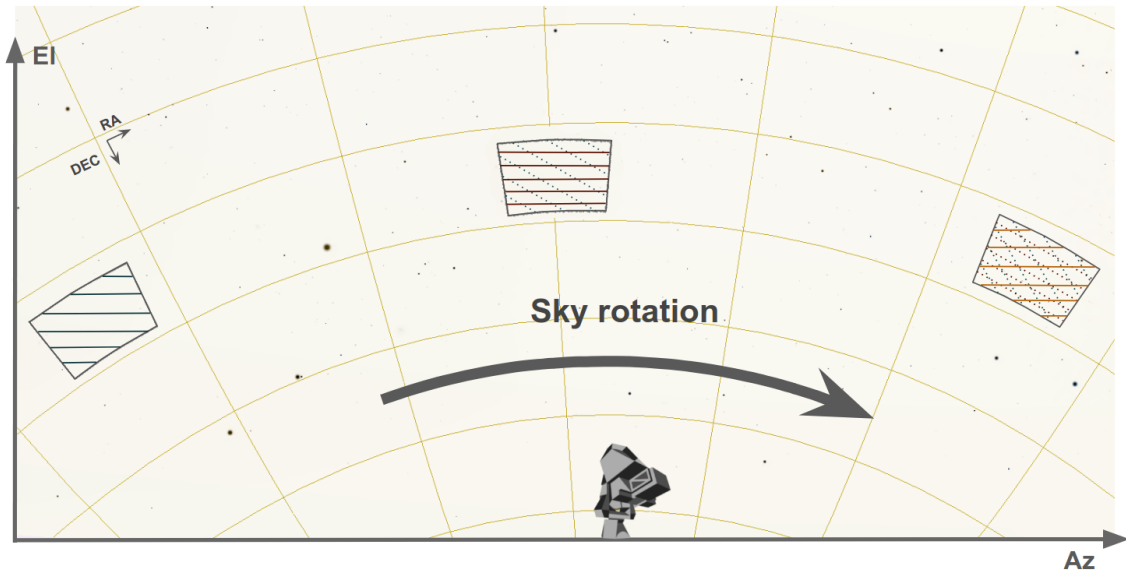


Figure 5.6 – Illustration of the cross-linking resulting from scanning the same sky patch at different moments of the day at different elevations. The solid lines trace the attack angle of the scan occurring at the current position of the sky patch, while the dotted lines trace the attack angles of previous scans.

Science results Following the first season of observations of small patch observations, POLARBEAR reported the first direct measurement of B-modes power spectrum at sub-degree angular scales, ruling out the null hypothesis of no lensing B-modes with a confidence of 97.1% [354]. Another result was the first detection of the lensing power spectrum, with a statistical significance of 4.2σ using only CMB polarization data [42]. In addition, a measurement of CMB lensing, via cross-correlation of the polarization data with the CIB from the Herschel satellite, with a statistical significance of 4.0σ was also reported [41]. Furthermore, E-mode and B-mode CMB measurements from the first season also allowed to derive constraints on anisotropic cosmic birefringence (<93 nG, at the 95% CL, expressed in terms of primordial magnetic field strength equivalent) and primordial magnetic

fields (<3.9 nG at the 95% CL) [37].

Data from the first season was later re-analyzed in combination with the second season data, through two distinct pipelines, one of them performing a filter-and-bin map-making procedure (Cf. the third section of this chapter) thereby producing biased sky maps and correcting for the bias at the power spectrum level, the second pipeline, which is computationally more expensive, produces unbiased maps of the sky by accounting for the filtering simultaneously with the map projection. The polarization map depths of the three sky patches observed after the two seasons is between 5 and 7 μK -arcmin. This analysis led to a twofold improvement in the B-mode band power uncertainties, enabling the rejection of the null-hypothesis of no B-modes at a confidence of 3.1σ [33].

Since I joined the collaboration in 2018, we conducted a new set of studies. In a recent analysis of the two season polarization data [191] we reject the null hypothesis of no lensing with a high significance of 10.9σ , including both statistical and systematic uncertainties, with a measurement of the lensing amplitude A_L (which parametrizes the amplitude of the lensing autospectra, $A_L = 1$ in the fiducial ΛCDM model)

$$A_L = 1.33 \pm 0.32 (\text{statistical}) \pm 0.02 (\text{systematic}) \pm 0.07 (\text{foreground}). \quad (5.1)$$

We also conducted cross-correlation analyses with the small patch measurements. One such analysis examined the cross-correlation between CMB lensing and cosmic shear from data of the Subaru Hyper Suprime-Cam (HSC) survey. The measured lensing amplitude of the cross-spectrum is $A_{\text{lens}} = 1.70 \pm 0.48$, rejecting the null hypothesis with a 3.5σ confidence [347]. We performed another analysis to measure the cross-correlation of the CMB lensing convergence reconstructed from polarization data collected in the first two observing seasons of POLARBEAR, and the galaxy overdensity field from the Herschel-ATLAS survey. The result is a 4.8σ rejection of the no-lensing hypothesis and an inferred best-fit galaxy bias parameter, $b = 5.76 \pm 1.25$ [192]. Finally, small patch observations also enabled the first demonstration of internal B-mode delensing using only CMB polarization data [28], and achieving a B-mode power variance reduction of 22% using iterative delensing methods [128, 129]. This also represents the first application of these methods to real data.

We also published results from large patch observations. These were inferred from the analysis of three seasons of data reaching an effective polarization map noise level of $32\mu\text{K}$ -arcmin over the 670 deg^2 sky patch. In one paper we report the measurement of the degree scale B-mode power spectrum shown to be consistent with ΛCDM lensing and a single dust foreground model [26]. The data allows to set an upper limit on the tensor-to-scalar ratio of $r < 0.90$ at a 95% confidence level. More efforts are currently in progress to improve this result. Another result we reported is the measurement of the E-mode power spectrum over a multipole range $500 \leq \ell < 3000$ shown to be consistent with ΛCDM [27].

POLARBEAR-2 and Simons Array The POLARBEAR experiment is currently being upgraded to prepare for the next generation, higher sensitivity measurements. A recent milestone in this direction is the deployment of the new POLARBEAR-2a instrument which achieved the first light in January 2019 [273]. Its focal plane consists of 7 wafers with a total of 7,588 detectors, which represents

a sixfold improvement on the number of detectors of the POLARBEAR instrument. Instead of dual slot antennas, pixels consist of broadband sinuous antennas coupled to superconducting transition edge sensors (TES) bolometers [436]. This technology allows POLARBEAR-2a pixels to be simultaneously sensitive to two frequency bands, 95 GHz and 150 GHz. The left panel of figure 5.7 shows a picture of the sinuous antennas and the TES bolometers. An ambient temperature continuously rotating half-wave plate (HWP) is mounted for polarization modulation [243]. The receiver is expected to reach a noise equivalent temperature of $5 \mu\text{K}_{\text{CMB}}\sqrt{s}$ in each of the two frequency bands. The right panel of figure 5.7 shows a planet observation of Venus from first light observations in January 2019. These observations are used for calibration tasks such as characterization of the beam shape.

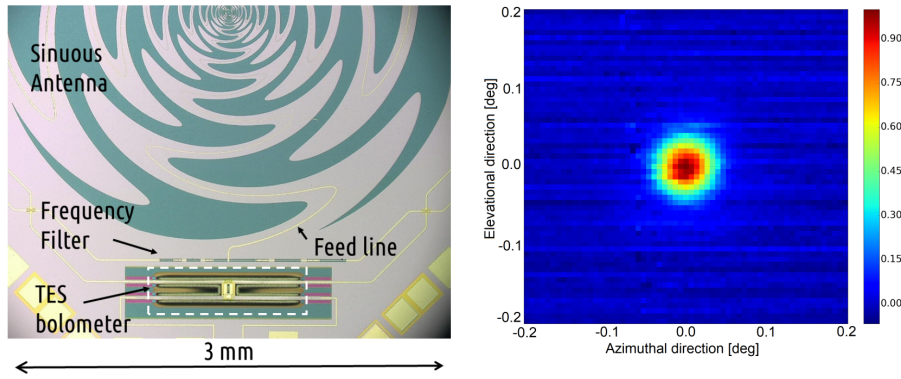


Figure 5.7 – *Left*: Microscopic image of sinuous antenna and TES detector of POLARBEAR-2a. *Right*: Venus map observed with 150 GHz TESs in one detector module. *Credit*: POLARBEAR collaboration [273].

The deployment of POLARBEAR-2a is the first step towards the deployment of the full scale Simons Array which should feature two other similar telescopes and receivers. POLARBEAR-2b should also be observing in the same frequency bands as POLARBEAR-2a, 95 and 150 GHz, while the POLARBEAR-2c receiver, which is expected to replace the old POLARBEAR receiver on the same telescope, will observe at 220 and 270 GHz to control foreground contamination. Both POLARBEAR-2b and POLARBEAR-2c deploy cryogenic HWPs [242]. When operating at full capacity, the three telescopes will thus have 22,764 detectors observing the sky in four frequency bands with a full array sensitivity of $2.5 \mu\text{K}_{\text{CMB}}\sqrt{s}$ [427]. The sky survey is targeted towards the low foreground regions, avoiding the galactic disk, and covers roughly 10% of the sky. An illustration is given in figure 5.8. The survey is also optimized to overlap with other surveys for cross-correlation studies. In particular, with surveys of the Simons Observatory Small Aperture Telescopes, which we will present in more details in the following paragraphs, as well as with the Hyper Suprime-Cam (HSC) on the Subaru Telescope [336], the Dark Energy Spectroscopic Instrument (DESI) [185], and the Vera C. Rubin Observatory [317]. Simons Array is expected to deliver competitive cosmological constraints, enabling detection of primordial gravitational waves for $r > 0.01$ with $> 5\sigma$. It will also constrain the spectral index n_s with an uncertainty $\sigma(n_s) = 0.0015$, and the total mass of neutrinos with $\sigma(\sum m_\nu) = 40 \text{ meV}$ through lensing measurements combined with DESI BAO data [427].

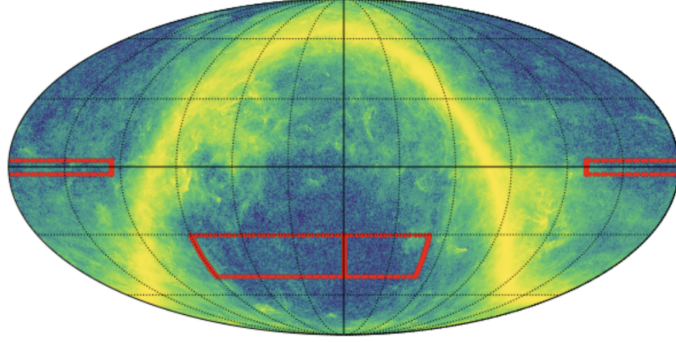


Figure 5.8 – The expected Simons Array sky survey shown on top of the Planck 353 GHz map. The survey will target the low dust regions to minimize the effect of foregrounds. *Credit: Aashrita Mangu.*

Simons Observatory

The Simons Observatory is the result of a joint effort between the ACT and POLARBEAR collaborations as well as many other scientists, gathering more than 300 researchers from many institutions across the globe. The collaboration is preparing a novel CMB experiment, which is currently being built in the Chilean Atacama desert on Cerro Toco, the same site that is currently hosting the ACT, POLARBEAR, and CLASS [184] telescopes. The Simons Observatory will aim at measuring CMB temperature and polarization anisotropies from few-degree scales down to arcminutes, thereby targeting a broad set of science goals. For this purpose, two types of telescopes will be deployed. A 6 m Large Aperture Telescope (LAT), with about 30,000 detectors, for high resolution measurements, and three 42 cm Small Aperture Telescopes (SATs), also featuring 30,000 detectors, which will perform lower resolution measurements targeted for the large scale science. The total 60,000 detectors deployed, representing an order of magnitude improvement on the current generation of experiments, will thus allow CMB measurements with unprecedented sensitivities. The first light for both the LAT and SATs telescopes is expected soon between mid-2022 and early 2023.

The LAT has a 6 m diameter primary and secondary mirrors with a crossed Dragonite design and a field of view of 7.2° at 90 GHz [209]. The telescope receiver (LATR) can host up to 13 optics tubes, seven of which will be populated in the initial deployment, each supporting three detector wafers. Four optics tubes will be sensitive to the mid-frequency bands (MF), 90/150 GHz, two to the ultra-high frequency bands (UHF), 220/270 GHz, and one to the low frequency bands (LF), 27/39 GHz. The SATs are 42 cm aperture refractors, each hosting seven detector wafers. Initially two MF SAT and one UHF SAT will be deployed. One of the MF receivers will later be replaced by a LF SAT. Each SAT has a large field of view of 35° and a continuously rotating cryogenic half-wave plate (CHWP). Each SAT is mounted on a platform and surrounded by a ground shield to minimize spurious signals from the ground. Figure 5.9 shows an illustration of the telescopes design along with up-to-date pictures from ongoing tests and integration of the receivers. The Simons Observatory is adopting two different sensor technologies, also shown in figure 5.9: lenslet coupled sinuous antennas developed at the University of California, Berkeley (UCB), and the horn coupled orthomode trans-

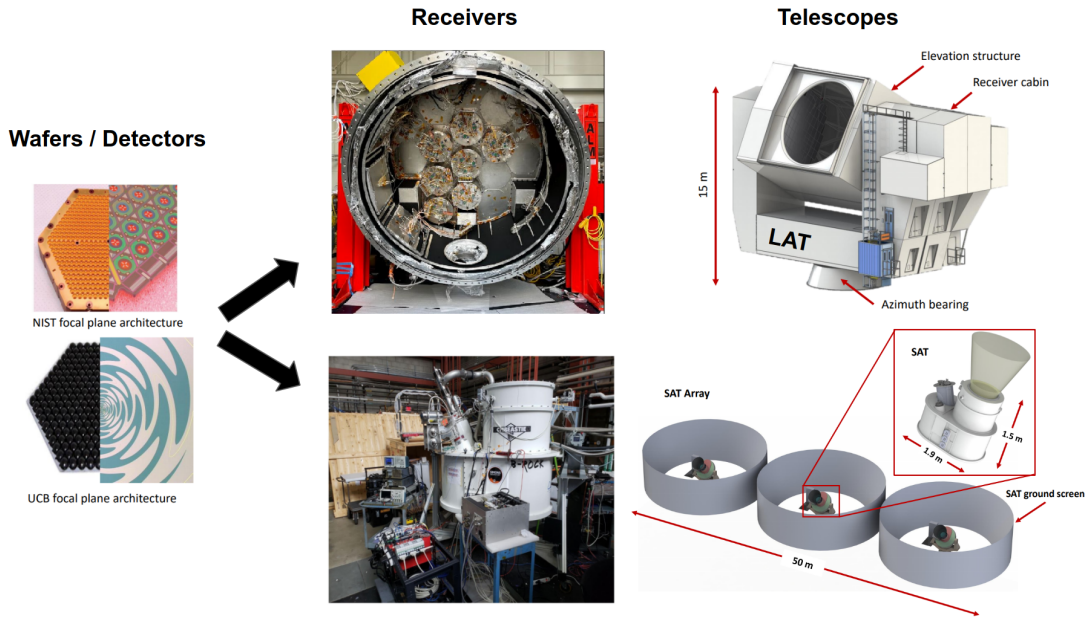


Figure 5.9 – Illustration of the Simons Observatory instruments designs.

ducers (OMTs) produced at the National Institute of Standards and Technology (NIST).

The survey of the SATs will cover roughly 10% of the sky with more weight in the Southern sky, overlapping with the sky patch of the BICEP/Keck array and focusing on regions with low galactic dust emission, in addition to limiting the range of elevations in order to minimize atmospheric loading. For the LAT the survey will cover around 40% of the sky optimized to maximally overlap with DESI and the Rubin Observatory surveys, and avoid the brightest regions of the galaxy. Figure 5.10 shows the footprint of the two sky surveys. The resolutions and expected temperature white noise levels of the two surveys at different frequency bands are summarized in table 5.1.

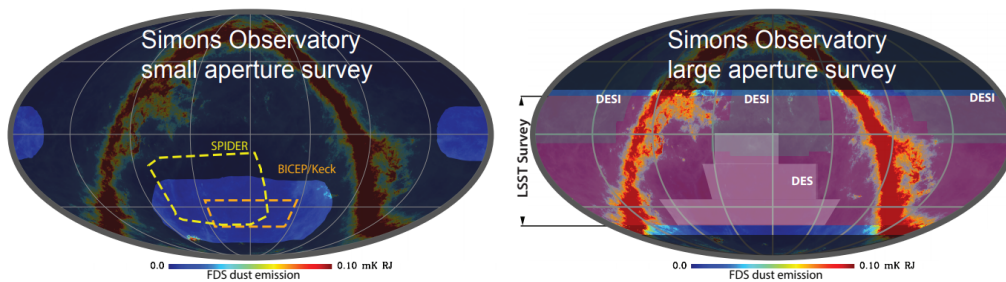


Figure 5.10 – Simons Observatory SATs and LAT surveys footprints on top of galactic dust emission. *Credit: Simons Observatory collaboration [36].*

The Simons Observatory data will bring unique insights on many scientific questions. These include constraining models of the early Universe through the search for primordial gravitational waves, characterizing the shape of the primordial spectrum, and setting constraints on non-Gaussianity. It will inform the search for light relics by setting constraints on the effective number of relativistic species. It will enable setting tight constraints on the total mass of neutrinos

Center frequency [GHz]	27	39	93	145	225	280
SAT						
FWHM [arcmin]	91	63	30	17	11	9
Noise [μ K-arcmin]	35	21	2.6	3.3	6.3	16
LAT						
FWHM [arcmin]	7.4	5.1	2.2	1.4	1.0	0.9
Noise [μ K-arcmin]	71	36	8.0	10	22	56

Table 5.1 – Summary of anticipated resolutions and sensitivities of SO instruments. Taken from Tab. 1 in [36].

and their mass hierarchies. The CMB lensing measurements will probe the large-scale distribution of matter, which in turn informs on the nature of dark matter and dark energy. Moreover, the LAT survey will also be sensitive to the thermal and kinetic Sunyaev-Zel’dovich effects which will be leveraged to help improve our understanding of galaxy formation and evolution, in addition to enabling complementary methods to be used to constrain the same cosmological parameters. This will be useful for consistency checks and for breaking degeneracies. Furthermore, Simons Observatory data will also constrain several characteristics of reionization. Table 5.2 shows the forecasted uncertainties of the primary science goals that will be pursued by the Simons Observatory. The LAT survey will also deliver a legacy catalog of 16,000 galaxy clusters and more than 20,000 extragalactic sources. The high cadence of the scan strategy will also allow to monitor variable sources, in particular AGNs and afterglows of gamma-ray bursts (GRBs), for transient science [36].

Parameters		Current	SO-Baseline	So-Goal	Method
Primordial perturbations	r	0.03	0.003	0.002	BB + external delensing
	$e^{-2\tau}\mathcal{P}(k = 0.2 \text{ Mpc})$	3%	0.5 %	0.4 %	TT/TE/EE
	$f_{\text{NL}}^{\text{local}}$	5	3	1	$\kappa\kappa \times \text{LSST-LSS} + 3\text{-pt}$
			2	1	kSZ + LSST-LSS
Relativistic species	N_{eff}	0.2	0.07	0.05	TT/TE/EE + $\kappa\kappa$
Neutrino mass	$M_\nu(\text{eV}, \sigma(\tau) = 0.01)$	0.1	0.04	0.03	$\kappa\kappa$ +DESI-BAO
			0.04	0.03	tSZ-N \times LSST-WL
	$M_\nu(\text{eV}, \sigma(\tau) = 0.002)$		0.03	0.02	$\kappa\kappa$ +DESI-BAO+LB
			0.03	0.02	tSZ-N \times LSST-WL+BL
Deviations from Λ	$\sigma_8(z = 1 - 2)$	7%	2%	1%	$\kappa\kappa$ +LSST-LSS
			2%	1%	tSZ-N \times LSST-WL
	$H_0(\text{km/s/Mpc}, \Lambda\text{CDM})$	0.5	0.4	0.3	TT/TE/EE+ $\kappa\kappa$
Galaxy evolution	η_{feedback}	50-100%	3%	2%	kSZ+tSZ+DESI
	p_{nt}	50-100%	8%	5%	kSZ+tSZ+DESI
Reionization	Δz	1.4	0.4	0.3	TT (kSZ)

Table 5.2 – Summary of the primary science goals of the Simons Observatory. Taken from [36, 448]. Forecasting methods are explained in detail in [36].

5.1.3 Future experiments

In this section we discuss two major experiments which are part of the long term efforts of the CMB community. These are projected to start observing in the late 2020s and will achieve yet another leap in sensitivity.

CMB Stage IV

The CMB Stage IV (or CMB-S4) experiment is a joint effort involving a significant portion of the international CMB community to build the ultimate ground-based CMB observatory [14]. Similarly to the Simons Observatory, the experiment is designed to be sensitive to a wide range of angular scales, for both low- ℓ and high- ℓ science, and will thus feature multiple large aperture telescopes and small-aperture telescopes. Their design will be informed by the previous expertise accumulated by the community in the development of the previous generations experiments and their different technological choices. The telescopes will be built at two different sites: the Cerro Toco site in Chile next to ACT, POLARBEAR, CLASS and the upcoming Simons Observatory. The second is the South Pole site currently hosting BICEP/Keck and SPT.

CMB-S4 will conduct two complementary surveys. One is an ultra-deep survey covering 3% of the sky in a very low foreground region. The map noise level is expected to be below $1 \mu\text{K-arcmin}$ enabling a very sensitive search for primordial gravitational waves which will be conducted with low resolution measurements (around 30 arcminutes) with SATs combined with high resolution measurements (around 1.5 arcminutes) with LATs for delensing. Another is a high resolution (less than 1.5 arcminutes) deep and wide survey covering 60% of the sky with a noise level of $1 \mu\text{K-arcmin}$. This will be targeted towards the lensing measurements to map the distribution of matter and constraining the dark sector as well as the mass of neutrinos, in addition to the search for light relics, and delivering legacy catalogs with clusters, extragalactic sources, along with monitoring variable sources for transient science. Overall, the CMB-S4 telescopes will host more than 500,000 detectors observing the millimeter-wave sky in multiple frequency bands between 20 and 270 GHz. More detailed reviews of the CMB-S4 science goals and forecasts can be found in [14, 16].

LiteBIRD

The Light satellite for the study of B-mode polarization and Inflation from cosmic background Radiation Detection (LiteBIRD) was selected by the Japanese space agency (JAXA) in May 2019 as a strategic large-class (L-class) mission, for an anticipated launch in 2028. LiteBIRD will conduct full-sky surveys of the CMB polarization anisotropies with unprecedented sensitivity from the Sun-Earth L2 Lagrangian point, in 15 frequency bands between 40 and 402 GHz. The satellite will host three telescopes, the Low frequency telescope (LFT), the mid frequency telescope (MFT), and the high frequency telescope (HFT). A continuously rotating HWP will be mounted on each telescope. The three telescopes will count around 4500 detectors in total. Figure 5.11 shows an illustration of the design of the instruments. LiteBIRD will measure the polarization signal in angular scales inaccessible from the ground covering the multipole range $2 \leq \ell \leq 200$. This will

allow to measure or otherwise set very tight upper limits on the tensor-to-scalar ratio, with a total uncertainty of $\sigma(r) < 10^{-3}$ without delensing, and including residual foregrounds and instrumental systematics [240], thereby severely constraining the physics of inflation. LiteBIRD will also be able to provide cosmic variance limited measurements of the low- ℓ E-modes, thus setting constraints on the optical depth to reionization τ_{reio} , with an unprecedented precision of $\sigma(\tau_{\text{reio}}) = 0.002$. This will, in turn, be crucial for measurements of neutrino mass [116, 54]. LiteBIRD data will also serve to derive constraints on cosmic birefringence and non-Gaussianity. In addition, the multi-frequency large-scale polarization maps delivered by LiteBIRD will be of significant importance for the modelling of galactic foregrounds.

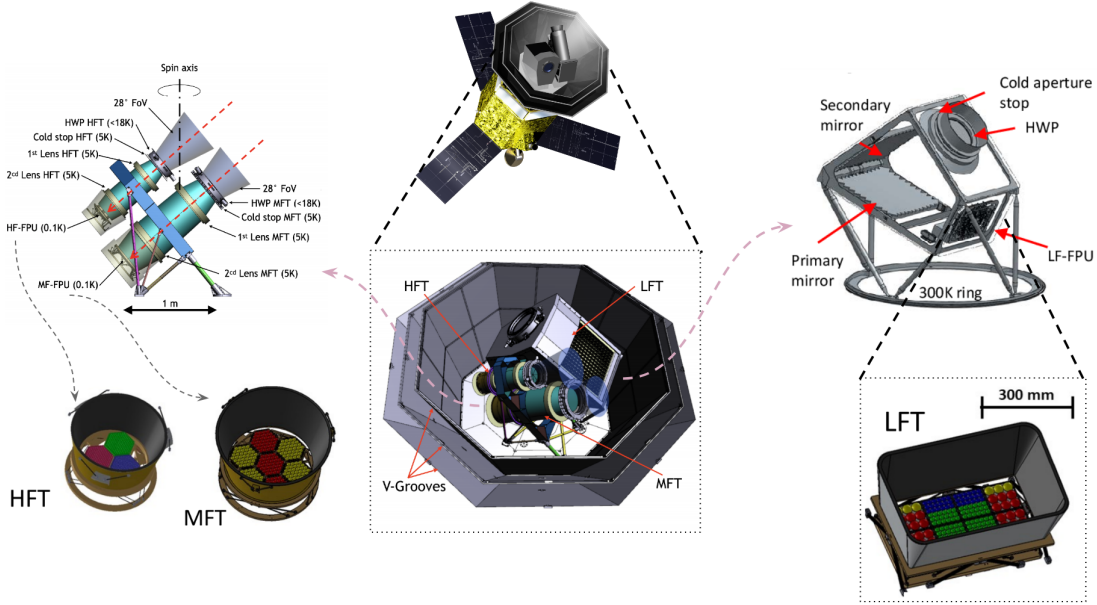


Figure 5.11 – Global design of the LiteBIRD payload module, and the onboard instruments and focal plane units. Figure compiled from 3D models given in [471, 339].

5.2 Challenges of CMB observations

Fulfilling the ambitious scientific goals of current and future CMB experiments requires us to reach levels of sensitivities where many unwanted effects become relevant and need to be taken into account, thereby complexifying the data analysis process. In this section, we give a brief overview of some of the major CMB contaminants that need to be properly modelled and subtracted from the data. Figure 5.12 shows an illustration of the typical configuration of CMB ground-based observations, emphasizing the different signals that combine into the data.

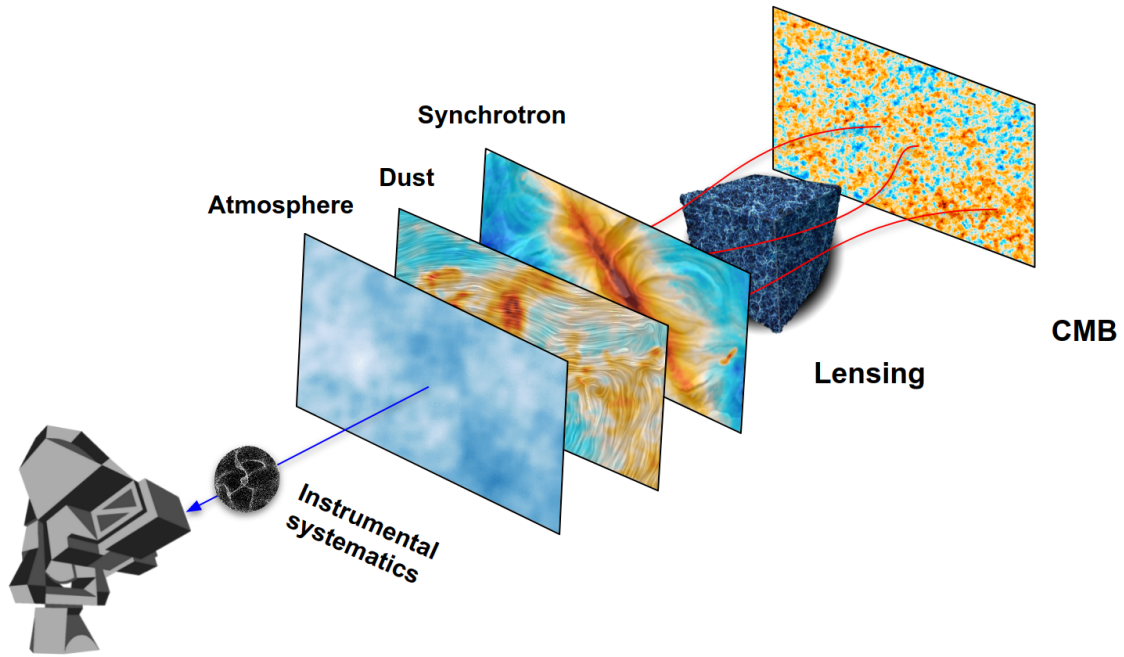


Figure 5.12 – Illustration of the different signals that intervene along the line-of-sight of a typical CMB ground-based experiment. Note that we mention only dust and synchrotron foregrounds, which are the most relevant diffused and polarized contaminants. Other contributions exist, however, at the smallest scales as well as for temperature data.

5.2.1 Systematics

The presence of systematics stemming from non-ideal instruments or calibration errors, in addition to the presence of contaminant signals from the environment where the telescope is operating, is a major challenge that needs to be conquered in order to successfully deliver robust scientific results. In the following, we cover some of the systematic effects that are particularly relevant in the first stages of the data analysis, when we are dealing with the raw time-domain data.

Atmosphere and $1/f$ noise

The challenges posed by atmosphere for ground experiments are twofold. The first one is its transmission profile which limits frequency coverage from the ground to a limited number of bands, as shown in figure 5.13. In particular, measurements in frequencies above 300 GHz are very challenging although not strictly inaccessible. Observing the sky signal in these frequencies allows monitoring polarized dust emission from the galaxy, and is thus useful for foreground subtraction. This is an area where synergies between ground-based and space-based experiments are exploited.

The second challenge is the atmospheric emission which is a considerable source of correlated noise in the measured signals both between time samples and between detectors. Figure 5.14 shows an example of simulated atmospheric emission signal measured by a single detector of the Simons Observatory Small Aperture Telescope at 150 GHz, along with its estimated power spectral density (PSD). In Fourier domain, the atmosphere behaves as $1/f$ noise in the low fre-

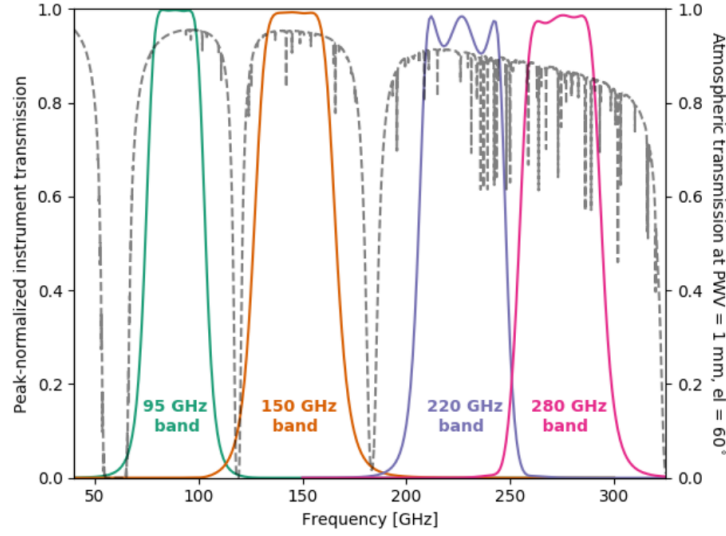


Figure 5.13 – Simons Array bandpasses on top of the atmospheric transmission profile at the Chajnantor plateau during good weather conditions. *Credit:* John Groh.

quencies, and has thus similar effects to instrumental noise. To correct for both effects, we typically down weight the lowest frequencies, or alternatively, we filter specific templates modelling the slowly varying drift of the signal. It is worth mentioning that the continuum emission of the atmosphere scales as a power law with observing frequency, $\propto \nu^\beta$ [183], therefore it may also be important to leverage the multi-frequency observations to mitigate its impact, by performing a joint estimation of the sky signal in several frequency bands. Atmosphere has only a small intrinsic polarization [82], and is therefore not expected to significantly affect the polarization measurements directly. However, the presence of ice crystals in tropospheric clouds can generate horizontally polarized signals by scattering thermal emission from the ground [455, 371]. This effect can produce polarized bursts increasing the polarization fraction up to 10% as observed for instance by the POLARBEAR experiment [439]. In addition to this effect, the unpolarized atmospheric signal coupled with other instrumental systematics responsible for intensity-to-polarization leakage, can have a relevant impact on polarization and needs to be investigated as well.

Atmospheric emission, or transmission, depend on the amount of precipitable water vapor (PWV), which is measured in millimeters, and is sensitive to properties of the observing site such as climate, and other variables such as season or weather. To minimize these effects, telescopes are typically built in the driest Earth locations such as the Chajnantor plateau in Chile or the South Pole.

Scan synchronous signals

Another prominent effect, particularly for ground experiments, is the presence of spurious signals that are synchronous with the telescope scan. These signals can be due to the response of the detectors to the local magnetic field, or to the sidelobes of the beam picking up emission from the ground, the galaxy, the Moon, the Sun, and Planets. Experiments are designed to minimize these effects,

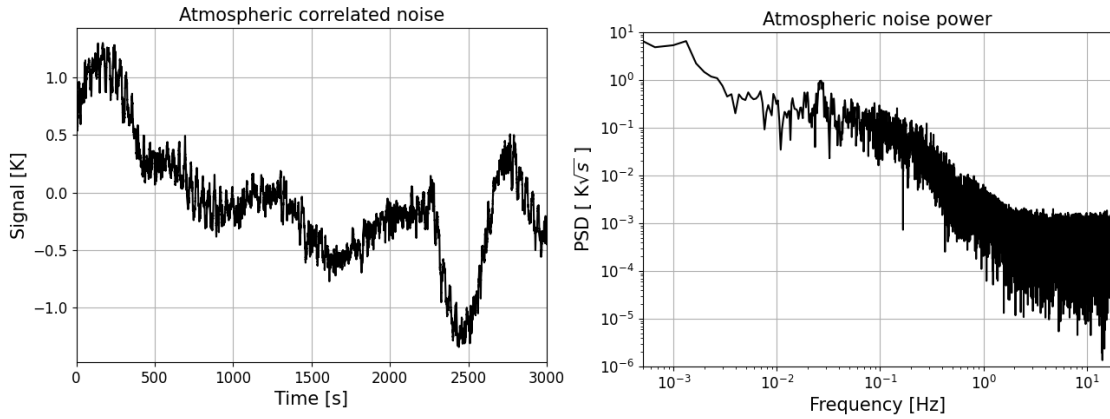


Figure 5.14 – *Left*: Example of a simulated timestream dominated by atmospheric contamination, for a single detector of the Small Aperture Telescope at the 150 GHz band. *Right*: Estimated PSD of the signal, showing a $1/f$ behavior at frequencies lower than ~ 2 Hz. The simulations are generated using the TOAST^a software package.

^a: <https://github.com/hpc4cmb/toast>

by including ground shields and baffles for example [209]. However, given the faint nature of the CMB signal we aim to measure, even small contributions cannot be neglected and need to be accurately modelled and subtracted in the data reduction.

Half-wave plate synchronous signals

As we have seen previously, a number of experiments deploy a HWP in order to modulate the polarization signal and avoid the $1/f$ noise at low frequencies, in addition to its key role in mitigating beam systematics and leakage effects. However, the HWP also introduces spurious signals of its own which are synchronous with its rotation. These signals can arise from differential transmission or reflection of unpolarized light as it propagates through the optical chain, or from polarized emission of the HWP itself. Placing the HWP at cryogenic temperatures helps minimizing its thermal emission, however the lower the temperature the lower the HWP needs to be located in the temperature stages of the optical chain which in turn worsens the effects of HWP synchronous signals (HWPSS), since any $I \rightarrow P$ leakage arising in front of the HWP is directly modulated at $4f_{\text{HWP}}$, the center of the science band [403]. For the Small Aperture Telescopes of the Simons Observatory, the HWP will be placed at the 40 K temperature stage. Figure 5.15 shows an example of simulated HWPSS for these telescopes based on realistic models developed in [403], emphasizing different contributions from transmission, reflection and emission at different frequencies of the HWP rotation. These spurious signals can be handled through time-domain data processing techniques such as demodulation procedures [266, 304], or can be estimated and subtracted simultaneously with the map reconstruction as we will discuss in the next chapter.

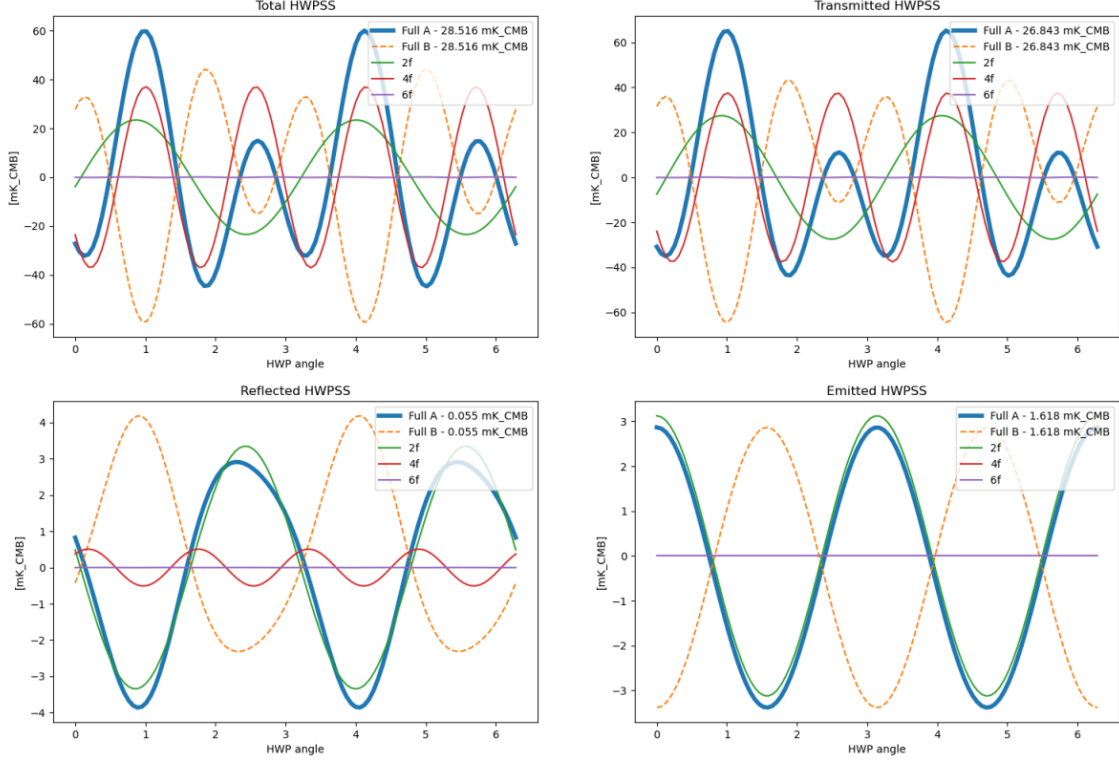


Figure 5.15 – Realistic TOAST simulation of the HWPSS signal expected in the Simons Observatory Small Aperture Telescopes, based on [403]. The different panels show the HWPSS template as a function of the HWP angle for the 2f, 4f and 6f harmonics as well as their sum for two detectors of an orthogonal pair.

Each panel represents either the total spurious signal, or the component corresponding to either transmission, reflection or emission (see titles above each panel). *Credit:* Reijo Keskitalo.

Intensity-to-Polarization leakage effects

In addition to the optical $I \rightarrow P$ leakage effects mentioned above, which arise from differential transmission or reflection of unpolarized incident light off the optics, another source of $I \rightarrow P$ leakage arises when we differentiate the signals between orthogonal pairs of detectors in the presence of gain, bandpass or beam mismatch [24, 34]. These lead to a stronger polarized $1/f$ noise limiting our constraining power at large angular scales. Polarization modulation with the HWP can handle these effects, provided we are able to efficiently mitigate the spurious effects of the HWP itself.

5.2.2 Astrophysical foregrounds

One of the most important challenges of CMB observations, particularly in the context of the search for primordial B-modes, is the presence of polarized diffuse foreground emission from the galaxy dominating over the sought-after primordial B-mode signal at all scales, as shown in the left panel of figure 5.16. Multi-frequency observations as well as accurate characterization of these foregrounds is thus crucial for their removal enabling high sensitivity measurements of the CMB polarization signal. In the following, we give an overview of two

of the major sources of these polarized foregrounds and discuss some of their properties.

It is worth mentioning, however, that these are not the only sources of foregrounds, as there is also free-free emission, the Anomalous Microwave Emission (AME) believed to be sourced by spinning dust grains, in addition to CO lines which are typically localized in frequency and thus avoided by the design of the bandpasses. However, these foregrounds are, at least in current experiments, not very relevant for polarization since their polarization fractions are typically very low.

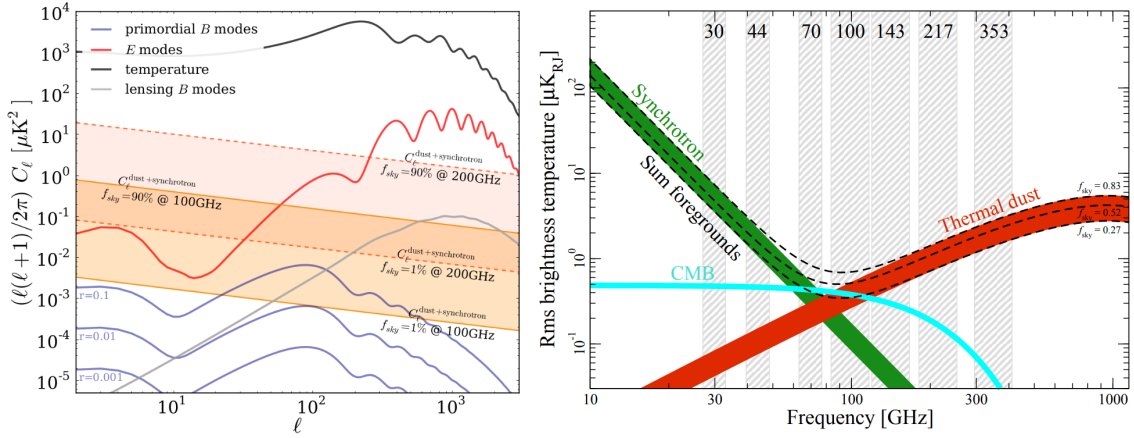


Figure 5.16 – *Left*: Expected B-mode foregrounds contamination (dust+synchrontron) on the cleanest 1-90% of the sky at 100 and 200 GHz. *Right*: RMS of the polarization amplitudes of foregrounds components as a function of frequency, compared with CMB. *Credit*: Josquin Errard / Planck Collaboration [376].

Thermal dust emission

Interstellar dust grains tend to form non spherical and elongated aggregates, behaving as electric dipoles and partially aligning along their short axis with the galactic magnetic field lines. As they are heated by unpolarized starlight they re-emit thermal radiation preferentially aligned along their long axis, thus polarized perpendicularly to the magnetic field lines. Figure 5.17 illustrates this process.

This polarized emission is dominant at frequencies higher than ~ 70 GHz (Cf. right panel of figure 5.16), therefore sky maps observed at high frequencies, such as the 353 GHz frequency channel of Planck, can be used as a useful tracer of the dust emission structure. Extrapolating the emission profile to other frequencies is crucial for foreground removal. For this purpose, the spectral emissivity of polarized thermal dust is usually modelled as a modified black body spectrum

$$A^{\text{dust}}(\nu) \propto \nu^{\beta_d - 2} B_\nu(T_d), \quad (5.2)$$

where β_d and T_d are respectively the dust spectral index and temperature, and $B_\nu(T) \equiv \frac{2h\nu^3/c^2}{e^{h\nu/kT} - 1}$ is the black body spectrum. The Planck satellite has measured a mean dust spectral index, $\beta_d = 1.53 \pm 0.02$, and a mean dust temperature, $T_d = 19.6 \pm 1.3$ K [35, 52]. However, this is a rather simple model which does

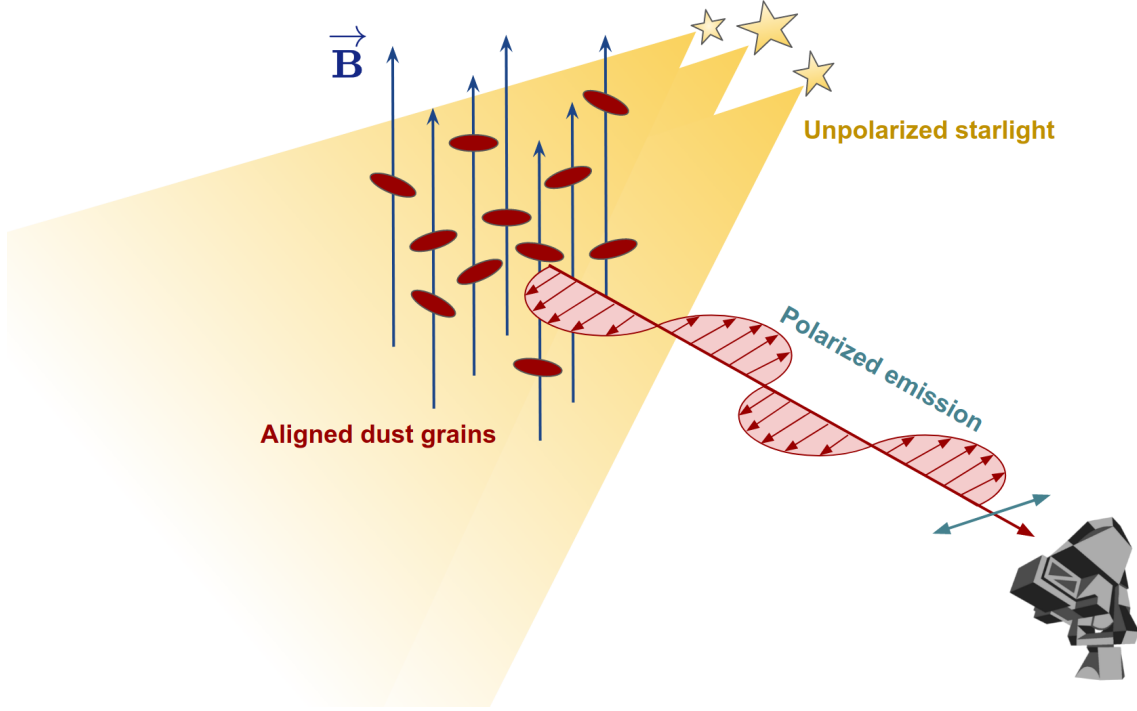


Figure 5.17 – Elongated dust grains aggregates, aligning preferentially along the perpendicular direction to the magnetic field of the galaxy, absorb unpolarized starlight and re-emit polarized emission aligned with their long axis.

not capture the full complexity of the spectral emission properties of polarized thermal dust. Planck has found evidence of spatial variations of the dust spectral indices [374, 52]. These variations can lead to decorrelation between different frequencies. A typical model of the thermal dust B-mode cross power spectra between two observing frequencies ν_1 and ν_2 can be written as [52]

$$\ell(\ell + 1)C_{\ell}^{BB}(\nu_1, \nu_2) \propto (\nu_1 \nu_2)^{\beta_d - 2} B_{\nu_1} B_{\nu_2} f_d(\delta_d, \nu_1, \nu_2), \quad (5.3)$$

where δ_d is the *decorrelation parameter* and f_d is defined as

$$f_d(\delta_d, \nu_1, \nu_2) \equiv \exp(-\delta_d (\ln(\nu_1/\nu_2))^2). \quad (5.4)$$

However, Planck data shows no evidence for dust decorrelation [52, 413] and the derived upper-limits imply that there should be no impact of decorrelation on primordial B-mode measurements for $r > 0.01$ [52]. More involved modelling of thermal dust also takes into account multiple dust components with distinct properties, e.g. each dust grain type with its own spectral index, temperature and polarization angle [241]. Errors in the modelling of foregrounds give rise to biases in the component separation, therefore efforts to improve the foreground models are essential to the ongoing quest for primordial B-modes, particularly for the sensitivities targeted by the forthcoming CMB experiments.

Synchrotron radiation

Synchrotron radiation arising from charged particles—electrons in particular—spiraling around the magnetic field lines of the galaxy is another

source of foreground contamination. Contrary to thermal dust, synchrotron emission dominates at lower frequencies as shown in figure 5.16. The frequency scaling is modelled by a power-law

$$A^{\text{sync}}(\nu) \propto \nu_s^\beta, \quad (5.5)$$

where β_s is the synchrotron spectral index, measured in the frequency range 2.3 and 33 GHz to be around $\beta_s \approx -3.2$, combining data from the S-PASS survey, Planck and WMAP [300]. Improved models of synchrotron radiation need to take into account the spectral curvature β_s , i.e. its variation with frequency [288], in addition to the spatial variation across the sky as is the case of thermal dust [289].

5.2.3 Delensing

If the measurement of the lensing signal is a primary science target of current and future CMB experiments, since it provides a wealth of information on cosmology as we have discussed in Chapter 3, it is also a major contaminant at large scales, overshadowing the primordial B-mode signal. The lensing induced E to B mixing manifests as white noise at the $5 \mu\text{K-arcmin}$ level (see figure 5.18), therefore experiments which aim at reaching an instrumental noise level below this figure are going to be dominated by lensing. It is thus crucial, for these experiments, to accurately characterize the lensing signal at small scales and perform a delensing procedure [285, 281, 410, 420, 414, 129, 332] which can help reduce the amplitude of lensing by large factors enabling higher sensitivity in the measurement of primordial B-modes as shown in figure 5.18.

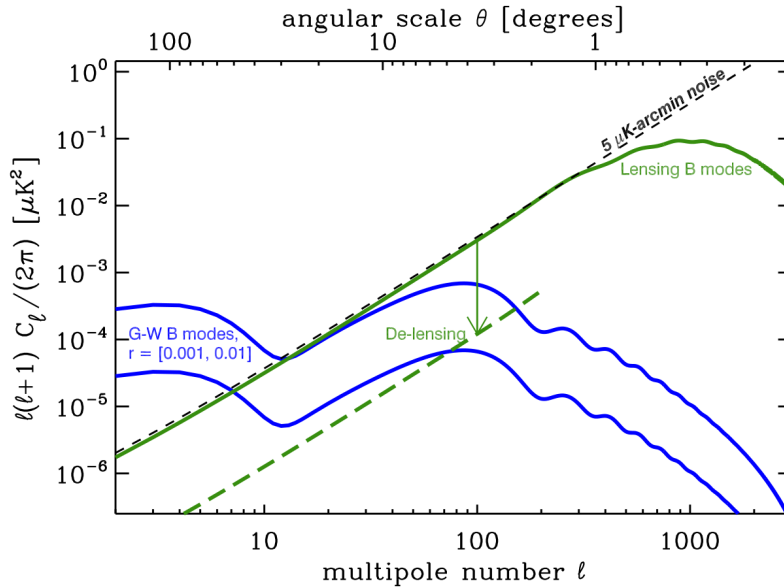


Figure 5.18 – The green curve shows the lensing induced B-modes which coincides on large scales with the gray dashed line representing $5 \mu\text{K-arcmin}$ white noise. Delensing can significantly suppress this effect as shown in the green dashed line, resulting in lower effective noise in the B-mode map. *Credit: CMB-S4 Science Book [16].*

The lensing field itself can be constructed using the quadratic estimator [254]. This estimator capitalizes on the fact that lensing induces correlations between different Fourier modes, and reconstructs the deflection field out of the weighted sum over pairs of CMB modes in harmonic space. Similarly the lensing field can also be constructed from external data, using tracers of the large scale structure such as the CIB, which is strongly correlated with CMB lensing. The reconstruction of the lensing potential is highly dependent on high-resolution and high-fidelity measurements of the polarization maps, as information on degree scale lenses can be contained in the E and B fields down to arcminute scales ($\ell \sim 2000$) [16].

Delensing can also improve the reconstruction of the lensing field itself allowing higher sensitivity measurements of cosmological parameters to which the lensing is sensitive, such as neutrino mass. This is the case for low noise experiments where the lens reconstruction from the EB pair can be improved after delensing the B-mode signal. This leads to a better estimate of the lensing field, which in turn can be used to improve the delensing of B-modes, giving rise to an *iterative delensing* procedure which can drastically improve the estimation of the lensing and cleaned B-mode signals [410, 129, 28].

5.3 Data analysis pipelines for modern CMB experiments

Extracting cosmological information from CMB observations requires processing the raw data collected by the detectors through multiple steps, where the data is increasingly compressed until we end up with a handful of cosmological parameters. These steps put together constitute what is referred to as *the data analysis pipeline*. In this section, we give an overview of the main procedures conducted in each step of the analysis. The technical details may change for different experiments according to their design and adopted methodologies, but the objective of each step remains the same.

Pre-processing

Before feeding the time-ordered data to the reduction and science analysis pipelines, a number of calibration tasks need to be conducted. In this context, observations of planets and radio sources are typically used for pointing reconstruction as well as for beam calibration. Polarized sources such as the Crab nebula remnant, Tau A [69], are used to calibrate the polarization angles, while the relative detector gains can be calibrated using chopped thermal sources or short scans in elevation referred to as “el-nods”. Comparisons with other CMB measurements can be used for absolute calibration. Absolute gain calibration can also be achieved with measurement of the CMB dipole, but this requires large sky coverage and is typically only accessible for balloon-borne or space-based experiments.

When the instrument is properly calibrated, the time-ordered data need to be checked for quality and cleaned before one can reconstruct the sky maps. This includes flagging bad quality data, such as corresponding to bad weather condi-

tions, or abnormally noisy timestream segments, in addition to glitches that may be caused by cosmic ray hits for instance. Data from the turnarounds are also typically discarded since the scan speed is not constant inducing irregular sampling of the sky, in addition to featuring a high level of microphonics, i.e. spurious signals caused by mechanical vibrations of the telescope.

Map-making

After the pre-processing of the time-ordered data, estimates of the sky signal at each observing frequency can be reconstructed. This represents a significant compression of the size of the data set: for current and forthcoming CMB experiments, the total size of raw time-domain data is expected to reach $\mathcal{O}(10^{12-14})$ time samples, this is typically compressed to sky maps of $\mathcal{O}(10^{5-8})$ pixels. Ideally, this compression should be performed with no loss of cosmological information. This is an important step in the analysis, where special care needs to be taken to accurately model the noise and the different systematics in the data, as any uncertainties or biases incurred at this step are subsequently propagated throughout the data analysis pipeline stages.

Formally, map-making is typically a linear operation mapping the time-domain measurements, \mathbf{d} , to a map in the pixel domain, \mathbf{m} , with some operator \mathbf{L} [445]

$$\mathbf{m} = \mathbf{L}\mathbf{d}. \quad (5.6)$$

We assume a data model of the following form

$$\mathbf{d} = \mathbf{P}\mathbf{s} + \mathbf{n}, \quad (5.7)$$

where \mathbf{s} is the true pixelized sky, scanned with some pointing matrix \mathbf{P} , which encodes information on the pixels observed and the orientation of the instrument with respect to the sky coordinates at each time. The noise timestream, \mathbf{n} , is assumed to vanish on average. Under these assumptions, requiring the map estimator to be unbiased leads to a constraint on the linear operator \mathbf{L}

$$\mathbf{L}\mathbf{P} = \mathbf{I}. \quad (5.8)$$

A general form of unbiased estimators as solutions to the above equation can thus be written as

$$\mathbf{L} = (\mathbf{P}^T \mathbf{W}^{-1} \mathbf{P})^{-1} \mathbf{P}^T \mathbf{W}^{-1}, \quad (5.9)$$

where the choice of the positive definite weight matrix \mathbf{W} , although always leading to an unbiased estimator, determines the noise properties in the pixel domain.

Binning The simplest choice of the weight matrix is $\mathbf{W} = \mathbf{I}$. In this case the map estimator has the form, $\mathbf{m} = (\mathbf{P}^T \mathbf{P})^{-1} \mathbf{P} \mathbf{d}$, corresponding to simple averaging of the time samples falling within each pixel. While this is an unbiased estimator which, moreover, can be easily computed, it leads to a complex noise structure in pixel domain with stripes along the scan direction induced by the noise correlations.

Maximum likelihood Under the assumption of Gaussian noise with noise covariance, \mathbf{N} , an optimal choice of the weight matrix is $\mathbf{W} = \mathbf{N}^{-1}$. This solution

corresponds to the maximum likelihood and minimum variance estimator. The main challenge here is that the noise covariance is not a priori known, therefore one needs to estimate it from the data assuming some sufficiently accurate noise model. Besides, even with additional assumptions to make the solution computationally tractable, it remains expensive to compute.

Templates marginalization Another possible choice of the weight matrix is to build it from a set of predefined orthogonal temporal modes spanning the columns of a matrix \mathbf{T} . The solution $\mathbf{W} = \mathbf{I} - \mathbf{T}\mathbf{T}^T$ is then effectively filtering all the temporal modes of \mathbf{T} concurrently with the sky map estimation. One can then define specific models of different time-domain systematics and deproject them from the data while constructing an unbiased sky estimate.

We will review all these methods more thoroughly in the next chapter. Another possibility is to compute a biased estimator of the sky, generally of the form

$$\mathbf{m} = (\mathbf{P}^T \text{diag}(\mathbf{W}^{-1}) \mathbf{P})^{-1} \mathbf{P}^T \mathbf{W}^{-1} \mathbf{d}. \quad (5.10)$$

The bias is then corrected for later in the data analysis stages, typically at the power spectrum level [245]. This solution is referred to as *filter and bin* map-making. It has the significant advantage of being cheap to compute and has been successfully applied in many experiments [150, 406, 40, 354].

The pixel domain noise covariance is analytically given by

$$\mathbf{N}_p = \mathbf{L}\mathbf{N}\mathbf{L}^T, \quad (5.11)$$

however given that this matrix is generally dense, its computation is prohibitive except in cases of very small sky coverage or low resolutions leading to a limited number of pixels. In practice, the uncertainties propagated from map-making to the subsequent analysis stages are estimated via Monte Carlo simulations.

Component separation

As we have seen before, the presence of foregrounds represents a significant challenge to the search for primordial gravitational waves. In addition, they are also expected to have an impact on the quality of the lensing estimation [190]. As a consequence, separating the foreground components from the CMB signal is a crucial step in the analysis for both large scale and small scale science.

The output of map-making is a set of sky maps at each observing frequency, each representing a different linear combination of distinct astrophysical signals. We can thus write the underlying data model as the following

$$\mathbf{m} = \mathbf{A}\mathbf{s} + \mathbf{n}, \quad (5.12)$$

where \mathbf{m} is the reconstructed sky map in each observing frequency, \mathbf{A} is the *mixing matrix* which determines how the different astrophysical components combine together at each frequency band to form the measured sky signal, \mathbf{s} contains the amplitudes of the different components and \mathbf{n} is the pixel-domain noise. The aim of component separation methods is to invert the above equation, recovering an estimate of the cleaned CMB signal as well as foreground maps,

without full knowledge of the matrix \mathbf{A} . Generally, component separation methods can be classified into two main approaches. The first approach, referred to as *parametric methods*, defines parametric models of the foregrounds and fit for them before inverting Eq. (5.12). The second approach, so-called *blind methods*, makes minimal assumptions about the foregrounds and extracts the CMB by exploiting its spectral properties or its statistical independence from the foreground signals.

Parametric methods Assuming some spectral laws of the foregrounds, such as the ones we have outlined in the previous section, the mixing matrix can be parametrized by a set of spectral parameters $\mathbf{A} \equiv \mathbf{A}(\beta)$. One can then fit for these spectral parameters, typically using data in the pixel-domain. This is achieved by maximizing the so-called *spectral likelihood* [107, 181, 430] which reads

$$\mathcal{S}(\beta) = -(\mathbf{A}^T \mathbf{N}^{-1} \mathbf{m})^T (\mathbf{A}^T \mathbf{N}^{-1} \mathbf{A})^{-1} (\mathbf{A}^T \mathbf{N}^{-1} \mathbf{m}). \quad (5.13)$$

The obtained mixing matrix is then used to invert Eq. (5.12) using a generalized least square estimator. The primary drawback of these methods is the need to ensure accurate modelling of the foregrounds, which is a challenging astrophysical problem, as any mismatches between the true mixing matrix and the recovered estimate biases the procedure, leading to the presence of foreground residuals in the component maps.

Blind methods The only spectral law that is assumed to be known in blind methods is the blackbody spectrum of the CMB. A well known class of these approaches is the so-called *Internal Linear Combination* (ILC) methods [90, 447, 161, 261, 398]. The basic idea of ILC is to construct a linear combination of the frequency maps that preserves the CMB signal and minimizes the variance of the output. We can write the estimator of the cleaned CMB map as the following

$$\hat{\mathbf{s}}^{\text{cmb}} = \sum_{i=1}^{n_f} w_i \mathbf{m}(\nu_i) = \mathbf{s}^{\text{cmb}} + \sum_{i=1}^{n_f} w_i (\mathbf{s}^{\text{fg}}(\nu_i) + \mathbf{n}(\nu_i)), \quad (5.14)$$

where the weights w_i satisfy $\sum_{i=1}^{n_f} a_i w_i = 1$ with a_i corresponding to the CMB mixing elements. The minimum variance solution corresponds to the weights

$$\mathbf{w} = (\mathbf{a}^T \mathbf{C}^{-1} \mathbf{a})^{-1} \mathbf{C}^{-1} \mathbf{a}, \quad (5.15)$$

where \mathbf{C} is the data covariance matrix. One weakness of this method is that it produces biased maps due to the presence of empirical correlations between the CMB and the foregrounds. The second class of blind methods is referred to as *Independent Component Analysis* (ICA) [72, 325, 125]. The idea here is to use statistical independence of the components and the fact that foregrounds have non-Gaussian distributions. A linear transformation of the data from different frequency channel is constructed

$$\mathbf{y} = \mathbf{W} \mathbf{m}. \quad (5.16)$$

According to the central limit theorem, any combination of statistically independent random variables is more-Gaussian, therefore by choosing the linear

operator \mathbf{W} such that we maximize some measure of non-Gaussianity, we also maximize the independence of the components of \mathbf{y} and hence we are able to reconstruct estimates of the original independent components combining into the sky signal. Finally a third class of blind or semi-blind methods is the so-called *Template fitting* approach, where the low and high frequency channels are used to build templates of thermal dust and synchrotron, the data is then used to fit for and subtract these templates to recover the cleaned CMB maps.

We also note that on top of the usual map-domain representations, these methods also allow to manipulate the data in other domains. In fact, variants exist in harmonic space, as well as wavelet and spin-wavelet domains.

Angular power spectrum estimation

Once the cleaned CMB maps have been obtained, the angular power spectra of the 6 pairs of fields: TT, TE, TB, EE, EB, and BB can be derived. In this stage of the analysis several effects need to be corrected for, including partial sky coverage, beam smoothing effects as well as possible asymmetries not accounted for in the map-making, in addition to biases introduced by the filtering. Three general approaches can be identified:

Maximum likelihood methods This approach relies on maximizing the full likelihood and has been successfully applied since the early CMB measurements [101]. If one writes the reconstructed sky map as the sum of the CMB signal and some noise contribution, $\mathbf{m} = \mathbf{s} + \mathbf{n}$, then assuming that the signal and the noise are uncorrelated and both following a Gaussian distribution, one can write the likelihood as the following

$$P(\mathbf{m}|\mathbf{C}_\ell) \propto \exp \left(-\frac{1}{2} \mathbf{m}^T (\mathbf{S} + \mathbf{N})^{-1} \mathbf{m} - \text{Tr}(\ln(\mathbf{S} + \mathbf{N})) \right), \quad (5.17)$$

where \mathbf{N} is the pixel-domain noise covariance, and \mathbf{S} is the signal covariance which has a simple form for temperature

$$S_{pp'} = \sum_{\ell} \frac{2\ell + 1}{4\pi} B_\ell^2 C_\ell^{TT} P_\ell(\cos \theta_{pp'}), \quad (5.18)$$

with B_ℓ the window function characterizing the beam and pixelization effects, P_ℓ Legendre polynomials and $\theta_{pp'}$ the angle between the two pixels p and p' . The general form involving polarization and cross-spectra can be found in [446]. Computing the power spectra estimates and determining their error bars requires the evaluation of the likelihood or a quadratic approximate which is very expensive, scaling as $\mathcal{O}(\mathcal{N}_{\text{pix}}^3)$, as it requires the inversion of the data covariance for each evaluation [103]. As a consequence, this method is mainly effective for low resolution cases targeting low multipoles [174].

Sampling methods These methods arose as an alternative to the expensive maximum likelihood approaches. They allow to iteratively sample the C_ℓ values from the posterior $P(C_\ell|\mathbf{m})$ typically using Markov Chain Monte Carlo (MCMC)

methods [264, 182]. One well known variant, based on *Gibbs sampling*, has been successfully applied to the three generations of CMB satellites, COBE [463], WMAP [352] and Planck [43]. The idea is to draw samples from the conditional distributions $P(\mathbf{s}|C_\ell, \mathbf{m})$ and $P(C_\ell|\mathbf{s})$ following

$$\begin{aligned} \mathbf{s}^{i+1} &\leftarrow P(\mathbf{s}|C_\ell^i, \mathbf{m}) \\ C_\ell^{i+1} &\leftarrow P(C_\ell|\mathbf{s}^{i+1}) \end{aligned}$$

the procedure converges to drawing samples from the joint density $P(C_\ell, \mathbf{s}|\mathbf{m})$, from which we derive $P(C_\ell|\mathbf{m})$ by marginalizing over \mathbf{s} . The computations are dominated by spherical harmonic transforms and thus scale as $\mathcal{O}(\mathcal{N}_{\text{pix}}^{3/2})$. One major challenge of sampling methods is that their performance may suffer from high dimensionality, which may limit their effectiveness in handling the large data sets of the near-term and long-term future CMB observatories.

Pseudo- C_ℓ methods In this approach [462, 245, 229, 133, 110], biased spectra (referred to as pseudo-spectra) are directly computed from spherical harmonic transforms of the maps. Formally, one defines a pseudo basis ${}_s\tilde{Y}_{\ell m}$ typically built from the basis of spin-weighted spherical harmonics, ${}_sY_{\ell m}$ ($s = 0$ for temperature and $s = \pm 2$ for polarization) as

$${}_s\tilde{Y}_{\ell m} = W {}_sY_{\ell m}, \quad (5.19)$$

where W is a *window function* defining the partially observed region of the sky and smoothed with some apodization function to avoid ringing effects. Pseudo-multipoles are then computed from this basis as

$${}_s\tilde{a}_{\ell m} = \int {}_s\tilde{Y}_{\ell m}^* f(\mathbf{n}) d\Omega, \quad (5.20)$$

where $f(\mathbf{n})$ corresponds to either the temperature or the polarization fields depending on s . Using the same E-B decomposition we have shown in Chapter 3, we can then define the pseudo power spectra for a pair of fields $X, Y \in \{T, E, B\}$ as

$$\tilde{C}_\ell^{XY} \equiv \sum_{m=-\ell}^{\ell} \frac{{}_X\tilde{a}_{\ell m}^* {}_Y\tilde{a}_{\ell m}}{2\ell + 1}. \quad (5.21)$$

Finally the pseudo-spectra can be related to the full sky angular power spectrum C_ℓ through their ensemble average via

$$\langle \tilde{C}_\ell \rangle = \sum_{\ell'} M_{\ell\ell'} F_{\ell'} B_{\ell'}^2 \langle C_{\ell'} \rangle + \langle \tilde{N}_\ell \rangle. \quad (5.22)$$

where $M_{\ell\ell'}$ is a mixing matrix, depending only on the geometry of the window function, describing the mode coupling induced by the cut sky. F_ℓ is the transfer function describing the effects of filtering, that corrects for the bias introduced in the context of filter and bin approaches, while it should be ideally close to unity for unbiased maps. These transfer functions along with the average power spectrum of the noise bias $\langle \tilde{N}_\ell \rangle$ are estimated using Monte-Carlo simulations. Once these are estimated, one can invert the above linear system to get an estimate of

the C_ℓ values along with their error bars. Pseudo- C_ℓ estimators are fast to compute, dominated by the cost of the spherical harmonic transforms $\mathcal{O}(\mathcal{N}_{\text{pix}}^{3/2})$, however they require a careful characterization of the instrument and the noise which is a resource consuming process demanding a large number of Monte-Carlo simulations to be performed.

Cosmological parameters estimation

The final step in the data analysis consists in the estimation of the cosmological parameters of underlying cosmological models, along with their confidence intervals, as well as conducting goodness-of-fit tests to assess whether the best-fit model is actually a good description of the data. One should also note that it is possible to include external data sets and perform a joint analysis which generally yields more powerful constraints as we have seen in Chapter 4.

In general, the power spectra estimated in the previous step of the analysis are binned in ℓ -space. The sizes of the bins $\Delta\ell$ (also referred to as bandpowers) are chosen such that the levels of noise and foreground contaminations do not vary significantly over the corresponding ℓ -segments, in addition in the pseudo- C_ℓ framework they also need to be sufficiently large to minimize the mode coupling effects and allow the inversion of the mode mixing matrix. For each bandpower, a $n_{\text{freq}} \times n_{\text{freq}}$ covariance matrix, $\hat{\mathbf{C}}_b$, can be constructed from all the estimated auto and cross spectra in the n_{freq} observing frequencies. Assuming underlying isotropic and Gaussian fields as well as full-sky measurements, $\hat{\mathbf{C}}_b$ follows the so-called *Wishart distribution* with $\sum_\ell \mathbf{P}_{b\ell}(2\ell+1)$ degrees of freedom per band power, where $\mathbf{P}_{b\ell}$ is the binning operator which defines the band power window functions. For partial sky measurements, one can assume that $\hat{\mathbf{C}}_b$ follows the same distribution but with an effective number of degrees of freedom per bandpower, ν_b given by the fractional uncertainty of the auto spectra. The likelihood of the data covariance $\hat{\mathbf{C}}_b$ can then be written assuming some underlying model covariance, \mathbf{C}_b , which should account not only for the cosmological model, but also for noise, foreground residuals, and other sources of systematics

$$-2 \ln \mathcal{L} = \sum_b \nu_b \left[\text{Tr}(\hat{\mathbf{C}}_b \mathbf{C}_b^{-1}) - \ln \det(\hat{\mathbf{C}}_b \mathbf{C}_b^{-1}) \right] \quad (\text{up to a constant}). \quad (5.23)$$

One can then minimize the above log-likelihood to find the best-fit parameters, or marginalize over nuisance parameters as well as determine confidence intervals. Given the complexity of the data sets, these procedures can be computationally heavy. One possible approach, is to discretize the parameter space, these are the so-called *gridding methods*, however the performance of these methods can quickly deteriorate with a high number of parameters as they scale as $\propto \mathcal{N}_{\text{params}}^2$. A faster alternative, which has become the standard approach in many CMB experiments, is to draw representative samples of the likelihood using MCMC methods [144, 312]. However, as we have previously mentioned these methods also suffer from convergence issues when handling large data sets.

In this chapter, we briefly went over the experimental efforts that have been deployed to measure CMB observables and extract key scientific insights that

inform models of cosmology as well as their connections with other areas of fundamental physics. I gave a more detailed review of two specific experiments that I have been involved with in my thesis, before conducting a general discussion of the scientific exploitation of CMB measurements and their associated challenges.

In my thesis work, I have specifically focused on developing and implementing methods suitable for the map-making stage of the analysis. This work is exposed in the next three chapters as the following: in Chapter 6 I give a detailed review of the map-making problem, and provide a general approach for reconstructing sky maps while deprojecting several time-domain systematics that may be present in the data. I particularly focus on a specific set of systematics that are relevant for the forthcoming generation of CMB experiments such as the Simons Observatory. In Chapter 7, I present a general framework that adopts the methodology discussed in Chapter 6 and implements different linear solvers techniques in a massively-parallel software, I discuss aspects of the implementation and evaluate its computational efficiency. Finally, in the last chapter I apply these techniques to different simulations of satellite and ground-based experiments for purposes of validation and to build an understanding of the sensitivity to different map-making parameters. In this context, I also investigate the effects of atmosphere on the polarization maps of the Simons Observatory Small Aperture Telescopes, using different map-making approaches, and assess the impact of some of the intensity-to-polarization leakage effects such as caused by gain mismatch.

CHAPTER 6

Unbiased map-making in the presence of time-domain systematics

Contents

6.1	The map-making problem	128
6.1.1	Data model	128
6.1.2	Map-making solution	129
6.1.3	Pointing matrix	131
6.2	Maximum likelihood map-making	133
6.3	Templates marginalization map-making	135
6.3.1	Atmosphere and $1/f$ instrumental noise	136
6.3.2	Scan synchronous signal	136
6.3.3	Half-wave plate synchronous signal	137
6.4	Degeneracy issues	138
6.4.1	Scanning strategy sourced degeneracies	138
6.4.2	Templates linear dependencies	139
6.4.3	Templates and sky degeneracies	139

In this chapter, I present a mathematical formalism for unbiased reconstructions of sky maps, accommodating not only projections of time-domain data in map-domain but also deprojections of time-domain systematics insofar as they can be modeled as linear contributions to the time-ordered data. In the first section, I give a detailed presentation of the map-making problem, considering a generic data model of CMB experiments, from which I derive a general expression for unbiased sky maps estimators. The next two sections are dedicated to a detailed discussion of two main approaches to estimating unbiased maps which employ different models of the noise and are adapted to deal with selected time-domain systematics. In the last section, I briefly present the different sorts of

degeneracies encountered in the map-making problem, along with prescriptions on how they can be handled.

The methodology adopted in this chapter, as well as the different models of time-domain systematics I discuss, are of interest for current and forthcoming CMB experiments, and are implemented in a map-making software framework I developed in this thesis. This framework is the object of a prepared publication: Hamza El Bouhargani et al. *MAPPRaiser: A massively parallel map-making framework for multi-kilo pixel CMB experiments* [180], essentially covering the content of this chapter and the next, as well as significant parts of the last one. As such, large parts of these chapters are verbatim excerpts from this paper.

6.1 The map-making problem

6.1.1 Data model

As explained in Chapter 5, the raw data of the detectors are properly calibrated and cleaned in the pre-processing stage, before they can be handed over to the map-makers. The data are taken at a fixed time step, and for each detector, ordered by the time at which they were collected. The data from different detectors are concatenated together in a single data vector. We denote this data vector as \mathbf{d} , and use a linear model to describe the data representing them therefore as,

$$\mathbf{d} = \sum_i \mathbf{T}_i \mathbf{x}_i + \mathbf{n}. \quad (6.1)$$

Here, the index i goes over different contributions, whatever their origins, and the columns of each matrix \mathbf{T}_i define a time-domain template describing the effect. \mathbf{n} denotes then a stochastic contribution, which is characterized only statistically and assumed for most of the application to be piece-wise time-stationary. We refer to this term generically as noise.

A key assumption behind the model is that all relevant contributions can be discretized with sufficient precision and described by a limited number of amplitudes stored in the vectors \mathbf{x}_i , and the corresponding templates, i.e. columns of the matrices \mathbf{T}_i , which are assumed to be known.

For definiteness we assume that $i = 0$ term describes the actual sky signal. The corresponding template matrix, \mathbf{T}_0 , is then typically referred to as a pointing matrix and denoted as \mathbf{P} . The vector \mathbf{x}_0 stands then for a pixelized sky signal, typically a sky map, and we denote it as \mathbf{m} , ($= \mathbf{x}_0$). On concatenating all the remaining template matrices, \mathbf{T}_i , and the amplitude vectors, \mathbf{x}_i , together, and naming the combined objects respectively as \mathbf{T} and \mathbf{x} , we can rewrite Eq. (6.1) as,

$$\mathbf{d} = \mathbf{P} \mathbf{m} + \mathbf{T} \mathbf{x} + \mathbf{n} \equiv \mathcal{P} \mathbf{y} + \mathbf{n}. \quad (6.2)$$

Here, for future convenience we have introduced a generalized pointing matrix, \mathcal{P} , and a generalized map, \mathbf{y} , which combine all the template matrices and amplitude vectors, respectively, including those describing the sky signal.

In the context of modern experiments with multiple thousands of detectors, the size of the time ordered data can reach $\mathcal{N}_t = \mathcal{O}(10^{12-15})$, while the number of the amplitudes is typically much smaller. In particular, the number of sky

amplitudes is usually on order of $\mathcal{N}_{\text{pix}} = \mathcal{O}(10^{5-8})$. The data are therefore hugely redundant reflecting the fact that features of the actual sky signal of the current interest are very subdominant as compared to the noise of the measurements.

6.1.2 Map-making solution

The map-making problem consists of estimating the sky signal, \mathbf{m} , by compressing the redundant information contained in the raw data and deprojecting the undesirable effects defined by the combined template matrix \mathbf{T} . This is done given some assumptions about the statistical properties of the noise, \mathbf{n} . The most general unbiased solution of the generalized map-making problem defined in Eq. (6.2), compressing the data with a weight matrix, \mathbf{W} , can be written as (see the map-making section in Chapter 5)

$$\hat{\mathbf{y}} = (\mathcal{P}^T \mathbf{W}^{-1} \mathcal{P})^{-1} \mathcal{P}^T \mathbf{W}^{-1} \mathbf{d}, \quad (6.3)$$

where the generalized map $\hat{\mathbf{y}}$ containing both sky signal and templates amplitudes can then be truncated to retain only the sky signal part [429]. Using block-wise inversion (Cf. Appendix A for more details), we can derive a concise form for the latter, that is the most general form for a map estimator fulfilling the above conditions [429, 378],

$$\hat{\mathbf{m}} = (\mathbf{P}^T \mathbf{F}_T \mathbf{P})^{-1} \mathbf{P}^T \mathbf{F}_T \mathbf{d}, \quad (6.4)$$

where

$$\mathbf{F}_T \equiv \mathbf{W}^{-1} \left(\mathbf{I} - \mathbf{T} (\mathbf{T}^T \mathbf{W}^{-1} \mathbf{T})^{-1} \mathbf{T}^T \mathbf{W}^{-1} \right) \quad (6.5)$$

is a filtering and weighting operator, which simultaneously filters all the contributions contained in the subspace spanned by the columns of the template matrix, \mathbf{T} , i.e. $\mathbf{F}_T \mathbf{T} = 0$, and weights modes orthogonal to it by the weight matrix, \mathbf{W} . Consequently, the estimator given by Eq. (6.4) is explicitly unbiased as long as the system matrix, $\mathbf{P}^T \mathbf{F}_T \mathbf{P}$, is non-singular and the assumed templates span the subspace affected by the unwanted signals. Indeed,

$$\begin{aligned} \hat{\mathbf{m}} &= (\mathbf{P}^T \mathbf{F}_T \mathbf{P})^{-1} \mathbf{P}^T \mathbf{F}_T (\mathbf{P} \mathbf{m} + \mathbf{T} \mathbf{x} + \mathbf{n}), \\ &= \mathbf{m} + (\mathbf{P}^T \mathbf{F}_T \mathbf{P})^{-1} \mathbf{P}^T \mathbf{F}_T \mathbf{n}, \end{aligned} \quad (6.6)$$

and the last term vanishes when averaged over the statistical ensemble of noise realizations as $\langle \mathbf{n} \rangle_{\text{noise}} = 0$.

We note for completeness that if some prior information concerning the sought-after amplitudes of the templates is available it can be incorporated in the map estimation modifying the filtering operator, which in such a case reads,

$$\tilde{\mathbf{F}}_T = \mathbf{W}^{-1} \left(\mathbf{I} - \mathbf{T} (\Theta^{-1} + \mathbf{T}^T \mathbf{W}^{-1} \mathbf{T})^{-1} \mathbf{T}^T \mathbf{W}^{-1} \right), \quad (6.7)$$

where Θ is the covariance matrix of the sought after amplitudes, \mathbf{x} . This matrix may be rank deficient if the prior information is available only for some of the amplitudes and not others. No prior information at all corresponds to $\Theta \rightarrow \infty$

and $\Theta^{-1} \rightarrow 0$ and we recover Eq. (6.5). Note that the templates for which the prior is available are not completely filtered out from the data during each map-making procedure, i.e., $\tilde{\mathbf{F}}_T \mathbf{T} \neq 0$, rather they are Wiener-filtered instead. The sky signal estimates are therefore not unbiased when averaged over the noise ensemble, however, their overall uncertainty over the ensemble of the noise and template amplitudes is minimized given the assumed priors [112].

Alternately, one can opt for a two step approach (also explicitly derived in Appendix A) where the amplitudes of the undesirable signals, \mathbf{x} , are first estimated as,

$$\hat{\mathbf{x}} = (\mathbf{T}^T \mathbf{F}_P \mathbf{T})^{-1} \mathbf{T}^T \mathbf{F}_P \mathbf{d}, \quad (6.8)$$

and the sky map is then recovered on the second step via,

$$\hat{\mathbf{m}} = (\mathbf{P}^T \mathbf{W}^{-1} \mathbf{P})^{-1} \mathbf{P}^T \mathbf{W}^{-1} (\mathbf{d} - \mathbf{T} \hat{\mathbf{x}}). \quad (6.9)$$

This approach is then essentially equivalent to the destriping technique originally proposed for Planck [113, 162, 324], but later also proposed in the context of ground-based experiments [434, 435, 378].

Both these approaches have been successfully employed in the past. In the realm of the multi-kilo pixel experiments, the number of instrumental, detector-specific contributions in the data is typically exceeding by far the number of the observed sky pixels. Thus, the system matrices in Eq. (6.3) and (6.8) are significantly larger than the one appearing explicitly in Eq. (6.4). Consequently, solving Eq. (6.4) may look as potentially more computationally efficient. For this to be indeed the case, there must be a computationally efficient way to construct the template orthogonalization kernel $(\mathbf{T}^T \mathbf{W}^{-1} \mathbf{T})^{-1}$, which in general can be very computationally cumbersome. Fortunately, in many relevant applications the kernel is usually very structured and sparse, thanks to the fact that most of the templates of interest have compact and disjoint support, which is moreover known a priori. This can be capitalized on to save on the calculations making this approach not only computationally feasible but also preferred in many applications. I therefore implement this approach in the framework I propose in the next chapter.

Eq. (6.4) generalizes and encapsulates as special cases many of the standard map-making techniques, which correspond to a specific choice of either the weight matrix or the template matrix or both. The specific cases include: (1) the standard maximum likelihood (minimum variance) map-making, where the templates are often dropped and the weights are taken to be given by the inverse covariance of the noise term, \mathbf{n} ; (2) the binned map-making where the weights are taken to be diagonal and equal to the noise rms and the templates neglected; (3) the standard destriper, where the templates are taken to be piece-wise constant offsets and the weights diagonal and set by the noise rms; (4) the generalized destriper, where the templates can be more general and the covariance of their amplitudes could be included as prior information.

Eq. (6.4) is also very flexible allowing for diverse choices for the pointing, the template, and the weight matrices with the only practical constraint being numerical considerations related to their application. As I discuss in the following section this permits applications of this equation in many cases of practical interest.

6.1.3 Pointing matrix

The sky signal measured by an instrument is a convolution of the actual sky signal, expressed with the three Stokes parameters I, Q and U , with an instrument response matrix, $\mathcal{R}(\delta\gamma, \psi, \nu)$. This matrix will in general depend on both the sky angle, $\delta\gamma$, measured with respect to the observation direction, γ , the orientation of the instrument with respect to the sky coordinates, ψ , and the frequency, ν . Therefore, a measurement made in ν_c frequency band, defined by a bandpass, $\mathcal{W}(\nu, \nu_c)$, can be written in general as,

$$\mathbf{d}_{\nu_c}(t) = \int d\nu \int d\delta\gamma \mathcal{W}(\nu, \nu_c) \mathbf{e}_i^T \mathcal{R}(\delta\gamma, \psi(t), \nu) \times \mathbf{s}(\gamma(t) + \delta\gamma, \nu) + \mathbf{n}(t), \quad (6.10)$$

where $\gamma(t)$ is an observation direction at time t and \mathbf{e}_i is a unit vector which defines whether we ultimately measure total power only, $i = 0$, or Stokes Q , $i = 1$, or U , $i = 2$, parameter. $\mathbf{s}(\gamma, \nu)$ is a vector of the three Stokes parameters as defined in the direction $\gamma(t)$ and at frequency ν .

It is frequently possible to represent the general instrument response matrix as a product of a Mueller matrix describing the detection chain of the instrument and a matrix defining the beam effects,

$$\mathcal{R}(\delta\gamma, \psi(t), \nu) \equiv \mathcal{M}(\boldsymbol{\iota}(t), \nu) \mathcal{B}(\delta\gamma, \nu) \mathbf{R}(\psi(t)), \quad (6.11)$$

where the matrices, \mathcal{M} and \mathcal{B} , are defined with respect to instrument coordinates and the matrix $\mathbf{R}(\psi(t))$, denotes a rotation between the sky and instrument coordinates. We note that the Mueller matrix, \mathcal{M} , may in general depend on time via some of the instrument parameters, such as, phase of a polarization modulator, as denoted by parameter $\boldsymbol{\iota}(t)$.

Assuming that the beam matrix is diagonal, axially symmetric, i.e. depending only on the magnitude of the angle, $\delta\gamma$, and not on its orientation, and that the beams for the Stokes parameters Q and U are identical, we can commute the rotation operator and the beam matrix, rewriting Eq. (6.10), as,

$$\mathbf{d}_{\nu_c}(t) = \int d\nu \mathcal{W}(\nu, \nu_c) \mathbf{e}_i^T \mathcal{M}(\boldsymbol{\iota}(t), \nu) \mathbf{R}(\psi(t)) \times \underbrace{\int d\delta\gamma \mathcal{B}(\delta\gamma, \nu) \mathbf{s}(\gamma(t) + \delta\gamma, \nu)}_{\equiv \bar{\mathbf{s}}(\gamma(t), \nu)} + \mathbf{n}(t), \quad (6.12)$$

where $\bar{\mathbf{s}}(\gamma(t), \nu)$ stands for a beam-smoothed sky signal, which can be therefore characterized by a finite set of pixelized sky signal amplitudes, $\bar{\mathbf{s}}(p(t), \nu)$, where $p(t)$ is a sky pixel observed at time t .

Furthermore, if the bands defined by $\mathcal{W}(\nu, \nu_c)$ are sufficiently narrow so that the beam-convolved sky signal is essentially constant across the band, we can write,

$$\begin{aligned} \mathbf{d}_{\nu_c}(t) &\approx \left[\int d\nu \mathcal{W}(\nu, \nu_c) \mathbf{e}_i^T \mathcal{M}(\boldsymbol{\iota}(t), \nu) \mathbf{R}(\psi(t)) \right] \\ &\quad \times \bar{\mathbf{s}}(p(t), \nu_c) + \mathbf{n}(t), \\ &\equiv \mathbf{w}(t) \bar{\mathbf{s}}(p(t), \nu_c) + \mathbf{n}(t). \end{aligned} \quad (6.13)$$

The vector, $\mathbf{w}(t)$, defines the weights with which the different Stokes parameters are coadded together to give the measured signal, while $p(t)$ defines a sky pixel number observed at time t . Both of these can be encoded in a single matrix, the pointing matrix \mathbf{P} , which is tall and skinny with the number of rows given by the total number of measurements and the number of columns given by the total number of observed pixels multiplied by the number of observables measured in each pixel. The data model can be then expressed conveniently in a vector form as,

$$\mathbf{d} = \mathbf{P} \mathbf{m} + \mathbf{n}. \quad (6.14)$$

In each row of \mathbf{P} corresponding to time t , the only non-zero elements are those in columns corresponding to pixel p observed at the time t . These elements are given by the respective elements of the weight vector \mathbf{w} . Typically, there are therefore either 1 – for total intensity measurements, or 2 – for Q and/or U polarization measurements, or 3 – for polarization-sensitive measurements, non-zero elements in each row. Consequently the pointing matrix is typically very sparse and thus manageable in spite of its huge size.

Important examples include the case of a fixed single-slot polarizer, when the measurements can be expressed as,

$$d_t = I_{p_t} + \cos(2\varphi_t)Q_{p_t} + \sin(2\varphi_t)U_{p_t} + n_t, \quad (6.15)$$

where φ denotes an angle of the polarizer with respect to the sky coordinates. Similarly, for an idealized frequency-independent, half-wave plate, used to modulate the incoming signal, the data model reads,

$$d_t = I_{p_t} + \cos(2\varphi_t + 4\phi_t)Q_{p_t} + \sin(2\varphi_t + 4\phi_t)U_{p_t} + n_t, \quad (6.16)$$

where ϕ_t stands for the HWP orientation with respect to the instrument coordinates. In both these cases, there are three non-zero entries per row of the pointing matrix, \mathbf{P} . As mentioned earlier this number can be lower, e.g. for total intensity measurement, or in the case of pair-differencing approaches discussed in Chapter 8 (see Appendix C). It can also be larger, for instance, in the case of a more realistic model of half-wave plate, where the induced rotation of the incoming polarization depends on the frequency, the simplification leading to Eq. (6.13) are not always appropriate and one may need to introduce as many as 5 different values to characterize signal in each sky pixel, resulting therefore in 5 non-zeros per row of \mathbf{P} [459]. The number of zeros per row may be even larger in the cases of the generalized pointing matrix, \mathcal{P} , and would have to be taken into account explicitly whenever the generalized vector of amplitudes, \mathbf{y} , is to be recovered as in Eq. (6.3).

In all these cases the pointing matrix is nevertheless extremely sparse and the number of non-zero values per row is in many relevant applications very limited.

Note that while common, this is not always the case. A notable exception is the case of asymmetric beams, i.e. when \mathcal{B} depends explicitly on a vector $\delta\gamma$ and not just its length. If beams happen to be strongly asymmetric to the extent that it is necessary to correct for it on the map-making stage, the data model in Eq. (6.14) would still apply and though the pointing matrix would be in such a case dense, it would be also structured, permitting specialized algorithms for

the map-making problem as have been indeed proposed [65, 233, 279]. In the framework developed in this thesis, the focus has been specifically on the cases with an arbitrary but sparse pointing matrix. And while the framework readily allows for extensions which would accommodate also cases with asymmetric beams, this requires however further development that is left for the future and is not considered in this work.

6.2 Maximum likelihood map-making

The estimator in Eq. (6.4) is unbiased whenever the templates are appropriately chosen, given the anticipated contaminants, and the system matrix is non-singular. This property is independent of the choice of the assumed weighting. The weighting is however a key in obtaining an estimate with as high signal-to-noise ratio as possible.

A maximum likelihood estimate (which is as well a minimum variance one) is indeed given by Eq. (6.4) with the weights corresponding to noise covariance,

$$\mathbf{W} = \mathbf{N} \equiv \langle \mathbf{n} \mathbf{n}^T \rangle. \quad (6.17)$$

Given the sizes of the current and forthcoming data sets, such covariance matrices are readily unmanageable in their fully general form. However, the noise properties of any single detector data can be typically adequately characterized as piece-wise stationary. The corresponding noise covariance is then block-diagonal where each block corresponds to a different stationary interval and has structure of a Toeplitz matrix. This means that for i th interval, the correlation between the noise at times t and t' depends merely on the time interval, $|t - t'|$, and not actual values of t and t' . Consequently, the i th block of the covariance matrix satisfies the condition,

$$\mathbf{N}^{(i)}(t, t') = \mathbf{N}^{(i)}(0, |t - t'|). \quad (6.18)$$

The noise in the CMB detectors is typically correlated all the way to very long timescales and the Toeplitz blocks of the noise covariance are in general dense matrices. However, at least for well-optimized scanning strategies, accounting for the correlations on the largest time scales is usually not necessary, as they affect predominantly only signals varying on angular scales too large to be well-constrained by the experiment anyway and therefore manifest themselves merely as an offset of the map (or a submap) [438]. Consequently, it is justifiable, and desirable, to assume that the correlations are band limited and set to zero beyond some characteristic correlation length parameter, λ . The value of this parameter should be fixed according to properties of the noise and the adopted scanning strategy. I discuss this in the case of simulated data considered in this work in Chapter 8 Section 8.2.3.

In the map-making problem, Eq. (6.4), we need an inverse of the weight matrix. The inverse of a general Toeplitz matrix is not Toeplitz and it is not easily computable. The inverse Toeplitz blocks once computed would have to be stored in the computer memory, what would quickly undermine the numerical feasibility of the procedure. However, for a banded-Toeplitz matrix, in particular with

a rather narrow band, its inverse can be very well-approximated by a Toeplitz matrix. The approximate inverse is given by,

$$\mathbf{N}^{(i)-1}(t, t') \simeq \begin{cases} \mathbf{C}^{(i)-1}(t, t'), & \text{if } |t - t'| \leq \lambda \\ 0, & \text{else} \end{cases} \quad (6.19)$$

Here $\mathbf{C}^{(i)}$ is a circulant matrix embedding the Toeplitz block $\mathbf{N}^{(i)}$. The inverse of a circulant matrix is also circulant and can be easily calculated with help of Fast Fourier Transforms (FFTs) at low numerical cost [213] as discussed in the next chapter Section 7.3.

The noise covariance has to be derived typically from the same data as the one used later on for the map-making [195, 429, 465]. As a consequence even for the case of Gaussian noise in the time-domain, the noise of the resulting map is not strictly Gaussian and some care may need to be taken in minimizing or accounting for it in particular at the largest angular scales. For most of the applications this effect is however thought to be largely irrelevant.

One specific choice which is of interest, in particular because of its numerical efficiency, is that of a diagonal weight matrix. This is a limiting case of the block-Toeplitz case discussed above with $\lambda = 0$, however, the application of the weights in such cases can be done directly in the time-domain without any need for Fast Fourier Transforms. From the point of view of the quality of the estimated maps this is however a rather bad choice which will typically result in the estimates being dominated by the low frequency modes. These modes can be however often treated as templates as it is for instance the case of the destriper techniques [113, 162, 324, 393, 276]. I discuss choices of templates suitable for this task in the following section.

Whenever a prior can be imposed on the amplitude of the templates, as in Eq. (6.7), introducing the templates is equivalent to adding a low-rank correction of the weight covariance, so in such a case the weights, \mathbf{W} , are replaced by, [429],

$$\mathbf{W} \rightarrow \mathbf{W} + \mathbf{T} \Theta \mathbf{T}^T. \quad (6.20)$$

Conversely, any low rank correction to the weight matrix can be rephrased as a set of templates and implemented via the projection operator, \mathbf{F}_T .

The need for the templates is not limited to the cases when the weight matrix is diagonal. In particular, it is often desirable to use templates to mitigate some systematic effects in the maximum likelihood map-making, when the weights are block-Toeplitz. In practice, the number of such templates is expected to be lower than the number of templates needed to model the low-frequency noise correlations in the case of diagonal weights. It is therefore computationally advantageous to extend the solution vector by adding the amplitudes of the required templates and solve this extended system using standard maximum likelihood technique with $\mathbf{F}_T = \mathbf{W}$ as in Eq. (6.3). This is equivalent to the meta-pixel approach of [429]. Consequently, in the map-making framework I propose, two options are implemented: (1) a maximum likelihood map-making with an arbitrary (but sparse) pointing matrix and block-Toeplitz weights, and, (2) cases of diagonal weights with templates which are directly included in the operator \mathbf{F}_T , as in Eq. (6.5) or Eq. (6.7). However, using the functionality of the libraries I developed in this work, other cases can also be implemented.

All the considerations above apply directly to data of a single detector. In the realm of multi-kilo pixel arrays, we have to deal with data collected concurrently by many detectors. It is therefore overly optimistic to assume no correlations between the measurements of different detectors. As pointed out earlier, neglecting those will not bias our sky signal estimates however it may easily render them very noisy. Sources of the correlations can be diverse. They could be either due to some common noise source seen by multiple detectors, such an atmospheric signal unavoidably present in the case of the ground-based experiments, or can be induced in the detection chain of the instrument, for instance, in its readout system. In the case of the maximum likelihood map-making such correlations could be in principle included as an off-diagonal block of a full weight (noise correlation) matrix, which includes weights for all detectors simultaneously. However, this would lead to significant numerical overhead rendering such a map-making hardly useful in practice. Consequently, the proposed framework does not provide this functionality at this time as a default.

Some workaround approaches which are more manageable but can render comparable performance are however frequently available. One is to use templates to account for those correlations or to capitalize on differencing scheme for data of some of the detectors devised to remove the correlated component.

6.3 Templates marginalization map-making

The templates can be in principle arbitrary reflecting the diversity of possible non-cosmological contributions which may be present in the data. Whether those can be deprojected from the data on the map-making stage will in general depend on three facts: (1) whether we can define efficiently a subspace, as spanned by the columns of the matrix \mathbf{T} , which contains all possible undesirable contributions; (2) whether this can be done while retaining limited dimensionality of the subspace; and (3) whether the subspace is orthogonal to the subspace spanned by the sky signal. If (1) is not met then some unaccounted for residuals will be present in the estimated sky signal. Failing on (2) may affect the quality, signal-to-noise, of the map, but the estimate will not be biased. Higher dimensionality of the deprojected subspace increases typically the chance that (3) is not met either. If (3) is not fulfilled the system matrix in Eq. (6.4) is singular and therefore some of the modes of the sky signal may not be recoverable. Note that if such modes are properly identified the map-making procedure can still be often applied estimating all the modes which can be estimated. The information about the non-recovered modes would have to be then passed on and accounted for on all subsequent stages of the analysis.

As mentioned earlier, from the computational point of view some templates may be difficult to implement as they may lead to significant computational overhead, in particular in the construction of the orthogonalization kernel. This issue can be often sidestepped by introducing multiple independent copies of the same template specific to different, disjoint time intervals. This has a clear downside of unavoidably increasing the dimensionality of the deprojected subspace and therefore leading to loss of precision. This however may be often acceptable. Crucially, the procedure makes all the computations more local in the computer memory, thus avoiding extra need for extensive global communication between

the compute nodes. At the same time the increase of the size of the template orthogonalization kernel is mitigated by the fact that it is limited to additional diagonal blocks.

While the proposed framework allows in principle for any user-defined templates, some templates are implemented directly in the package. These are constructed following previous CMB data analysis experience. I describe them below highlighting the orthogonalization kernel structure they imply.

6.3.1 Atmosphere and $1/f$ instrumental noise

As have been discussed previously, in a typical CMB experiment the noise is correlated due to low frequency excess noise of instrumental origin, and the atmospheric emission in the case of ground experiments. We model this low frequency noise by a set of Legendre polynomials, and hence build a filtering operator which deprojects polynomial trends from the data. A subdominant part of the deprojected signal would contain long term modes of the sky signal, however given that the signal-to-noise ratio at low frequencies is very low, we expect that filtering the sky signal at these frequencies would not result in a significant loss of cosmological information.

We define a set of time intervals, $\bigcup [t_i, t_{i+1}]$, over which the amplitudes of the polynomial trends for each detector are considered to be constant. In practice, these time intervals can span the duration of each sweep in the constant elevation scans (CESs) of a ground telescope for example. The number of polynomial orders, n_{poly} , is fixed. We do not consider here common detector modes, which would model correlations between detectors. We can define each template as a compact support function, which represents one polynomial order for a given detector during one time interval $[t_i, t_{i+1}]$ (e.g. one sweep). In summary, for a detector j , there exists a set, $(\alpha_{j,k})_{k \in \llbracket 0, n_{poly}-1 \rrbracket}$ of real amplitudes, such as the $1/f$ noise component in the i -th time interval would be given by:

$$n_j(t) = \sum_{k=0}^{n_{poly}-1} \alpha_{j,k} P_k(t), \quad (6.21)$$

where P_k is the Legendre polynomial of degree k . Each of these templates represents one column of the global template matrix \mathbf{T} . By virtue of this construction and the orthogonality of Legendre polynomials, the columns would be orthogonal, making the polynomial templates kernel $(\mathbf{T}^T \mathbf{W}^{-1} \mathbf{T})^{-1}$ diagonal, simplifying the computational cost of the filtering operator.

6.3.2 Scan synchronous signal

These templates model any spurious signals which are synchronous with the telescope scans. For ground experiments, these would be dominated by ground pickup sourced by the far side lobes of the beam sensing the surrounding terrain. We model this signal with a template parametrized by the azimuth angle of the boresight pointing.

We define an azimuth grid, $\bigcup_{i=1}^{n_{az}} [\varphi_i, \varphi_{i+1}]$ for each CES, fixing *a priori* a number of azimuth bins n_{az} . Then for each detector j we consider a ground

structure, $(g_{j,i})_{i \in \llbracket 1, n_{az} \rrbracket}$, which is stable over the duration of the full CES. The template marginalization performance is in such a case contingent upon having enough redundancy in the scanning strategy to separate between the ground and sky signal over the duration of the CES. The scan synchronous signal is then given by the following,

$$\forall t : \varphi(t) \in [\varphi_i, \varphi_{i+1}] \Rightarrow SSS_j(t) = g_{j,i}. \quad (6.22)$$

In the template matrix formulation as in Eq. (6.1) the amplitude vector stores the bin amplitudes, $g_{j,i}$ and there is one column of the corresponding scan-synchronous template matrix for each azimuthal bin. The columns have consequently non-zero values only in rows corresponding to times at which the telescopes azimuth falls into the corresponding bin. As only one bin contributes to the overall measurement at any given time the template matrix has only one non-zero per row and is therefore very sparse and the columns are by construction orthogonal. The non-zero values of the matrix can be typically set to unity, however, they may be replaced by some uniquely defined function of time, such as one describing gain drifts. For the diagonal noise weights, the corresponding diagonal block of the orthogonalization kernel, $(\mathbf{T}^T \mathbf{W}^{-1} \mathbf{T})^{-1}$, will be diagonal. However, the scan-synchronous templates may not be and, in general, will not be orthogonal to the other templates, such as the polynomial templates discussed earlier.

Note that more involved models of the scan synchronous signal are readily possible. For instance, for a given grid of azimuths, $\{\varphi_i\}$, rather than binning we can linearly interpolate between the signal estimated in the grid nodes. This could be helpful whenever the scan-synchronous signal is changing rapidly with the azimuth and when the binning may require a very large number of azimuth bins in order to minimize the signal variation across each bin. In this case, the interpolation may be more efficient. There will be two non-zero values per each row of the corresponding template matrix corresponding to two azimuth grid points, φ_k, φ_{k+1} bracketing the telescope azimuth, $\varphi(t)$, and the two non-zero elements would be given by,

$$\mathbf{T}_{t_k} = \frac{\varphi(t) - \varphi_k}{\varphi_{k+1} - \varphi_k}, \quad \mathbf{T}_{t_{k+1}} = \frac{\varphi_{k+1} - \varphi(t)}{\varphi_{k+1} - \varphi_k}. \quad (6.23)$$

The corresponding kernel in such a case would be tridiagonal and therefore could be inverted efficiently.

6.3.3 Half-wave plate synchronous signal

As explained in the last chapter, many of the modern CMB experiments feature a half-wave plate which brings a significant advantage in reducing noise correlations and other systematics, but tends to introduce its own spurious signals which are typically synchronous with the HWP angle and can be modulated at different harmonics of its rotation speed [266, 304].

We could define templates for these signals in a model-independent way as in the previous section, this time relying on binning in the HWP angle. As the HWP synchronous signals tend to vary rather abruptly with the HWP angle, this may

however require a large number of bins even if some interpolation is invoked. Alternately we can model these spurious signals as a sum of harmonics of the HWP rotation frequency. This is the approach I have chosen in this work.

We define a set of time intervals, $\bigcup [t_i, t_{i+1}]$, with an *a priori* fixed length Δt , except for the remainder interval at the end of each CES. The amplitudes of the HWP harmonics in each detector are considered to be constant over each Δt . The number of HWP harmonics, n_{HWP} is fixed. In a similar manner to the polynomial templates, we can define each template as a compact support function, which represents either the cosine or the sine part of HWP harmonic for a given detector on time interval $[t_i, t_{i+1}]$. So, for a detector j , there exists two sets, $(\alpha_{j,k})_{k \in \llbracket 1, n_{HWP} \rrbracket}$ and $(\beta_{j,k})_{k \in \llbracket 1, n_{HWP} \rrbracket}$ of real amplitudes, such as the HWP synchronous signal in the i -th time interval would be given by:

$$\mathcal{H}_j(t) = \sum_{k=1}^{n_{HWP}} \alpha_{j,k} \cos(k\phi_t) + \beta_{j,k} \sin(k\phi_t), \quad (6.24)$$

where ϕ_t is the HWP angle at time t . The interval length Δt can be chosen so that we get an integer number of rotations of the HWP, in the ideal case where the rotation frequency is not fluctuating. This would simplify the structure of the HWP segment of the templates kernel. This is because each of the templates as we defined them above represents a column of the HWP segment of the templates matrix, and the proposed choice of the time intervals would make the columns orthogonal or near-orthogonal (except the last one in each CES span) by virtue of the orthogonality of cosines and sines in an integer number of periods (HWP rotations). Therefore the corresponding part of the kernel $(\mathbf{T}^T \mathbf{W}^{-1} \mathbf{T})^{-1}$ would be nearly diagonal (last column and row of each CES excluded), again speeding up its inversion and improving the performance of the filtering operator.

One could further generalize Eq. (6.24) by introducing time dependent coefficients α and β accounting on potential gain drifts. This would typically result in a full block-diagonal structure of the kernel, with each block corresponding to a different time interval and different detectors, and would therefore require a full numerical inversion.

6.4 Degeneracy issues

So far our pointing or generalized pointing matrices were considered to be full column rank, therefore ensuring that the system matrices are invertible. In practice, this assumption can be violated for different reasons which we explore in the following sections.

6.4.1 Scanning strategy sourced degeneracies

In order to solve for the three Stokes parameters I , Q and U , of each pixel in the sky, the scanning strategy should ensure, that each pixel is observed at least three times with sufficiently different orientations. The polarization modulation induced by the half-wave plate helps in this regard as it ensures that in each pixel crossing, sufficient samples are taken with different orientations to allow for the reconstruction of the pixel's Stokes parameters. However, for some pixels,

particularly in the periphery of the observed sky patch, it may happen that the pixel is only partially scanned, such that the duration of the pixel crossing is very short and does not allow its observation with sufficient redundancy.

In practice this issue is solved by examining the condition numbers of each block in the block diagonal matrix $\mathbf{P}^T \text{diag } \mathbf{W}^{-1} \mathbf{P}$. Each block corresponds to an observed sky pixel, which if observed with sufficient redundancy, should ideally have a condition number close to 2. The higher the condition number, the more the pixel noise gets boosted. We therefore set a threshold on the condition number of each pixel, above which the pixel, and its corresponding TOD samples, get effectively removed from the data.

6.4.2 Templates linear dependencies

The second types of degeneracies stem from linear dependencies between templates. One prominent example, is the offset polynomial template and the ground template which both filter the global offset of the TOD.

These degeneracies prevent the kernel matrix $\mathbf{T}^T \mathbf{W}^{-1} \mathbf{T}$ from being invertible. In practice, this does not pose any problem for the sky signal recovery and is easily addressed: we orthogonalize the templates by computing the Moore-Penrose pseudo-inverse of the kernel via a singular value decomposition (SVD). This is made possible because of the simple structure of the kernel matrix as discussed in Section 6.3.

6.4.3 Templates and sky degeneracies

This last type of degeneracies is due to linear dependencies between the filtered templates and some sky modes. When such degeneracies are present, the concerned sky modes are impossible to reconstruct and would therefore be missing from the maps, because they are filtered from the data by the deprojection operator \mathbf{F}_T .

In practice, this means that the matrix $\mathbf{P}^T \mathbf{F}_T \mathbf{P}$ would have singularities. Given the size, and the fact that this matrix is dense in general, the direct computation of these singular modes to regularize the inversion is not feasible. In an iterative approach such as the one adopted in this work, this manifests itself as a significant slow down in the convergence rate, led by the small eigenvalues in the eigenspectrum of the system matrix introduced by the singular or nearly singular modes. One way to address this issue, is to approximate these modes with an iterative method such as the Lanczos procedure, and deproject them from the system. This provides the basis for the *a posteriori* two-level preconditioning technique, that I will present in more details in the next chapter.

Massively parallel mapping and linear solvers techniques

Contents

7.1	Short introduction to parallel computing	142
7.2	Overview of the MAPPRaiser framework	143
7.3	MIDAPACK: the microwave data analysis package	144
7.3.1	MAPMAT: pointing algebra package	145
7.3.2	Toeplitz algebra package	146
7.3.3	Templates algebra package	149
7.4	MAPPRaiser: the MIDAPACK parallel iterative sky estimator	151
7.4.1	Mapping methods	151
7.4.2	MAPPRaiser solvers	152
7.4.3	Preconditioners for PCG	154
7.5	Tests and validation	159
7.5.1	Simulations description	160
7.5.2	Strong scaling tests and collective communication performance	161
7.5.3	Large run demonstration	163
7.5.4	Solvers and preconditioners performance	166

The sensitivity requirements for the forthcoming generation of CMB experiments are such that the data sets collected are expected to reach unprecedented volumes. For example, the anticipated Simons Observatory time-domain data is expected to reach the Petabytes scale. In fact, CMB data volumes have been following an exponential trend for the past decades, growing conjunctly with the so-called *Moore's law* describing the evolution of available computing capabilities

as shown in figure 7.1. To exploit the full potential of supercomputers performance, one needs to carefully design the softwares such that they can properly run on very large number of cores provided by the machines to process the huge amount of data, without detrimental degradation of their performance. This is in fact essential for map-making where the full size of the time-domain data needs to be processed, and particularly in the context of the upcoming and future CMB experiments which demand an exquisite control of systematics, thereby calling for increasingly more sophisticated data models. The challenge is also reinforced by the fact that the Moore's law for supercomputers is ending as can be seen in figure 7.1 by the dashed blue line which shows the projected actual performance, while the exponential trend for CMB data sets is expected to continue. As a consequence, significant efforts need to be spent on the development of massively parallel software, on its profiling and computational performance validation, as well as on the development of state-of-the-art numerical techniques enabling better overall efficiency. In addition, the collaborative nature of CMB projects also requires that the developed software be portable, and run with satisfactory performance on multiple platforms.

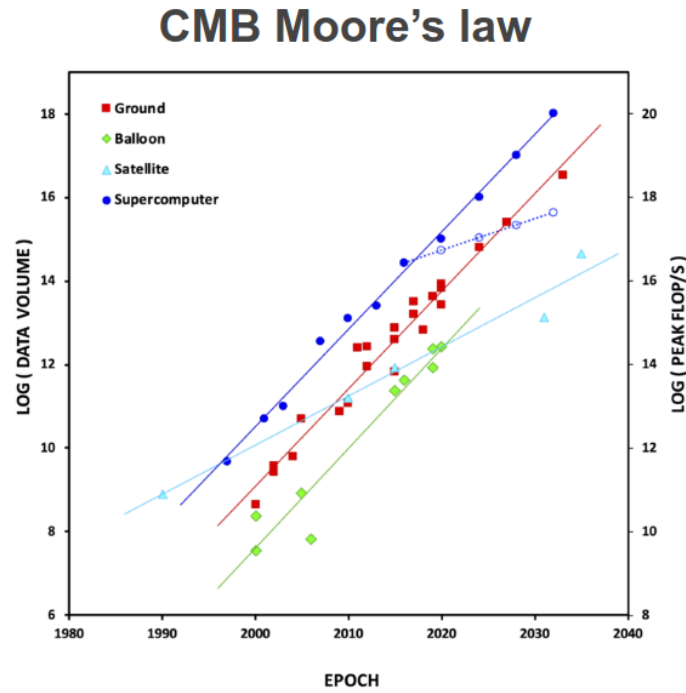


Figure 7.1 – Moore's law for CMB experiments and supercomputer performance. *Credit: Julian Borill [1].*

This is exactly the focus of the work presented in this chapter. I give a detailed description of the map-making software framework (<https://github.com/B3Dcmb/midapack>) I developed in this thesis, and go over important implementation details. I will first present an overview of the code with its overall architecture, then describe each of its building blocks and how they fit into the global framework, before presenting several demonstrations of the software performance and its implemented techniques on realistic examples.

7.1 Short introduction to parallel computing

We begin this chapter with a short introduction to parallel computing, focusing on CPU based architectures which are the ones used in the present work. For a full review one may refer to [355].

A supercomputer is typically made of multiple shared memory units, referred to as *nodes*, each hosting one or few CPUs with multiple cores. The nodes are connected via a network to allow exchange of data whenever needed. A parallel program defines a set of instructions (a sequential program) to be executed by multiple computing units at the same time. Depending on the programming model, these computing units can be either a single core, or multiple cores within a single node or multiple nodes. We can distinguish three major programming models:

- **Distributed memory model:** each sequential program, also called *process*, has only access to its local memory space. Network communication is used to exchange data between different processes. This network communication needs to follow a protocol to synchronize the messages sent and received by a large number of processes at the same time. The standard protocol used is the so-called Message Passing Interface (MPI).
- **Shared memory model:** in this model a single process is divided into multiple simultaneous tasks, called *threads*. Threads can share the same memory space, code and the hardware resources used by the process. These can be hardware threads of the same compute core (e.g. Hyper-Threading) or multiple cores sharing the same memory of a node. The standard application programming interfaces (API) used for shared memory models are OpenMP and POSIX Threads.
- **Hybrid model:** this last case makes use of the two previous programming models, featuring multiple processes, typically corresponding to different nodes, communicating via the network. Each process is divided to a number of threads typically corresponding to the compute cores available in each node. This model is more flexible, and can easily adapt with modern supercomputers architectures.

It is worth mentioning that parallel programs that make extensive use of communication over the network, can suffer from performance degradation when using a large number of processes. Therefore, studying the scaling of these programs is very important to ensure that they are able to efficiently use the computing resources to process large amounts of data. This is not the case for the so-called *embarrassingly parallel* applications where the program can be naturally parallelized by dividing the data over the allocated processes, which perform the totality of their tasks completely independently without the need to communicate.

In the context of map-making applications, the data is divided in time-domain, which results in a non-trivial distribution of the map-domain objects, with significant overlap between the processes. As a consequence, one needs to use collective communication operations, i.e. a data transfer between all processes operating in the run, when projecting the time-domain data in map-domain, particularly if one wants to take into account correlations between time

samples and detectors. This operation is not scalable, and becomes the bottleneck of map-making at large numbers of MPI processes.

7.2 Overview of the MAPPRAISER framework

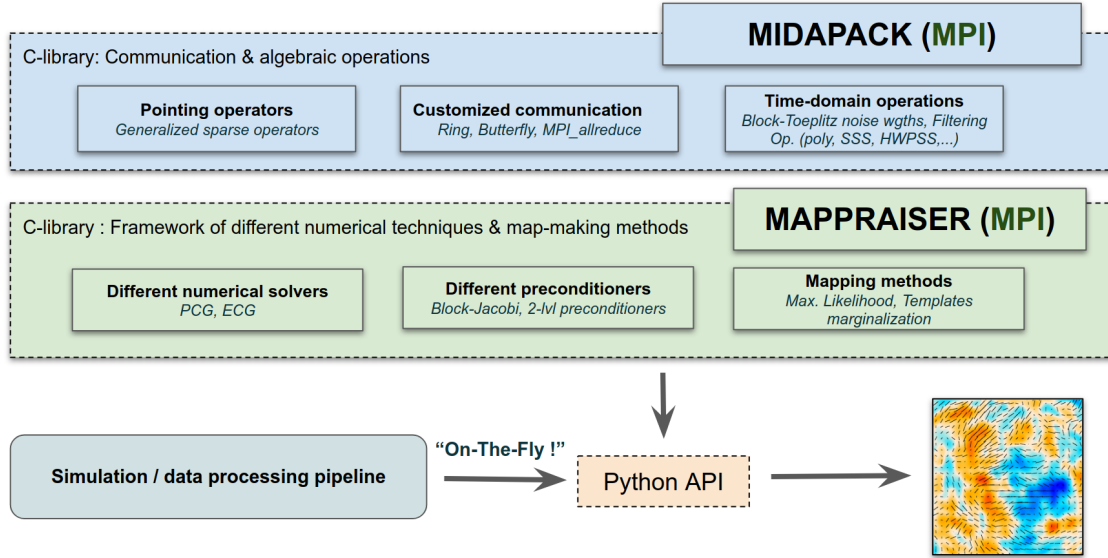


Figure 7.2 – Overview of the software top-level architecture.

A summary of the high-level code architecture, is given in figure 7.2. The software is comprised of two main C-libraries:

- The first library, called MICrowave Data Analysis PACKAge (MIDAPACK), provides low-level operations needed to perform high-level procedures in the data analysis pipelines. Such operations include, sparse linear algebraic operations for generalized pointing and template operators, customized communication schemes for distributed data reduction, operators for structured matrices, such as Toeplitz matrices, or block-diagonal templates kernels, necessary to perform some time-domain operations such as templates filtering or noise weighting. These operations can be relevant for all main stages of the data analysis pipeline, but in accordance with the scope of this thesis I mainly focus on map-making applications. Part of this library has been developed prior to this thesis, in particular the customized communication schemes and the Toeplitz algebra operations. The documentation of this old version can be found in [7].
- The second library, called MidAPack PaRAllel Iterative Sky Estimator (MAPPRAISER), builds on the first library to provide different numerical techniques, solvers and map-making methods to compute unbiased maps, and filter time-domain systematics concurrently with the map projection.

The implemented routines follow the MPI programming model and are therefore able to use massively parallel platforms with memory distributed between a large number of compute nodes. Some routines also make use of multithreading,

but this option still needs to be optimized to be fully functional in the code. However these routines may provide the basis for a future extension of the framework to a full hybrid OpenMP/MPI model.

On top of this two-level core layer of the code, there is a Python wrapper, which allows the map-making code to interface easily with simulation or real data analysis pipelines. In particular, the Python wrapper can be used to perform the map-making run “on-the-fly” chaining the map-making code with simulation software or with a previous step of an actual data processing pipeline, thereby reducing I/O cost.

7.3 MIDAPACK: the microwave data analysis package

The MIDAPACK library consists of three main sub-packages:

- **The pointing operations package MAPMAT:** this package contains sparse matrix vector products operations to perform the projections from time-domain data to map-domain data and vice-versa. In their most general form, these matrices are rectangular with considerably more rows than columns, given the relative sizes of the time-domain and map-domain data and are referred to as *tall-and-skinny* matrices. Their most important property however is their sparsity as they feature at most few non-zero elements per row.
- **The Toeplitz algebra package:** this package provides routines to perform the product of Toeplitz-structured matrices with an arbitrary data matrix. I refer hereafter to matrices which are symmetric, block-diagonal with banded Toeplitz blocks, or matrices with all of the above properties along with some rows and corresponding columns missing as “Toeplitz-structured”. These special matrices arise naturally in CMB data analysis, as they encode the correlation properties of stationary time-domain processes such as instrumental $1/f$ -noise or atmosphere. The narrow band property is due to the finite noise correlation length which typically can be assumed shorter than the length of individual scans. These routines are used in noise weighting operations, i.e weighting the different frequency modes in the time-domain data with their corresponding inverse noise power.
- **The templates algebra package:** this package provides routines to handle multiple sets of templates modeling different sorts of systematics. The routines allow to perform deprojections of templates, i.e. time-domain filtering operations as well as to orthogonalize a set of different templates to regularize the deprojection operation.

In the following subsections, I will discuss the implementation details of each one of these sub-packages. Before doing so I need to define the distribution scheme of the timestreams. In a typical experiment, we would have multiple detectors observing the sky in a number of independent time-intervals where the noise is considered to be stationary, we hereafter refer to these time-intervals as “scans”. For each scan, the detector timestreams are stacked. Typically, we can consider

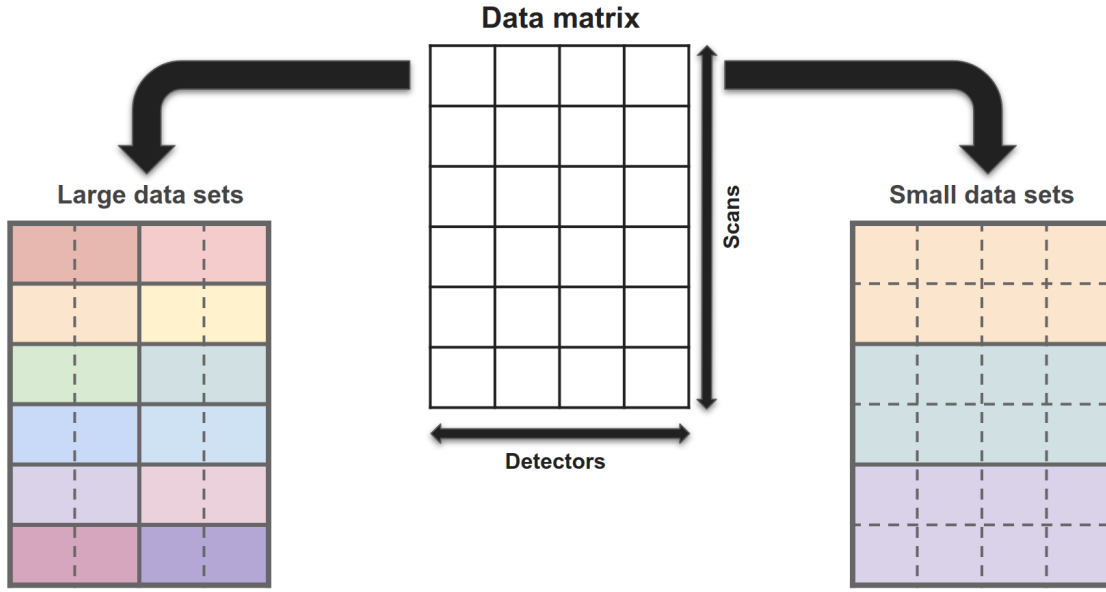


Figure 7.3 – Distribution scheme of the TOD depending on the processed data set size. Each cell in the data matrix corresponds to one detector time stream for the full duration of a scan. Depending on the size of the data set, each MPI process represented by the different colors in the diagram, will then handle a subset of detectors for a single scan or all detectors for a subset of scans. The lines separating the detector scans in a single process are dashed, since they actually hold a single contiguous block of memory.

two regimes depending on the size of the data. In the small data sets regime (small number of detectors, short scans, low sampling rate), each MPI process can hold the timestreams of all detectors for a given number of scans. In the large data sets regime (large number of detectors, long scans, high sampling rate), each MPI process will be holding the timestreams of a subset of detectors for the full duration of a single scan. To illustrate this description, we show a toy example of this data distribution in figure 7.3. The choice of not splitting single detector timestreams between multiple processes, is justified by the fact that this would add some communication between neighbouring processes in noise weighting or time-domain filtering operations, what could hinder the scalability of the application to large numbers of processors.

7.3.1 MAPMAT: pointing algebra package

Pointing data distribution & pixelization schemes

The pointing data distribution follows the same distribution scheme as adopted for the timestreams: each process will be assigned the pixels and pointing weights corresponding to the detector timestreams it is processing. The pixels indices are defined via two indexation schemes. A global pixel indexation, corresponding to the global pixel index on the sky, and a local pixel indexation, corresponding to a reordering of the pixels observed at least once by the local timestreams of each MPI process. The package is independent of the particular choice made for the sky pixelization.

Data structure

The pointing matrix is defined with a C-struct called `Mat`. Each MPI process holds an instance of this structure in memory. One integer, `nnz`, specifies the number of non-zero elements per row. The values and columns locations of these elements are specified with two arrays as is done in the ELL¹ storage format, for efficient memory consumption. It also contains the ordered set of locally observed pixels, and a precomputed set of arrays and parameters necessary to conduct the customized collective communication operations in an efficient way, along with a flag specifying the particular communication scheme chosen for these operations.

Projection operations & collective communication

The package features two main core operations:

- **The un-pointing operation:** This operation generates a time-domain vector, \mathbf{t} , from a given map-domain object, \mathbf{m} , and a pointing matrix \mathbf{P} : $\mathbf{t} = \mathbf{P} \mathbf{m}$. Given the data distribution scheme, this operation can be performed locally in each MPI process, as a simple sum of the products of the pointing weights with the corresponding map entries.
- **The pointing operation:** This is a projection operation, where given a time-domain vector, \mathbf{t} , and a pointing matrix \mathbf{P} , we compute a map vector \mathbf{m} , where each pixel is the sum of the time samples observing it: $\mathbf{m} = \mathbf{P}^T \mathbf{t}$. Given that the data are distributed in time-domain, the time samples observing a given pixel, are typically distributed over many MPI processes. Hence this operation is performed in two steps: we first compute local products of the pointing weights with the local time samples, \mathbf{m}_i , then we compute the sum of these local products over all processes to obtain the global map projection: $\mathbf{m} = \sum_i \mathbf{m}_i$. The second step requires a collective communication operation, given that these types of operations are not scalable, this represents a bottleneck, in the map-making procedure. In addition to the standard `MPI_Allreduce`, the MAPMAT package implements several customized communication algorithms such as Ring and Butterfly. One may refer to Appendix B for more details.

7.3.2 Toeplitz algebra package

Data distribution

For each detector timestream, we build a Toeplitz block that is completely defined by one array containing a number of elements equal to the half-bandwidth λ , i.e. the correlation length accounted for in the reconstruction. These blocks are therefore distributed following the time-domain data layout: each MPI process holding in memory the Toeplitz blocks corresponding to the detectors and the scans it has been assigned. In case, we want to include correlations between detectors, we would need to also build cross-correlation Toeplitz blocks, and hold

¹ a storage format originally used by the ELLPACK package, storing sparse matrices only via their non-zero entries and their corresponding column numbers.

them in memory, and eventually need a significant amount of communication depending on the spatial scale of the cross-correlations and the size of the data set.

Data structure

The Toeplitz-structured noise covariance is defined with a C-struct called `Tp1tz`. Each MPI process holds an instance of this structure in memory. It is parametrized with the total number of rows, the number of Toeplitz blocks both in total and in the local memory of the MPI process, the localization of the local segment in the global matrix, in addition to the size of the local timestreams, and an array of instances of a C-struct called `Block`. This latter structure defines each Toeplitz block, and contains the actual elements of the first row, the half bandwidth, the total size of the block and its localization in the global data structure. The `Tp1tz` structure is also assigned an MPI communicator, since the product operations can also handle data of a single scan that is split between many MPI processes, hence requiring some point-to-point communication. However, in practice, we avoid distributing the data in such way, so that we can limit the non-scalable communication cost.

Toeplitz product operations & numerical algorithms

The package provides routines to perform the product of a Toeplitz-structured matrix with a general data matrix, which is stored as vectors in the column-wise order. Three product operations can be performed:

- **stmm routines:** these concern the multiplication of a symmetric banded Toeplitz matrix by a data matrix of the same size. If the data are distributed over several MPI processes, then each process should have the same and full Toeplitz matrix stored in memory.
- **stbmm routines:** for the multiplication of a symmetric, block-diagonal matrix, with Toeplitz banded blocks by a data matrix. As explained above, such structure is defined with a list of instances of the C-struct `Block`, representing each a Toeplitz block in the block-diagonal matrix. Each block is then multiplied with the corresponding subblock of the data matrix. We also note, that the blocks need not be contiguous, i.e. we can have rows in the data matrix which correspond to no given block in the Toeplitz structure, such rows are copied unchanged implicitly assuming the corresponding entries in the Toeplitz structured matrix to be equal to 1. Given the data layout adopted in practice, each MPI process will have a number of full detector timestreams, and therefore all the corresponding Toeplitz blocks. However, in general, as was discussed before the package also handles, detector timestreams being split between multiple MPI processes, in this case each of those MPI processes should have the corresponding full Toeplitz blocks stored in memory, and communication is used to copy necessary data to enable the product operation locally.
- **gstbmm routines:** these are similar to stbmm routines, but with some sets of rows and corresponding columns removed from the Toeplitz-structured matrix. These are called *gaps*. The routines perform a product in a similar way

to the `stbmm` routines, but with neglecting the gaps, i.e. the corresponding rows in the output are set to zero.

So far we have discussed, which operations may be called according to the data structure, but we have not discussed how the product operation with the noise covariance blocks is performed in itself. The algorithm is based on a *shift-and-overlap* approach [383], where a product of a single band-diagonal Toeplitz matrix by a data matrix is done by dividing them into a set of overlapping blocks, and performing a sequence of products of the overlapping Toeplitz subblocks by the corresponding overlapping segments of the data matrix (Cf. figure 7.4). Each of the latter products is performed in turn by embedding the Toeplitz subblock in a minimal circulant matrix and performing the multiplication via Fast Fourier transforms. The edges of the resulting product of each of the subblocks is biased by contributions from the corners of the circulant matrix. The overlap between subblocks allows to circumvent this issue, as we can drop these edges without getting any gaps in the global product after combining all the subblocks products together. The edges of the input vector are padded with zeros. The size of the subblock is set appropriately to optimize the computation and is typically a multiple of the half bandwidth λ . The overall complexity of the operation is $\mathcal{O}(n \ln \lambda)$, where n is the size of the initial Toeplitz matrix. Performing the operation directly by embedding the full Toeplitz matrix into a circulant one that is twice as large and computing the product via FFTs would yield a complexity of $\mathcal{O}(2n \ln 2n)$.

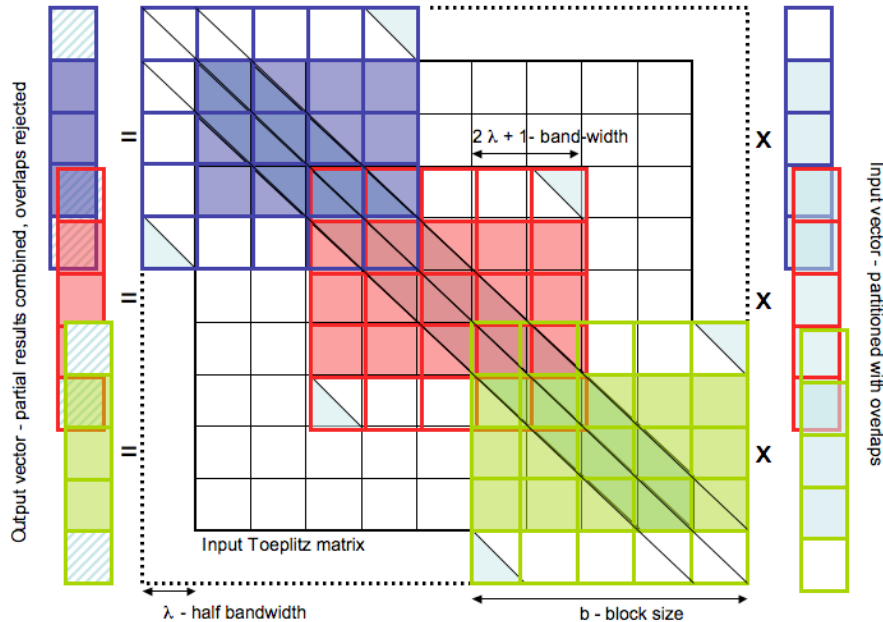


Figure 7.4 – Illustration of the shift and overlap algorithm implemented in the Toeplitz algebra package. *Credit:* MIDAPACK v1.1b documentation [7].

When communication is needed as in the cases discussed before, it remains local, involving only neighboring processes. Each process needs to send to and receive from a neighboring process a vector of data of the length defined by the half bandwidth of the Toeplitz block, λ , shared between them. This is sufficient to compute the part of the Toeplitz-vector product corresponding to the input data of each process locally. In particular, we note that all the FFT calls used by

the package are either sequential or threaded. The point-to-point communication pattern enabled by the package can be either *non-blocking*, using `MPI_Isend` and `MPI_Irecv` calls, or *blocking* with `MPI_Sendrecv` calls. The non-blocking communication allows computations and communication to overlap, leading to better performance at the cost of being less simple to use than the blocking communication. The choice is specified via a flag `FLAG_BLOCKINGCOMM` determined when constructing the `Tpltz` structure, and set by default to 0 for non-blocking communication.

7.3.3 Templates algebra package

Data distribution

The templates are defined in such a way so as to match with the time-domain data distribution. At the input of the map-making run, a number of templates classes, i.e. models of time-domain systematics of different orders, are defined. The detector timestreams are therefore modeled as a linear combination of the sky signal, these templates classes, and some noise. Each template is a compact support function that spans the duration of the full scan for one given detector. Hence, each MPI process is assigned the set of templates corresponding to the detector time streams it is processing. Contrary to the Toeplitz package, no splitting of a detector time stream between processes is supposed to happen, as no communication is performed and all operations are local.

Data structure

Models of time-domain systematics are defined with a C-struct called `TemplateClass`. The structure defines the time-domain range of the template, i.e. the time samples indices of the detector time stream it is associated with. Each template is discretized with regards to certain parameter: it can be time, for example, by splitting the scan to a number of time intervals over which the amplitudes of the templates are constant, boresight azimuth angle for scan synchronous signals, or other parameters depending on the nature of the signal modeled by the template. The C-struct defines the bin range of each template, in the total set of local bins (time intervals, azimuth grid, etc) spanned by the MPI process time streams. Additional flags define the type of template class considered (e.g. polynomial template, or scan synchronous template, and its order when it applies), information about the data segment it is associated to (detector, scan, season, etc). Finally two time-domain arrays, similarly to the pointing data storage format, specify the bins corresponding to each time sample, and their corresponding weights, i.e the template signature in time-domain. Given that templates are time-domain objects, they can have a significant memory footprint. Therefore, two operating modes are implemented for the templates: the weights and bins can either be fully stored in memory at the moment of templates construction, if the processed data set is sufficiently small with a manageable memory footprint, or they can be computed on-the-fly while conducting the product operations with the templates matrix, for large data sets, hence trading speed for memory efficiency. Each MPI process holds an array of instances of `TemplateClass` in memory, and as such would have all necessary information to perform the time-

domain filtering operations locally, without any need for communication with other processes.

Templates algebraic operations

The main utility of the routines provided by this package, is to allow the application of the filtering operator, F_T , to a time-domain vector. Three core routines form the basis for such an operation:

- **Templates projection in time-domain:** This operation generates a time-domain vector, t , from a given set of templates amplitudes, x , and a templates matrix, T : $t = T x$. Each process loops over the set of scans it is assigned, and for each scan, loads a set of templates weights, which typically depend only on boresight pointing parameters and therefore do not vary for individual detectors in the same scan. Higher order templates weights are constructed recursively, by using for example recursive relations between Legendre polynomials, or trigonometric functions. This allows us to avoid repeating the recursive computation for each detector time stream, which can significantly slow down the operation. Once the weights are stored in memory, we loop over the array of templates classes: for each template we construct the array of bins on-the-fly, point to the already loaded weights, perform the series of products with the corresponding segment of templates amplitudes, and finish with freeing the memory segment occupied by the bins, and move on to the next template class. Once all the templates of the scan have been spanned, we remove the loaded weights from memory, and repeat the process for the next scan, until all scans are processed.
- **Time-domain projection in templates space:** This operation projects a time-domain vector, t , in the templates amplitudes space, using a templates matrix T : $y = T^T t$. The operation is performed in a similar way to the one above, with loops over scans and templates classes, and templates bins and weights constructed on-the-fly before computing the products with the relevant segments of data.
- **Templates kernel construction and inversion:** What we refer to as templates kernel here, is the block-diagonal matrix $T W^{-1} T$. Each block of this matrix is constructed solely from the set of templates that apply to a single detector time stream, as long as the templates do not model any common detector modes. As such the construction of this kernel as well as its inversion are considerably simplified. We loop over the scans, and the detectors, and for each detector in a given scan, we perform a double nested loop on the templates classes associated with the detector time stream, weights and bins are computed on-the-fly depending on the template class, and we compute the kernel entries with a series of products between the different templates weights. The inverse of the kernel is simply the inverse of each block. The blocks are inverted with an SVD, where singular values below a certain threshold, typically 10^{-12} , are set to zero, effectively regularizing the inversion. This is equivalent to an orthogonalization procedure, and therefore, spurious effects due to linear dependencies between templates are neutralized. Note that it is possible to explicitly exploit the inner structure of the

kernel blocks to speed-up both their inversion and construction. For example, we know that Legendre polynomials or cosine and sine harmonics over a full period are orthogonal by definition, and would therefore yield a diagonal kernel, therefore we can ignore some off-diagonal entries while constructing the kernel, and exploit the structure of the kernel when performing the inversion. As in the current applications the kernel inversion does not constitute significant fraction of the total runtime such optimizations are not included in this work and are left for future extensions.

7.4 MAPPRaiser: the MIDAPACK parallel iterative sky estimator

7.4.1 Mapping methods

As discussed in the last chapter, Eq. (6.4) represents a broad class of plausible map-making approaches which result in an unbiased estimate of the sky signal. However, implementing it in its full generality is cumbersome and the resulting code likely to be rather inefficient. Instead, in the proposed framework I have implemented two highly optimized, specific cases of this general estimator, which taken together cover most of the cases of practical interest.

These two specific map-making methods are,

1. the case, where the weights, \mathbf{W}^{-1} are assumed to have the form of a Toeplitz-block-diagonal matrix and the deprojection/weighting operator, \mathbf{F}_T , contains no templates, i.e. $\mathbf{F}_T = \mathbf{W}^{-1}$. The templates can be included as part of the generalized pointing matrix, see Eq. (6.3). Note that this method allows for a (nearly) maximum likelihood solution for the map in cases when the time-domain noise is piece-wise stationary and uncorrelated between different detectors. Indeed, the imposed structure of the weights is flexible enough to allow for a very good approximation to the actual inverse noise covariance as discussed in Section 6.2. The MAPPRaiser package computes the weights using information derived directly from the collected data as in [429, 386, 465]. To do so, we consider that over each scan, every detector timestream is dominated by stationary noise with a power spectral density (PSD) of the form:

$$P(f) = \sigma_0^2 \left(1 + \left(\frac{f + f_0}{f_k} \right)^\alpha \right) \quad (7.1)$$

Assuming at first order that the data is only made of noise, we compute the average periodogram of each detector time stream, and fit it to the above model. The obtained PSD is then inverse Fourier transformed to obtain the noise autocorrelation which is later smoothed with a gaussian window $\mathcal{N}(0, \lambda/2)$, and cut at the *a priori* specified half bandwidth λ . The smoothing avoids any ringing effects, and the resulting autocorrelation represents the first row of the corresponding Toeplitz block of the noise covariance matrix. Once we have the noise covariance we then follow the steps outlined in Sect. 6.2 to compute the approximation to the inverse covariance. In general,

even if efforts are made to use as good an approximation as possible, the weights do not have to reflect the actual inverse noise covariance. In such cases the estimated map will be still unbiased but not optimal in the sense of ensuring the lowest possible uncertainty. Nevertheless, for simplicity, we refer to this approach hereafter as a maximum likelihood map-making; and,

2. a variant of the method which allows for arbitrary templates but assumes the weights to be diagonal. This method includes all the predefined templates provided in the software as described in Section 6.3. We call this method in the following a template map-making.

7.4.2 MAPPRaiser solvers

Linear systems can be solved using either direct or iterative methods. Direct methods typically compute some decomposition of the system matrix, and consequently are time consuming, scaling in general as a cube of the size of the system. Their advantage is that they give a (numerically) precise solution. Iterative methods form successive approximations that converge to the exact solution, and are particularly well-adapted to solving large sparse systems. They require only efficient routines for a product of the system matrix and an arbitrary vector and are very efficient in terms of any extra storage. This first property is particularly relevant in the map-making case where the inverse of the system matrix cannot be typically explicitly computed due to both memory and computational cycle issues. Instead, the form of the system matrix allows for an efficient computation of a product of the matrix by a vector. Indeed, this can be implemented [123] by performing the operations from right to left. In the case of the maximum likelihood method, this involves a calculation of the matrix-vector product in the form,

$$\mathbf{P}^T \mathbf{W}^{-1} \mathbf{P} \mathbf{x}, \quad (7.2)$$

what can be done by performing three subsequent operations solely on vectors: (1) unpointing; (2) a block-diagonal Toeplitz product; (3) pointing, all of which are implemented in the MIDAPACK library described earlier. In the case of the template map-making the respective product reads,

$$\mathbf{P}^T \mathbf{F}_T \mathbf{P} \mathbf{x}, \quad (7.3)$$

and again requires three operations performed one after another with the Toeplitz product above replaced now by an application of the deprojection/weighting operator, \mathbf{F}_T . This in turn requires a precomputation of the inverse kernel and then merely a series of template pointing and depointing operations all available in MIDAPACK.

MAPPRaiser provides two iterative solvers for the map-making equation. Both are based on the Conjugate Gradient approach. These are Krylov projection methods [213] suitable for symmetric (Hermitian) positive definite matrices adapted to the specificity of the map-making problem. I discuss them below.

Preconditioned conjugate gradient solver

The first one is the Preconditioned Conjugate Gradient (PCG), which is used to solve Symmetric Positive Definite (SPD) systems. This solver has already been extensively used in the literature to solve the map-making problem

[350, 477, 123, 438]. One can refer to these papers for the technicalities of the method. Below we illustrate the algorithm of PCG solver, where M is representing a preconditioner. The objective of applying a preconditioner is to reduce the condition number of the system matrix. In MAPPRaiser, two kinds of preconditioners are implemented. These are the Block Diagonal preconditioner (BD) and the two level preconditioner ($2lvl$). The block-diagonal preconditioner defines the current standard in the field. It is easy to compute and in many cases performs already very well. The two-level preconditioner has been proposed and demonstrated for the map-making problem in [438] and extended to component separation in [359]. It is more involved and requires a precomputation. To the best of my knowledge the MAPPRaiser implementation of this preconditioner is the very first and only existing implementation in actual map-making software.

```

Compute  $\mathbf{r}_0 = \mathbf{b} - \mathbf{A}\mathbf{x}_0$ ,  $\mathbf{z}_0 = M\mathbf{r}_0$  and  $\mathbf{p}_0 = \mathbf{z}_0$ ,
for  $k = 0, 1, \dots, k_{max}$  do
     $\alpha_k = \langle \mathbf{r}_k, \mathbf{z}_k \rangle / \langle \mathbf{A}\mathbf{p}_k, \mathbf{p}_k \rangle$ 
     $\mathbf{x}_{k+1} = \mathbf{x}_k + \alpha_k \mathbf{p}_k$ 
     $\mathbf{r}_{k+1} = \mathbf{r}_k - \alpha_k \mathbf{A}\mathbf{p}_k$ 
    if  $\|\mathbf{r}_{k+1}\|_2 < \varepsilon \|\mathbf{b}\|_2$  then stop
     $\mathbf{z}_{k+1} = M\mathbf{r}_{k+1}$ 
     $\beta_k = \langle \mathbf{r}_{k+1}, \mathbf{z}_{k+1} \rangle / \langle \mathbf{r}_k, \mathbf{z}_k \rangle$ 
     $\mathbf{p}_{k+1} = \mathbf{z}_{k+1} + \beta_k \mathbf{p}_k$ 
end

```

Algorithm 1: The PCG algorithm for the system $\mathbf{Ax} = \mathbf{b}$ with a preconditioner M [401].

Enlarged conjugate gradient solver

The second iterative solver is the Enlarged Conjugate Gradient (ECG) [216]. In this approach the matrix \mathbf{A} is partitioned into \mathcal{N} subdomains. Then the unknown vector \mathbf{x} is split into t vectors, $\mathbf{X}^{1 \leq i \leq t}$, such that we can still retrieve the original vector by summing them: $\mathbf{x} = \sum_{i=1}^t \mathbf{X}^{(i)}$. The parameter t is called the enlarging factor, and the method is in practice not sensitive to the choice of the splitting scheme [218]. Following these considerations, the method can be derived in a similar way to the standard CG solver, replacing the residuals and search directions by $\mathcal{N} \times t$ matrices, and the optimal step which is a scalar in CG becomes a $t \times t$ matrix. The algorithm is shown below.

The rate of convergence of the ECG is given by the following result demonstrated in [218]: if \mathbf{x}_k is the approximate solution given by the Enlarged Conjugate Gradient with an enlarging factor t at step k , and \mathbf{x}_* the true solution, then we have

$$\|\mathbf{x}_k - \mathbf{x}_*\|_{\mathbf{A}}^2 \leq C \left(\frac{\sqrt{\kappa_t} - 1}{\sqrt{\kappa_t} + 1} \right)^{2k}, \quad (7.4)$$

where $\kappa_t = \lambda_n / \lambda_t$, the ratio of the largest eigenvalue of \mathbf{A} , and its t smallest eigenvalue, and C a constant independent of k . As such the solver behaves as

Compute $\mathbf{r}_0 = \mathbf{b} - \mathbf{A}\mathbf{x}_0$, and split it into a matrix with t columns, \mathbf{R}_0^e , s.t.

```

 $\mathbf{r}_0 = \sum_{i=1}^t \mathbf{R}_0^{e(i)}$ ,
 $\mathbf{Z}_1 = \mathbf{M}\mathbf{R}_0^e$ ,
for  $k = 1, 2, \dots, k_{max}$  do
     $\mathbf{P}_k = \mathbf{Z}_k (\mathbf{Z}_k^T \mathbf{A} \mathbf{Z}_k)^{-1/2}$ 
     $\boldsymbol{\alpha}_k = \mathbf{P}_k \mathbf{R}_{k-1} \triangleright t \times t$  matrix
     $\mathbf{X}_k = \mathbf{X}_{k-1} + \mathbf{P}_k \boldsymbol{\alpha}_k$ 
     $\mathbf{R}_k = \mathbf{R}_{k-1} - \mathbf{A} \mathbf{P}_k \boldsymbol{\alpha}_k$ 
    if  $\|\sum_{i=1}^t \mathbf{R}_k^{(i)}\|_2 < \varepsilon \|\mathbf{b}\|_2$  then stop
    Construct  $\mathbf{Z}_{k+1}$  s.t.  $\mathbf{Z}_{k+1}^T \mathbf{A} \mathbf{P}_i = \mathbf{0}, \forall i \leq k$ 
end
 $\mathbf{x}_k = \sum_{i=1}^t \mathbf{X}_k^{(i)}$ .
```

Algorithm 2: The ECG algorithm for the system $\mathbf{A}\mathbf{x} = \mathbf{b}$ with a preconditioner \mathbf{M} [218].

a deflated-CG [400], and converges faster than CG. However, the rate of convergence is not a sufficient metric to assess the performance of the method. In practice, a more relevant figure of merit is instead time to solution. In the case of ECG the expanded map space leads to an increase of the number of operations required on each iteration of the solver by the enlarging factor t , therefore a trade-off needs to be found between the gain in the rate of convergence and the linear increase in the algorithm's complexity. In the scaling experiments I conducted whose results are given in figure 7.12, I adopt a naive implementation, where the product of the matrices \mathbf{A} and \mathbf{X} is computed by looping on the columns of the enlarged map, and performing the standard CG $\mathbf{A} \mathbf{x}$ operation for each column. This baseline implementation of the solver demonstrates the improvement in the number of iterations with increasing the enlarging factor as theoretically predicted. The gain in terms of the number of iterations is however sublinear and does not compensate for the increased cost of each iteration. Consequently, the overall run time increases. This can be expected to change with an improved approach to performing the product of the system matrix and expanded map. One particularly promising avenue is to capitalize on GPU-accelerators. This is left for future work. As is, the main interest of the ECG solver is in the fact that it does not require sophisticated preconditioning and is nevertheless very efficient in the cases when the system matrix includes a number of small eigenvalues.

7.4.3 Preconditioners for PCG

The objective of preconditioning is to improve the convergence and stability of the iterative solvers. The preconditioner, \mathbf{M} , can be thought of as an approximate to the inverse of the system matrix, \mathbf{A} . In preconditioned solvers, instead of the original problem, we solve a modified system with the system matrix replaced by a product of the system matrix and the preconditioner, i.e. $\mathbf{M} \mathbf{A}$. Unlike the actual inverse of the system matrix the preconditioner should be easy to compute, and the preconditioned system matrix, $\mathbf{M} \mathbf{A}$ should be better conditioned and its eigenspectrum more clustered. Both of these properties make the system

more stable and speed-up the convergence. The MAPPRaiser package offers two preconditioning options.

Block-Jacobi preconditioner

The first preconditioner is the standard Block-Jacobi preconditioner defined as,

$$\mathbf{M}_{BD} = (\mathbf{P}^T \text{diag } \mathbf{W}^{-1} \mathbf{P})^{-1}, \quad (7.5)$$

where the inner most matrix denotes a diagonal matrix given by the diagonal of the weight matrix. This preconditioner is block-diagonal with square diagonal blocks sizes of which are determined by the number of non-zero elements per row of the pointing matrix. The construction of this preconditioner is done in two steps. First, for each pixel we sum over the (non-zero) elements of all the rows of the pointing matrix corresponding to this pixel while weighting them by the diagonal noise weights. On the second step, we invert each of the blocks using a direct method. On this second step, we also evaluate a condition number of each block, and use it as a criterion to excise pixels with ill-conditioned blocks from the map-making.

Two level preconditioners

The first level of the two-level preconditioners considered in MAPPRaiser is based on the deflation technique. Its purpose is to suppress (deflate) an unwanted subspace of a matrix, so as a result it is contained in its null space. In the context of preconditioning this is ideally the subspace which contains all the smallest eigenvalues of the system matrix. The second level then consists of adding an extra contribution to the deflated matrix in order to shift the eigenvalues of the deflation subspace to unity. There are different variants of the two-level preconditioners which have been shown to be efficient in different applications. Here, we follow the proposal of [438] and focus on the so called “Adapted Deflation Variant 1” (A-DEF1) option in [441]. The preconditioner is defined as follows,

$$\mathbf{M}_{2lvl} = \mathbf{M}_{BD} (\mathbf{I} - \mathbf{A} \mathbf{Q}) + \mathbf{Q}, \quad (7.6)$$

$$\mathbf{Q} = \mathbf{Z} (\mathbf{Z}^T \mathbf{A} \mathbf{Z})^{-1} \mathbf{Z}^T, \quad (7.7)$$

where \mathbf{Z} is a deflation subspace matrix. Its columns span a subspace to be deflated. This is a tall and skinny matrix with a limited number of columns each of which are pixel-domain vectors. The number of columns of \mathbf{Z} defines the volume of the deflation subspace. \mathbf{A} stands for our system matrix. The first term in Eq. (7.6) defines the first level of the preconditioner. Indeed, we have $(\mathbf{I} - \mathbf{A} \mathbf{Q}) \mathbf{A} \mathbf{Z} = 0$. The second term is the second level correction.

A key element determining the performance of the two level preconditioner is a construction of the deflation subspace matrix, \mathbf{Z} . The MAPPRaiser implementation of the PCG solver with the two level preconditioner accepts any user provided matrix \mathbf{Z} . In addition, MAPPRaiser includes functionality which allows for calculating the deflation matrix based on two different approaches. I describe those in the following.

Note that the two-level preconditioner is in principle a dense, square matrix with the dimension determined by the number of the observed sky pixels. For this reason, we never explicitly compute it as in many applications this would be prohibitive both from the point of view of the required computations and the computer memory. Instead, for any available deflation matrix, \mathbf{Z} , we only precompute matrix $\mathbf{Z}' \equiv \mathbf{A} \mathbf{Z}$, the so-called coarse matrix, $\mathbf{E} \equiv \mathbf{Z}^T \mathbf{A} \mathbf{Z} = \mathbf{Z}'^T \mathbf{Z}'$, and its inverse. Computing \mathbf{Z}' is computationally demanding as we have to apply the system matrix to as many vectors as the number of columns of the matrix \mathbf{Z} . As each such product is also a dominant computational cost of each iteration of the PCG solver, this computation is as costly as performing as many iterations as the volume of the deflation subspace. As discussed later, this cost can be however effectively hidden in some of the procedures aiming at the calculation of the deflation matrix, \mathbf{Z} . This is indeed the case in one of the specific constructions implemented in MAPPRASER. The matrix \mathbf{E} is then computed by calculating $\mathbf{Z}'^T \mathbf{Z}'$, which entails negligible cost. The coarse matrix is in most of the cases a small, dense matrix. We therefore can compute its inversion with help of standard, highly efficient, dense linear algebra routines. If it is known a priori that the columns of \mathbf{Z} are all linearly independent, and therefore the coarse operator is positive definite this could be done very efficiently using the Cholesky decomposition. Otherwise, one has to use a pseudo-inverse, which relies on the calculation of the full eigenspectrum decomposition of the coarse matrix and is therefore appropriately more costly. In this implementation, the second approach is chosen, as the calculation of the inverse is in any case subdominant. I use the *dgels* LAPACK routine in the MKL implementation for this purpose. All these operations have to be performed only once at the beginning of the solution. In addition, we need to precompute the block-Jacobi preconditioner following the procedure outlined earlier at negligible cost.

The preconditioner has to be applied once at every step of the PCG iterations. This involves a computation of a product of the preconditioner and some arbitrary pixel-domain vector. For this purpose we represent the preconditioner in Eq. (7.6) as,

$$\mathbf{M}_{2lvl} = \mathbf{M}_{BD} - \mathbf{M}_{BD} \mathbf{Z}' \mathbf{E}^{-1} \mathbf{Z}^T + \mathbf{Z} \mathbf{E}^{-1} \mathbf{Z}^T. \quad (7.8)$$

We apply the preconditioner by multiplying all factors of each of its terms one-by-one, from right to left to a vector capitalizing on precomputed and stored in memory objects. All these objects are either pixel-domain or deflation-subspace objects and the required operations are respective matrix-vector products which require at most $\mathcal{O}(\mathcal{N}_{\text{pix}} \dim \mathbf{Z})$ floating point operations. Unless $\dim \mathbf{Z}$ is very large this is much less costly than application of the system matrix to a vector. The columns of the matrix \mathbf{Z} are distributed between the MPI processes in the same way as any other map. This saves the memory but requires one extra global communication call of the allreduce type per iteration. The size of the reduced data is given by the dimension of the deflation space and the time overhead therefore negligible. Consequently, the total cost of the application of the preconditioner is subdominant and the application of the system matrix to a vector remains the dominant cost of a single iteration. Therefore, the time of a single iteration is essentially the same for the two-level and block-Jacobi preconditioners. MAPPRASER provides two approaches to calculating the deflation subspace

matrix, \mathbf{Z} . These are,

- **2lvl a priori construction**

This construction is based on the structure of the data set, which is assumed to be divided into a number of time intervals (such as noise stationary intervals corresponding to the blocks of the noise covariance, \mathbf{N}). This construction is applicable only if total intensity is part of the data model. In this case, we calculate the number of observations of pixel p_i during the j -th interval, denoted as s_i^j . s_i is then the total number of observations of pixel p_i . The (i, j) element of the deflation matrix, \mathbf{Z} , which corresponds to the total intensity signal is given by the fractional number of observations of the pixel p_i in the j -th interval, i.e. $\mathbf{Z}_{ij}^{tint} = s_i^j / s_i$. Consequently every such row of \mathbf{Z}^{tint} represents a partition of the unity, since $\sum_j \mathbf{Z}^{tint}(p, j) = 1$. The rows corresponding to Q and U Stokes parameters are set to 0 [438].

- **2lvl a posteriori construction**

In this construction, we attempt to estimate the relevant eigenpairs of the system matrix, \mathbf{A} . We denote these as (θ, \mathbf{y}) standing for an eigenvalue and the corresponding eigenvector, respectively. We use the Lanczos algorithm [213, 401]—given explicitly below in algorithm 3—to calculate an m -by- m tridiagonal matrix, \mathbf{T}_m and an \mathcal{N}_{pix} -by- m column orthonormal matrix \mathbf{V}_m such as,

$$\mathbf{A}\mathbf{V}_m = \mathbf{V}_m\mathbf{T}_m. \quad (7.9)$$

The Lanczos algorithm is an iterative procedure and m defines the number of the iterations. It is bound to be smaller than the size of the system matrix. In our case, $m \ll \mathcal{N}_{\text{pix}}$. The algorithm is applicable to any symmetric positive-definitive (SPD) matrix and requires only matrix-vector products. Here, these are implemented using the same routines as used by the PCG algorithm. Given the structure and the limited size of the matrix \mathbf{T}_m , its eigen-structure can be straightforwardly calculated. Given Eq. (7.9), if (θ, \mathbf{x}) is an eigenpair of \mathbf{T} then $(\theta, \mathbf{V}_m \mathbf{x})$ is a corresponding eigenpair of \mathbf{A} . Therefore $\mathbf{V}_m \mathbf{x}$ define the columns of the deflation matrix, \mathbf{Z} . We take all the precomputed vectors, given that the cost of applying the two level preconditioner is subdominant.

The Lanczos algorithm computes internally products of the system matrix and columns of the matrix \mathbf{V}_m , which are usually discarded. We store those, i.e. the matrix $\mathbf{A} \mathbf{V}_m$, together with the matrix \mathbf{V}_m and use them to compute directly the matrices $\mathbf{A} \mathbf{Z}$ and \mathbf{Z} . This saves significant precomputation time needed to construct the preconditioner as discussed earlier.

This procedure is a special case of a more general procedure proposed in [438] and employed in [387], which uses the so-called GMRES algorithm, which is more numerically stable and applicable to a general matrix. I find the Lanczos procedure to be more computationally efficient and sufficiently robust in my applications. The main difference between the two approaches is that the present procedure is applied to the system matrix and not to the system matrix preconditioned by the block-Jacobi preconditioner as in [438, 387].

```

Compute  $\mathbf{r}_0 = \mathbf{b} - \mathbf{A} \mathbf{x}_0$ ,  $\beta_1 = \|\mathbf{r}_0\|_2$ , and  $\mathbf{v}_1 = 1/\beta_1 \mathbf{r}_0$ 
for  $k = 1, 2, \dots, m$  do
     $\mathbf{w}_k = \mathbf{A} \mathbf{v}_k - \beta_k \mathbf{v}_{k-1}$  (if  $k = 1$  set  $\beta_1 \mathbf{v}_0 \equiv \mathbf{0}$ )
    store  $\mathbf{A} \mathbf{v}_k$ 
     $\alpha_k = \langle \mathbf{w}_k, \mathbf{v}_k \rangle$ 
     $\mathbf{w}_k := \mathbf{w}_k - \alpha_k \mathbf{v}_k$ 
     $\beta_{k+1} = \|\mathbf{w}_k\|_2$ 
    if  $\beta_{k+1} < \varepsilon$  then  $m := k$  and break loop
     $\mathbf{v}_{k+1} = 1/\beta_{k+1} \mathbf{w}_k$ 
end
return  $\mathbf{T}_m = \text{tridiag}(\beta_i, \alpha_i, \beta_{i+1})$ ,  $\mathbf{V}_m = [\mathbf{v}_1, \dots, \mathbf{v}_m]$ , and  $\mathbf{A} \mathbf{V}_m$ .

```

Algorithm 3: The Lanczos algorithm [401] for the matrix $\mathbf{A} = \mathbf{P}^T \mathbf{W}^{-1} \mathbf{P}$. The ε threshold is set to machine precision, and was never reached in practice.

I compare advantages of different solvers in detail in the following sections. As a rule of thumb, each iteration of the PCG algorithm is as costly as applying the system matrix to a single column of the deflation matrix, \mathbf{Z} , and as costly as a single iteration of the Lanczos procedure. Denoting by n_{iter} the number of the PCG iterations, solving the map-making problem using the block-Jacobi preconditioner requires n_{iter} products of the system matrix times a vector. For two-level preconditioners, this amounts to $n_{iter} + \dim \mathbf{Z}$ where the second term is due to the Lanczos iterations for the a posteriori preconditioner and the application of the system matrix to the deflation matrix in the case of the a priori preconditioner. All the other costs are then typically subdominant. If the same system or similar systems (see [387]) need to be solved multiple times the respective costs are: $n_{solves} n_{iter}$ and $n_{solves} n_{iter} + \dim \mathbf{Z}$. Obviously the key parameter is the number of iterations required to reach desired precision for different types of the preconditioners and different dimensions of the deflation space. It can be argued on theoretical grounds that $n_{iter}^{BD} \geq n_{iter}^{2lvl}$ independently of the deflation matrix [438]. This is because the two-level preconditioner cannot increase the condition number of the preconditioned system matrix. The gain can be very significant if \mathbf{Z} encompasses the smallest eigenvalues of the system matrix. In such cases, the precomputation time can be offset very quickly if multiple, similar map-making problems have to be solved, leading to substantial overall net gain for the full computation. In fact, whenever the number of iterations per solve is known, we can estimate the minimal number of the map-making solves required to reach this regime. I investigate this in detail in the next section.

Note that the deflation matrix computed with the Lanczos algorithm encodes useful information about the system matrix and therefore the error covariance of the estimated map. This information can be useful on the subsequent steps of the analysis.

7.5 Tests and validation

In this section, I present several tests that I have performed to validate and demonstrate the computational and scaling performance of the software framework. Two main aspects drive my discussion throughout this section: the first is how well optimized are the routines used by the solvers to perform the various operations. Here I focus mainly on the time cost of each CG iteration. In particular, I investigate how well do they scale as I increase the number of cores processing the data. The second is how efficient are the solvers themselves, taking different metrics into consideration. In particular, I give a detailed discussion of the properties of the Enlarged CG and the preconditioning techniques presented earlier.

I run the simulations on two different HPC platforms:

- **Cori:** This is a Cray XC40 machine based at NERSC². It can reach a peak performance of about 30 petaflops. It has two partitions, the first one is made of 2,388 Intel Xeon "Haswell" processor nodes, and the second one is made of 9,688 Intel Xeon Phi "Knight's Landing" (KNL) nodes. In this work, I use the "Haswell" partition. Each Haswell node has two sockets, each of them populated with a 2.3 GHz 16-core Haswell processor, therefore allowing the use of 32 cores per node. Each core supports 2 hyper-threads, and has two 256-bit-wide vector units, allowing a theoretical peak of 36.8 Gflops/core. In terms of available memory, each node has a 128 GB DDR4 2133 MHz memory.

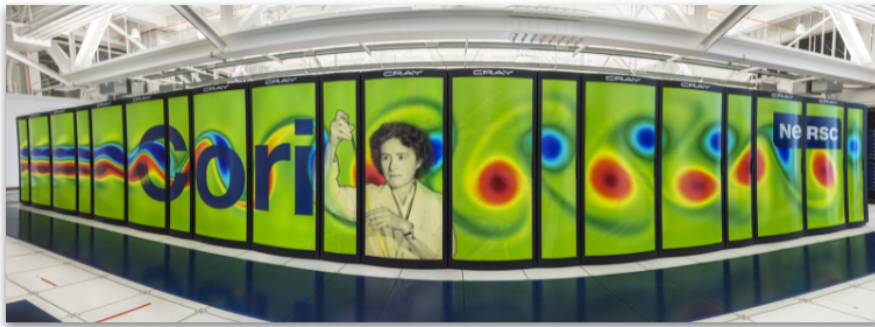


Figure 7.5 – The Cori machine at NERSC. *Credit: NERSC.*

- **Joliot-Curie:** This is a Bull Sequana machine based at TGCC³. The super-computer features four partitions, however, in this work I use a single one; the "SKL" partition, which is made of 1,656 nodes, each populated with 2.7 GHz Intel Skylake 8168 bi-processors, each processor hosting 24 cores, for a total of 79,488 cores and a total peak performance of 6.86 Pflops. Each node has a 192 GB DDR4 memory.

² National Energy Research Scientific Computing Center: <https://www.nersc.gov>

³ Très Grand Centre de calcul du CEA (The French Alternative Energies and Atomic Energy Commission Computing Center): <http://www-hpc.cea.fr/fr/complexe/tgcc.htm>

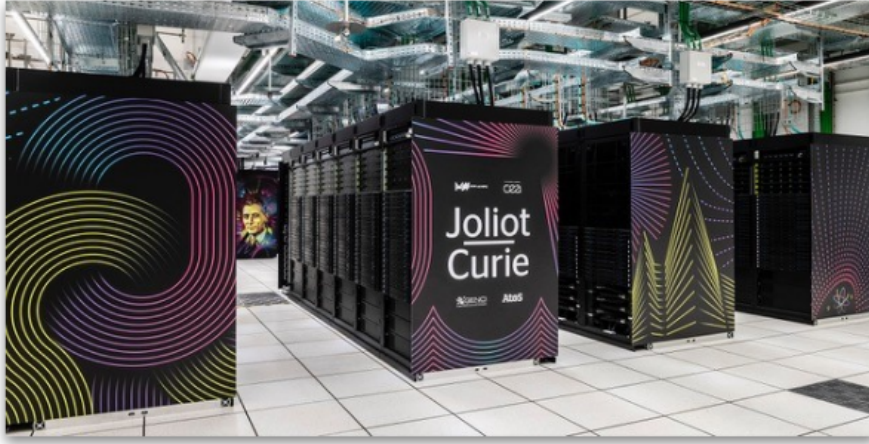


Figure 7.6 – The Joliot-Curie machine at TGCC. *Credit: Atos.*

7.5.1 Simulations description

I simulate CMB observations using the TOAST⁴ framework. In this section, I describe the results obtained with TOAST simulations following the Simons Observatory (SO) specifications. I simulate two sets of observations with the Small Aperture Telescope (SAT) in the 150 GHz channel, with 6,020 detectors:

- **S1:** corresponds to the observation of a $20 \times 10 \text{ deg}^2$ sky patch defined by the following sky coordinates: DEC-050 .. -030, RA+011.613 .. +023.226. The observations consist of 33 constant elevations scans (CESs), each between 45 min and 1 h long, totaling to about 24 hours of observation. The sampling rate is set to 132 Hz, yielding a total of $\mathcal{O}(10^{11})$ time samples, i.e. ~ 10 M time samples per detector. I use this simulation to evaluate the scaling performance of the software.
- **S2:** corresponds to the observation of a smaller sky patch of $10 \times 10 \text{ deg}^2$ defined by the following sky coordinates: DEC-035 .. -045, RA+040 .. +050. The observations consist of 95 CESs, spanning 5 consecutive days, each about 15 min long, for a total of 24 hours of observation. The sampling rate is set to 37 Hz, yielding a total of $\mathcal{O}(10^{10})$ time samples, i.e. ~ 1 M time samples per detector. I use this simulation to study the convergence of the two-level preconditioners.

The scanning speed is 1 deg/s in both simulations, and the half wave plate rotation frequency is set to 2 Hz. The input map of the simulations is pixelized in HEALPix format [214] with a resolution corresponding to $n_{\text{side}} = 512$, and includes a CMB realization from best-fit Planck cosmology and a Gaussian lensing potential. In addition, I simulate instrumental $1/f$ noise and atmosphere. The resulting sky patch from **S1** counts around $\sim 250,000$ pixels, while the one from **S2** counts about $\sim 140,000$ pixels. In addition, we simulate instrumental $1/f$ noise and atmosphere. The instrumental noise has a characteristic knee frequency of 50 mHz (f_k in Eq. (7.1)), a slope of ~ -1 (α in Eq. (7.1)), and a detector NET of $400 \mu\text{K}\sqrt{s}$. The native atmospheric simulation built in the TOAST framework

⁴ <https://github.com/hpc4cmb/toast>

is based on a 3d turbulence model [183] sensitive to parameters such as wind properties, weather conditions, and turbulence characteristics, calibrated on real data from experiments in Chile such as POLARBEAR [280] and ACT [450]. A 3d volume is generated and propagated according to the wind direction and speed as it is scanned by the detectors, thereby generating correlations between time samples as well as cross-correlations across the focal plane between detectors.

7.5.2 Strong scaling tests and collective communication performance

I perform several strong scaling tests, using the **S1** simulation on Cori and Joliot-Curie to demonstrate the ability of the code to run on a large number of cores and handle the data reduction of a large data set. Given that the Skylake processors (Joliot-Curie) are more powerful than the Haswells (Cori), we expect to see longer runtimes with the second machine for the same set of the sky reconstruction parameters. In the runs presented here, I choose the maximum-likelihood map-making approach with a Toeplitz-structured noise covariance and a half-bandwidth set to 2,048 samples. This corresponds a time-domain correlation length of ~ 15 s, given the adopted sampling rate. The solver converges in around ~ 1200 iterations for the case considered. Given that the two machines had different versions of TOAST installed, the noise simulations were not exactly the same resulting in slight changes in the total number of iterations before convergence between the two machines. However, these changes do not affect the computational performance metric adopted here, which is the mean time per iteration. I choose this metric because it depends only on the half-bandwidth and the total size of the problem which are exactly the same for all the tests performed on both machines. The results are shown in figure 7.7.

The `Ring` runs do not scale properly as we increase the number of MPI tasks handling the problem, and completely diverge when using a large number of processes, therefore we exclude it from any further analysis. The customized `Butterfly` scheme exhibits better mean times per iteration compared with the standard `MPI_Allreduce` on the Cori machine, particularly as we increase the number of MPI processes. On the Joliot-Curie machine, the `Butterfly` scheme shows similar or slightly better performance than the standard `MPI_Allreduce` depending on the number of MPI tasks. To get a more quantitative view of the scaling performance, we measure the scaling efficiency given by

$$\varepsilon = \frac{\delta t_1}{N \times \delta t_N}, \quad (7.10)$$

where δt_1 is the reference mean time per iteration to which the scaling efficiency is compared, theoretically this should correspond to the iteration time in a serial run (with a single process), and δt_N is the mean time per iteration for a run with N MPI processes. In practice, the size of the problem does not fit into a single node and hence we cannot perform serial runs, hence we set the reference to be the run with the lowest number of MPI processes, and N to the factor by which we increase the number of processes with respect to the reference. In these tests, this corresponds to 396 MPI processes for the `MPI_Allreduce` tests on Joliot-Curie and to 512 MPI processes for the remaining cases. The change of reference

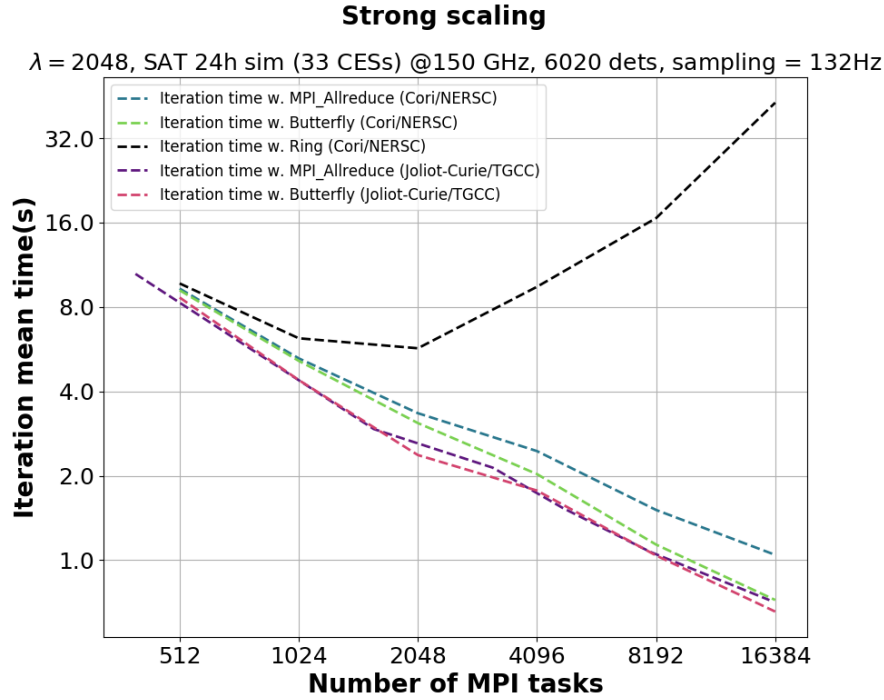


Figure 7.7 – Strong scaling of the mean time per iteration comparing 5 sets of runs: (1) runs executed at the Cori machine at NERSC using the `MPI_Allreduce` communication scheme; (2) runs executed at the Cori machine at NERSC using the `Butterfly` communication scheme; (3) runs executed at the Cori machine at NERSC using the `Ring` communication scheme; (4) runs executed at the Joliot-Curie machine at TGCC using the `MPI_Allreduce` communication scheme; (5) runs executed at the Joliot-Curie machine at TGCC using the `Butterfly` communication scheme.

point should not affect our conclusions: given the proximity between the two references one can reasonably assume that the scaling between the two is close to 100%. Results are shown in figure 7.8.

The `Butterfly` scheme performs better than `MPI_Allreduce` by a few percents on both machines. The scaling efficiency at ~ 16000 processes reaches around 40% for the `Butterfly` scheme, while it's below 30% for `MPI_Allreduce` on Cori. Nevertheless, in the considered cases and for the number of MPI processes used, the communication overhead is small as compared to the computation time, and the current implementation of the `Butterfly` scheme is restricted to numbers of processes which are powers of 2. While the algorithm can be straightforwardly extended, this requires more work, and is left for future extensions of the software. Instead, given the satisfactory performance of the standard `MPI_Allreduce` scheme, I select it as the default option for the rest of this work. I point out that both schemes, `MPI_Allreduce` and `Butterfly`, show good scaling for these types of applications on both machines, demonstrating good overall portability of the software.

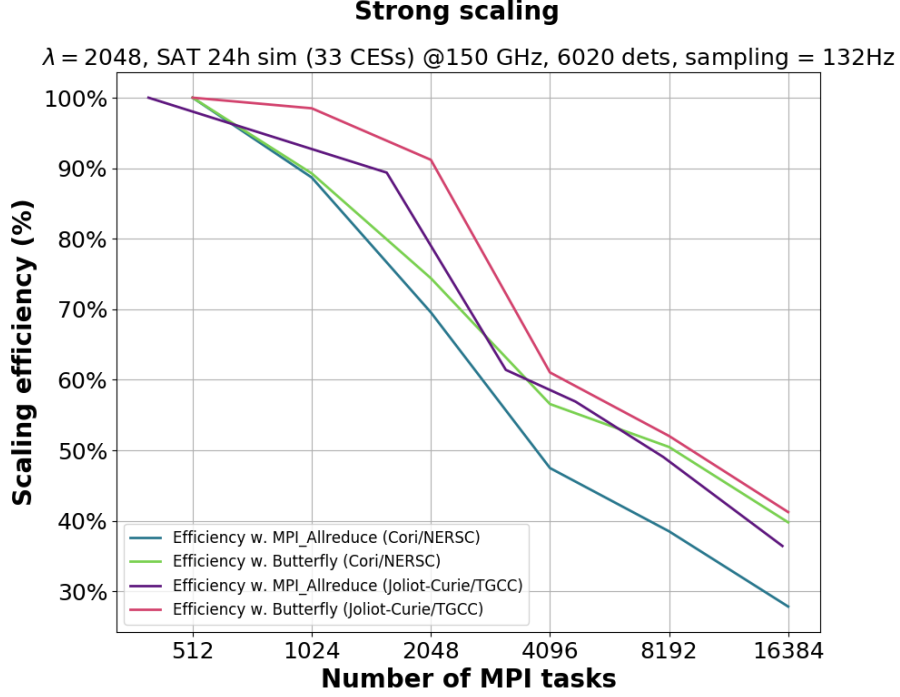


Figure 7.8 – Strong scaling efficiency comparing 4 sets of runs: (1) runs executed at the Cori machine at NERSC using the `MPI_Allreduce` communication scheme; (2) runs executed at the Cori machine at NERSC using the Butterfly communication scheme; (3) runs executed at the Joliot-Curie machine at TGCC using the `MPI_Allreduce` communication scheme; (4) runs executed at the Joliot-Curie machine at TGCC using the Butterfly communication scheme.

7.5.3 Large run demonstration

In this subsection, I present the largest run I performed to date with the MAP-PRAISER software. I simulate observations of the Small Aperture Telescopes with TOAST following a scan strategy constructed with the opportunistic scheduler [212]. The constraints are a minimum observing elevation of 55 deg, a Sun and Moon avoidance radius of 45 deg, and targeted tiles of 10×20 deg (RA,DEC), overlapping by 5 deg in right ascension. The resulting schedule has an observing efficiency of 64.8%, with an f_{sky} of about 35%, and all observing elevations are below 63 deg. It has been adopted in the mission scale simulations of the Simons Observatory, used to prepare the science pipelines. In this demonstration, I select the first 333 CESs of this schedule. Similarly to the **S1** simulation, each CES is between 45 and 1 h long. I process data from a single frequency channel, the 150 GHz band, corresponding to 6,020 detectors. The sampling rate is set to 132 Hz. The total size of the problem is about $\mathcal{O}(10^{12})$ time samples. The nside of the input map is set to 512, corresponding to a total number of pixels in the observed sky area of ~ 1 M pixels.

I perform the simulation and the data reduction on the Joliot-Curie machine, allocating 47,952 cores. Each core is processing around 16 M samples. The input signal contains CMB, instrumental as well as atmospheric noise. I choose the Toeplitz-structured noise covariance for the map reconstruction, with a half-bandwidth $\lambda = 16,384$ time samples, corresponding to ~ 2 min of correlations.

The solver stops after 1,600 iterations reaching a relative residual of 10^{-6} . Figure 7.9 shows the evolution of the relative residuals as a function of the wall clock time. The total run time is about 5,300 s, thereby consuming a total of 70,596 CPU-hours. The mean time per iterations is ~ 3.24 s, demonstrating that the scaling of the code does not significantly degrade even with a very large number of cores. The polarization maps obtained in this run are shown in figure 7.11. Figure 7.10 shows the hits map of the zoomed in $5 \text{ deg} \times 5 \text{ deg}$ sky patch. The number of hits exceeds 10^5 in all pixels located in that area. By selecting pixels for which the number of hits is above 10^5 , the obtained noise RMS of the Q and U maps is $2.22 \mu\text{K}$.

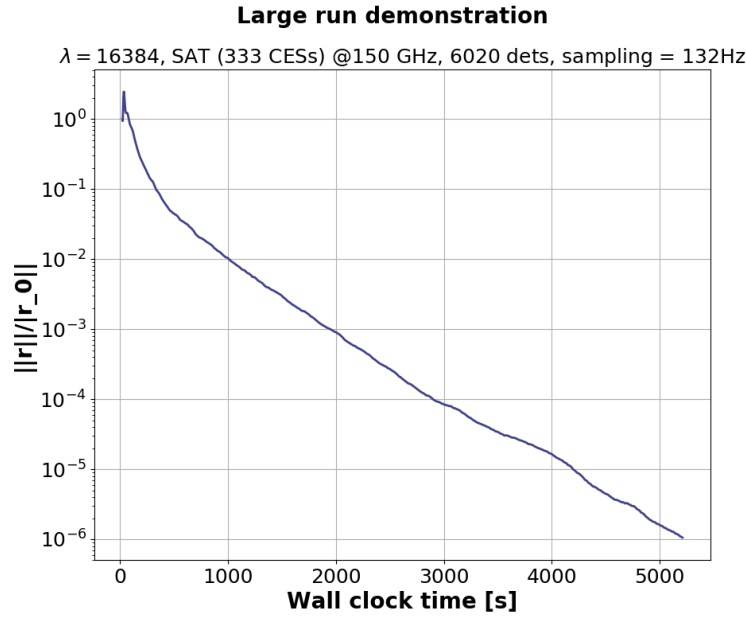


Figure 7.9 – Convergence curve showing the relative residuals as a function of the wall clock time.

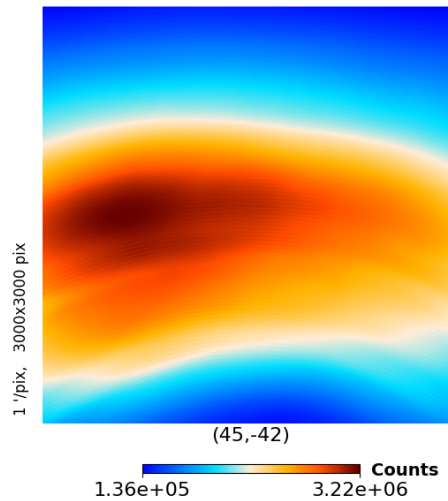


Figure 7.10 – Hits map of a $5 \text{ deg} \times 5 \text{ deg}$ sky patch located at 45 deg longitude and -42 deg latitude.

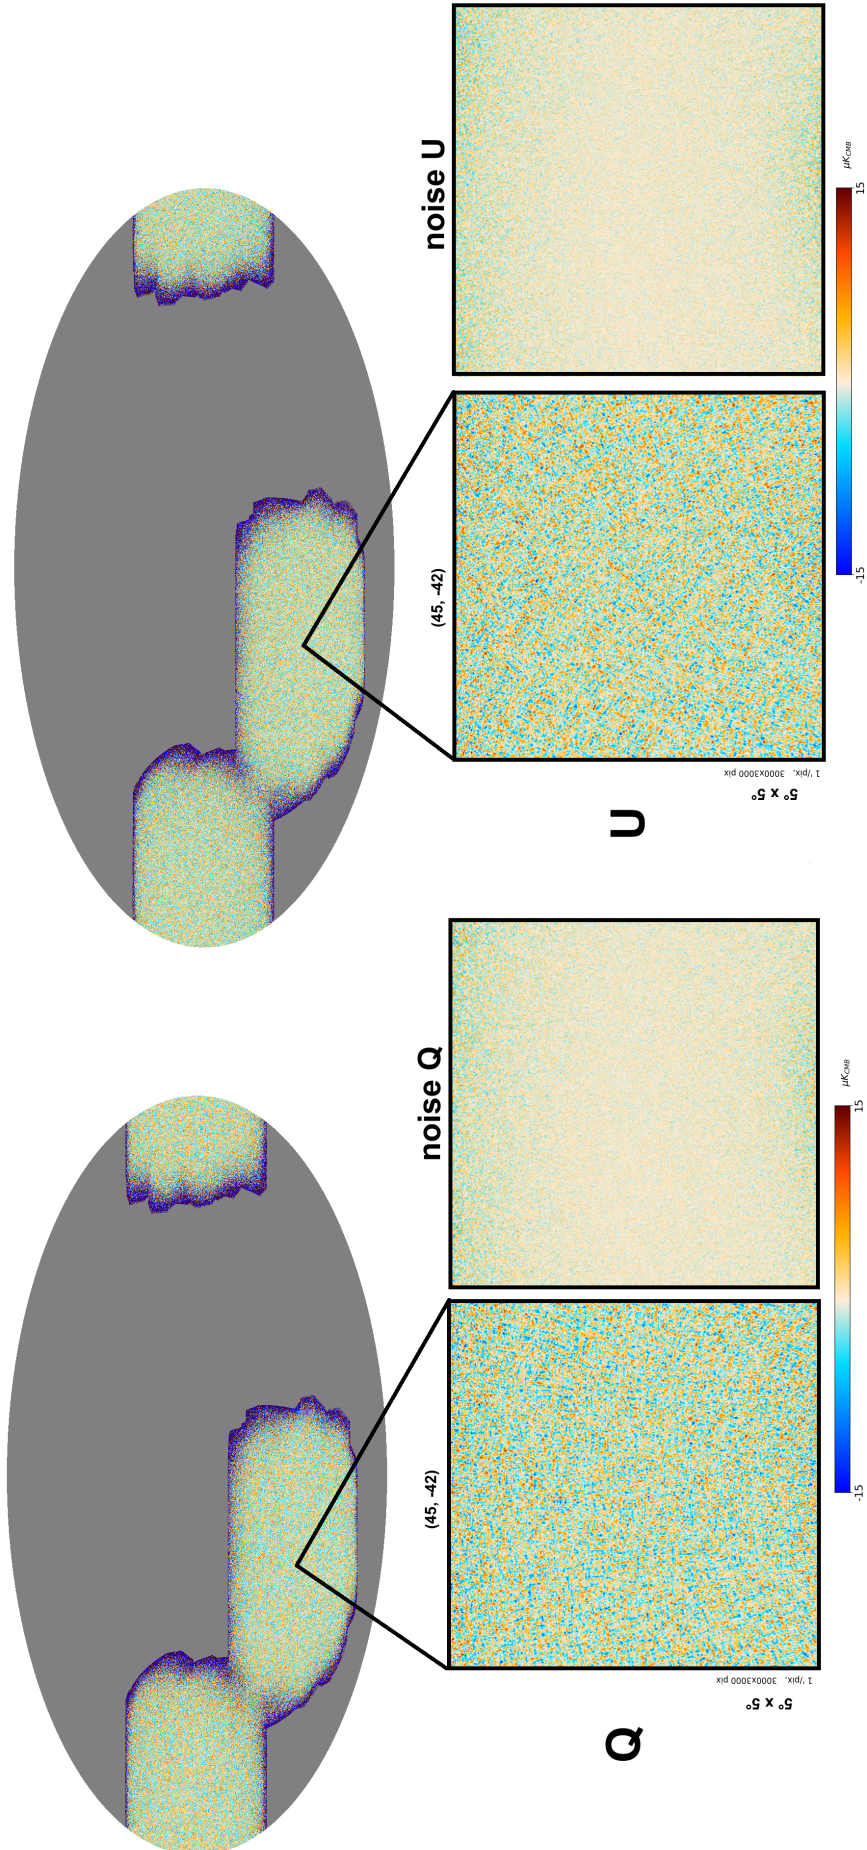


Figure 7.11 – Q and U polarization maps obtained from the large run demonstration.

7.5.4 Solvers and preconditioners performance

In this last section, I focus on evaluating the efficiency of the different linear solvers techniques discussed previously. I use the simulated data to evaluate their convergence properties and discuss their advantages and limitations.

Enlarged CG

In figure 7.12, I show a selection of tests of the ECG solver for different enlarging factors. The data used here constitute a 1 h subset of the **S1** simulation corresponding to a single constant elevation scan. I also exclude the atmospheric noise from the input, and artificially increase the instrumental noise f_{knee} to 1 Hz. I assume a half bandwidth, $\lambda = 2^{18}$, corresponding to a correlation length of about 33 minutes. This configuration results in a degraded conditioning of the system matrix showcasing the ability of the ECG solver to deal with such cases. I execute multiple ECG runs varying the enlarging factor from 1 to 16, and compare it to a standard PCG run. In the case of a unity enlarging factor, the ECG solver should behave exactly like the PCG, as the two solvers are theoretically equivalent. This provides an additional way to validate the implementation of ECG, in addition to basic checks with noiseless simulations and binned maps. I verify that this is indeed the case as the two convergence curves coincide as shown in figure 7.12.

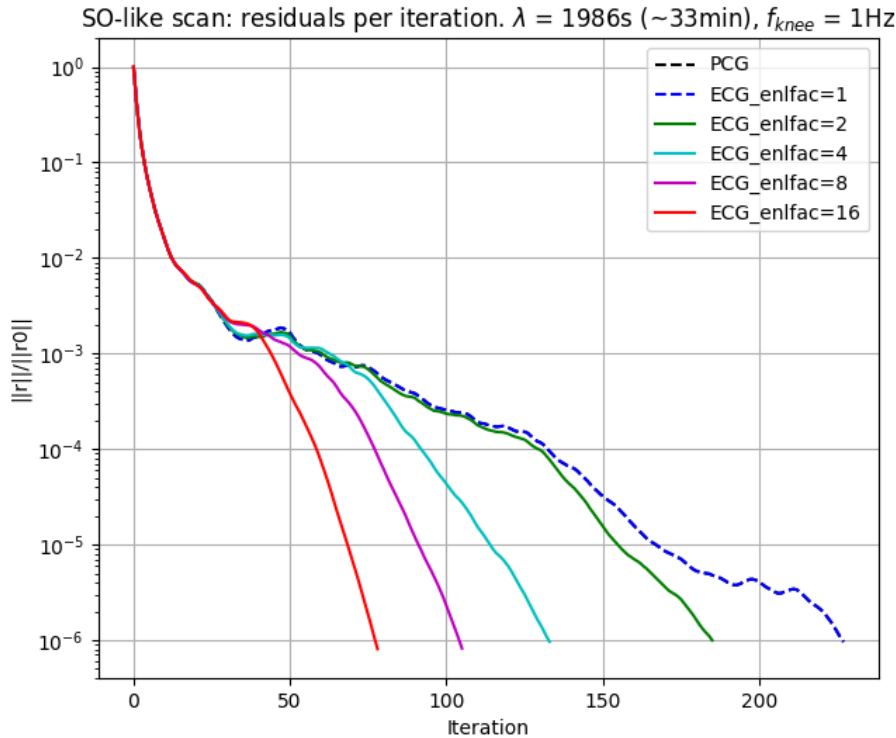


Figure 7.12 – Scaling of ECG iterations with the enlarging factor, using 1h of SO-SAT simulated data.

In addition to this, we observe that the number of iterations required to reach convergence (defined in the present case by a tolerance parameter of 10^{-6}) is monotonically decreasing with the enlarging factor. In particular, with an enlarging factor of 16, we gain by about a factor 3 with respect to standard PCG. This

behavior is also theoretically expected from Eq. (7.4). However, as previously mentioned this does not translate automatically to a net gain in the run time. For that purpose, one should ensure that the implemented algebraic routines used by the solver are sufficiently optimized to compensate for the increase in the cost of each iteration. This is not the case for the naive implementation I used in this test. The objective here is to provide a proof of concept that such solver can be applied in the map-making applications and provides sufficient speedup of the convergence rate, to prepare for future work directly targeting the optimization of the underlying routines for a net gain in the run time.

PCG with two-level preconditioners

In the present subsection, I compare the performance of the three preconditioners provided by MAPPRASER. For this purpose, I run an adapted version of the **S2** simulation on Cori, where I cut down the number of detectors by a factor 14, i.e. to 430 detectors, to run on a small number of processes (64 MPI tasks) and I scale down the noise to emphasize the sky signal and study the angular convergence as well. The convergence properties should not be affected by such changes as they don't depend on the noise normalization in general. I use the block-diagonal (BD) preconditioner, the *a priori* and *a posteriori* two-level preconditioners with various sizes of the deflation space, and compare the performance in terms of the number of iterations and wall clock time necessary to converge (figures 7.13, 7.14 and 7.16). In addition, I also show the TT and EE angular power spectra of the reconstructed maps at all iterations to assess the convergence of the different angular scales (figures 7.17 to 7.24). The BB spectra are noise dominated in this regime and therefore not considered here. The power spectra are estimated using the NaMaster library⁵ [56], correcting for the mask, however not for the noise bias. The *a priori* preconditioner tests summarized in figure 7.13 show only marginal gain on the number of iterations to convergence, demonstrating that, at least in the case studied, there is no significant benefit in using the *a priori* preconditioner over the standard BD preconditioner. In fact, the *a priori* preconditioner is included mostly for historical reasons. Indeed, the main early interest in this preconditioner was driven by the possibility of saving on the precomputation time needed to construct \mathbf{Z} [438, 217], which is essentially guessed and therefore computationally cheap. However, constructing the preconditioner still requires a computation of the $\mathbf{A}\mathbf{Z}$ matrix. As I show in the following for the novel *a posteriori* construction I propose in this thesis, the computation of \mathbf{Z} is as costly as that of $\mathbf{A}\mathbf{Z}$ and the cost of the latter can be hidden in the calculation of the former. Consequently, this new class of the *a posteriori* constructions is not only more performant but also requires less precomputations. Notwithstanding these considerations, the *a priori* construction may still be of interest, in particular, in the cases where there is limited number of nearly singular eigenvectors spanning a known low-dimensionality subspace. In the following, I focus solely on the convergence properties of the *a posteriori* preconditioner

Figure 7.14 shows that, for the *a posteriori* preconditioner, the number of iterations required for convergence monotonically decreases with the size of the deflation subspace, $\dim \mathbf{Z}$. For any sufficiently high $\dim \mathbf{Z}$ (here $\gtrsim 128$), the *a posteriori*

⁵ <https://github.com/LSSTDESC/NaMaster>

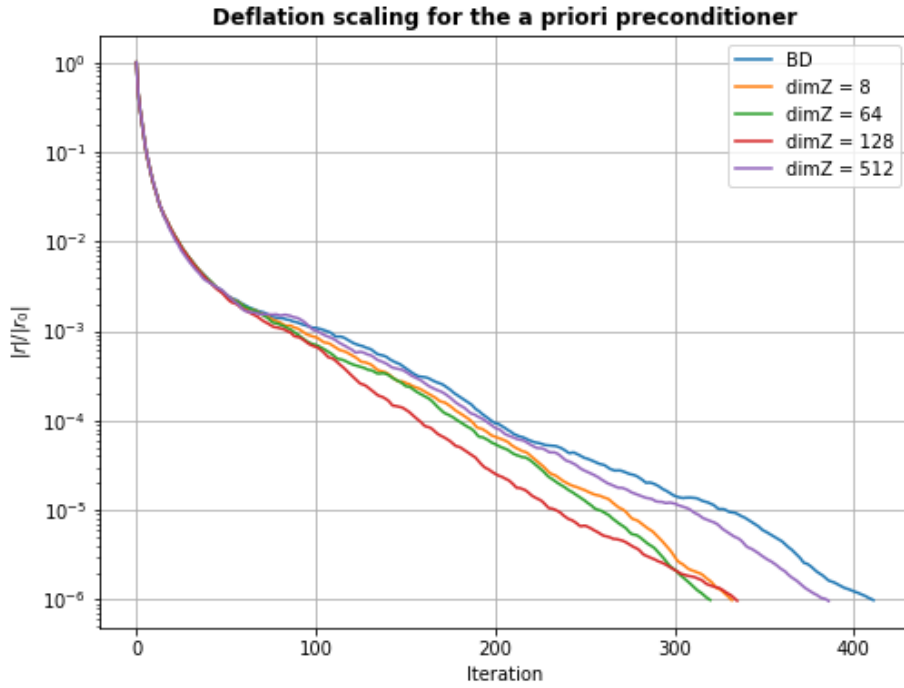


Figure 7.13 – Convergence in terms of the number of iterations of the BD preconditioned CG, compared with the *a priori* two-level preconditioned CG, for different dimensions of the deflation subspace.

two-level preconditioner performs consistently better than the *a priori* two-level preconditioner. This is expected since the former uses a more involved procedure to construct the deflation subspace, yielding better estimates of the eigenvectors of the system matrix, once the Lanczos procedure runs for a sufficient number of iterations.

Additionally, the convergence gain as shown in the figure mainly stems from a sharp drop of the residual at the very first CG steps of the solver. This can be understood by examining the convergence of the different angular scales in figures 7.17 to 7.24. The colorbar encodes the wall clock time from the start of the solver until convergence is reached, not including precomputation time. The colors are given in log scale and the curves correspond to the power spectra at different iterations. Comparing the various TT and EE plots, we can see that the higher the size of the deflation subspace, the wider the range of the angular scales recovered from the first CG step. In fact, the precomputation allows the solver to compress information encoded in the system matrix, \mathbf{A} , to a small set of modes contained in the deflation matrix, \mathbf{Z} . When running the PCG steps, that information allows to recover the smallest angular scales immediately in the first step, while subsequent steps are spent correcting the largest angular scales. For deflation subspaces of high dimensions, $\dim \mathbf{Z} \geq 512$, we can see that the power spectra of the different iterations are very close to each other and to the converged curve, hence in these cases we practically recover the full range of angular scales of the solution from the first step. In particular, the case $\dim \mathbf{Z} = 2048$ converges in two steps, meaning that we compress practically all of the information of the large system matrix, \mathbf{A} , in the relatively low number of modes of the deflation subspace. The *a posteriori* two level preconditioner thus not only improves con-

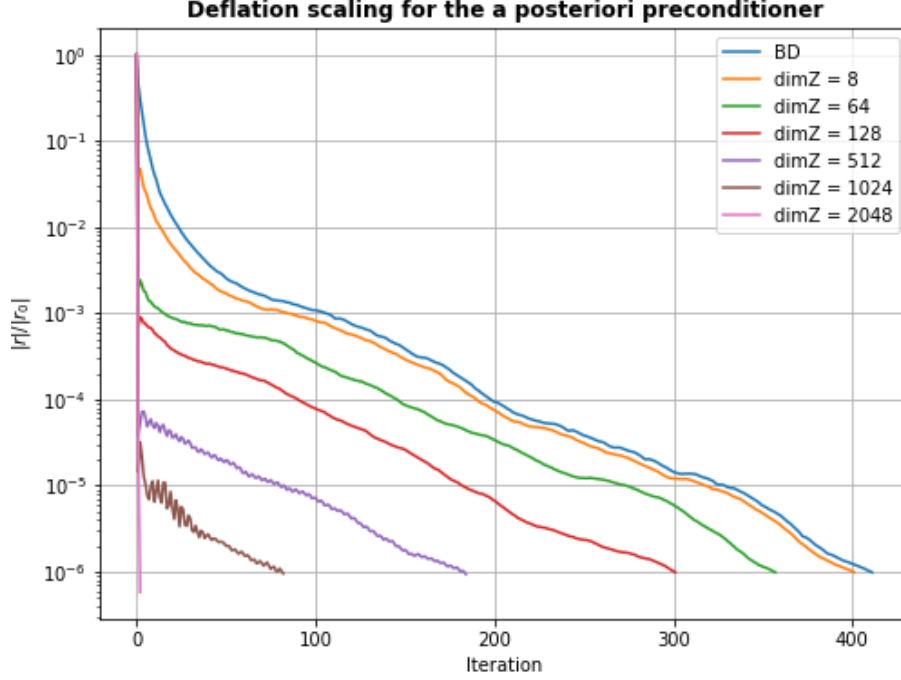


Figure 7.14 – Convergence in terms of the number of iterations of the BD preconditioned CG, compared with the *a posteriori* two-level preconditioned CG, for different dimensions of the deflation subspace.

vergence, it also gives us some insight into the noise covariance structure of the computed maps. However, for $\dim \mathbf{Z} = 512$ or 1024 , while the spectra reach their final shape in at most a couple of steps, the iterations continue for nearly 100 or 200 steps more in order to attain the formal convergence. This is related to the fact that the residual used to test for convergence is weighted towards the noisiest pixels often located at the outskirts of the observed sky patch and which matter little in the power spectrum computation. This suggests that if the power spectra are of ultimate interest for the analysis, a more suitable convergence metric could be proposed leading to significant performance gain.

The cost of precomputation for different sizes of the deflation subspace constructed with the Lanczos procedure is shown in the left panel of figure 7.15. This cost scales linearly with $\dim \mathbf{Z}$, and can be as costly as a few PCG runs. This makes this method mostly advantageous when we want to solve multiple systems, with roughly the same system matrix but different realizations of the right hand side, as it is the case for example when performing null tests, or Monte Carlo simulations. Then the preconditioner needs to be constructed only once, stored in memory, and all subsequent runs can benefit from the same convergence boost virtually for free. The cost of each iteration does not substantially change in a way that may impact the net gain in runtime as shown in figures 7.15 (right panel) and 7.16.

The slight increase in the average iteration cost seen for the cases $\dim \mathbf{Z} = 1024$ or 2048 , is due to the fact that operations start getting non negligible contributions from the $\dim \mathbf{Z}$ map-domain operations. However, for the case considered, this increase does not affect the results in any significant way, but may be of concern for high resolution maps, in the context of LAT observations for instance.

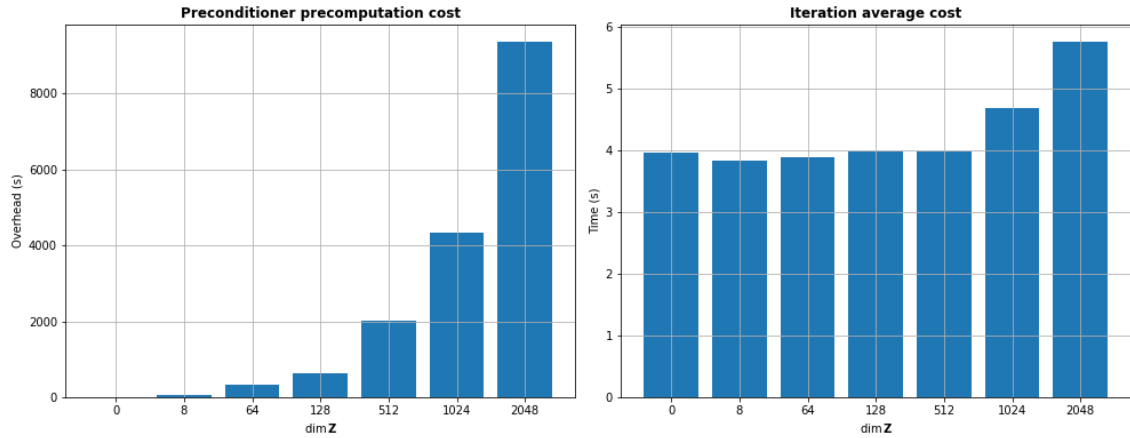


Figure 7.15 – *Left*: Precomputation cost of the *a posteriori* preconditioner as a function of the dimension of the deflation subspace matrix $\dim Z$. In practice this dimension is the number of iterations performed by the Lanczos procedure. *Right*: Average iteration cost with the same preconditioner as a function of the dimension of the deflation subspace matrix $\dim Z$.

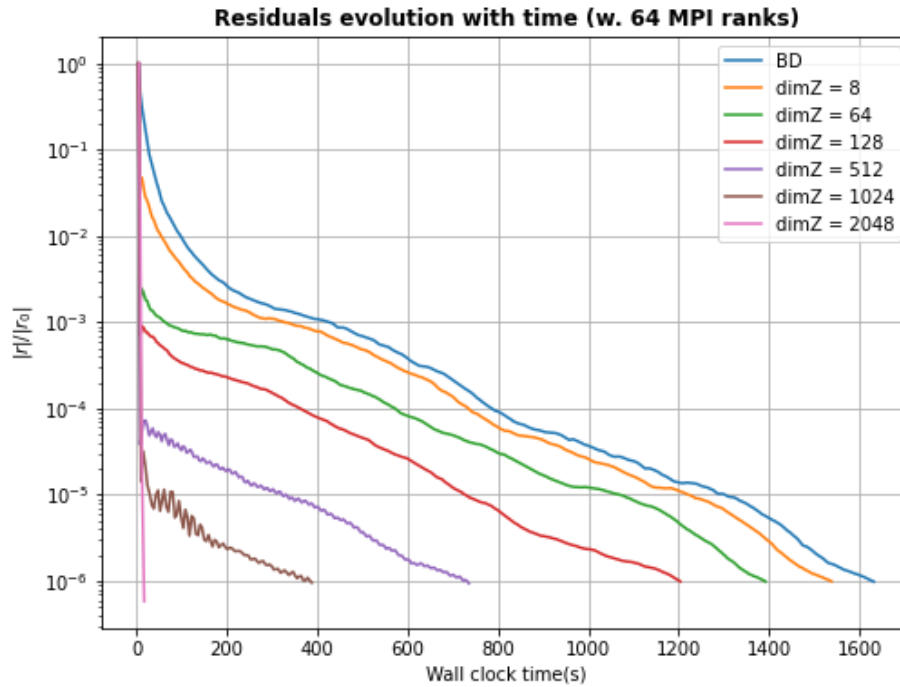


Figure 7.16 – Convergence in wall clock time of the BD preconditioned CG, compared with the *a posteriori* two-level preconditioned CG, for different dimensions of the deflation subspace. The precomputation time is not included here.

In this case, memory space is also a limiting factor and one may be forced to downsample the columns of Z for the method to work. In such a case, the small scale modes would not be captured by the deflation subspace, but these typically converge very well regardless. This is however beyond the scope of the present work, and may be the object of future investigation.

Finally, we can consider some numbers to illustrate the effective gain in run-

time enabled by this technique. In the present case studied, the precomputation cost with $\dim \mathbf{Z} = 2048$ is roughly ~ 5.5 PCG runs with the standard block-diagonal preconditioner, and the solving time is roughly 10 s. Solving 100 similar systems for example, would then cost around $\sim 9500 + 10 \times 100 = 10,500$ s, while it would cost around $1650 \times 100 = 165,000$ s running with the block-diagonal preconditioner. We gain about a factor ~ 16 solving 100 similar systems. Note that the gain is only going to increase as we increase the number of similar systems to solve, hence alleviating the cost of map-making in Monte Carlo runs for example, as we become only limited by the cost of the simulations.

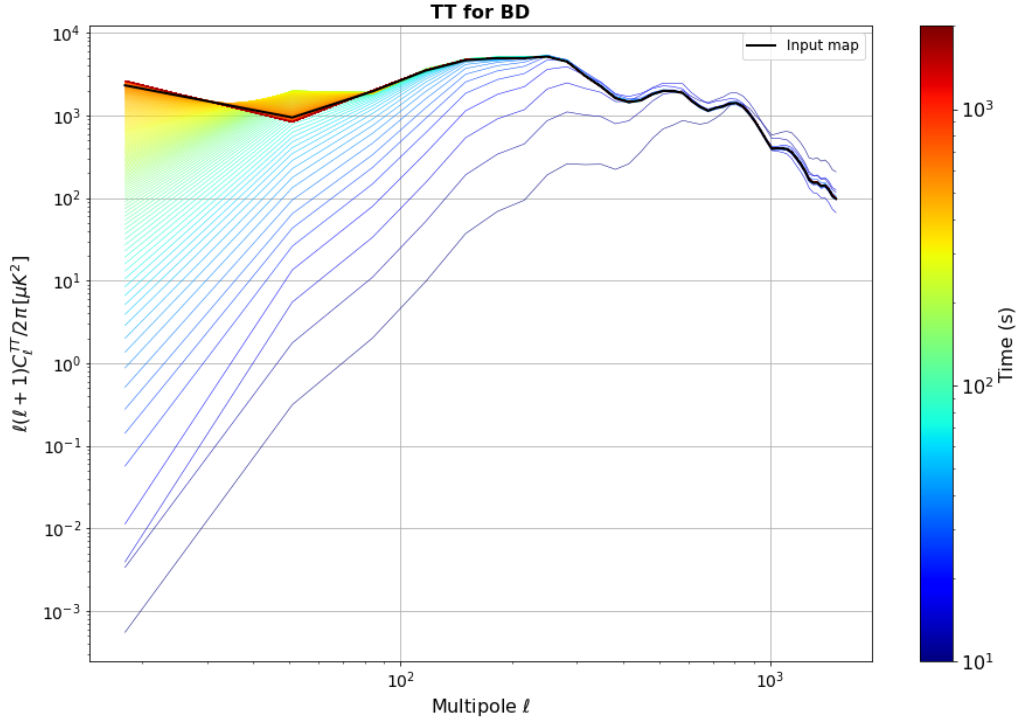


Figure 7.17 – Temperature angular power spectra of the maps at different stages of convergence of the BD preconditioned CG.

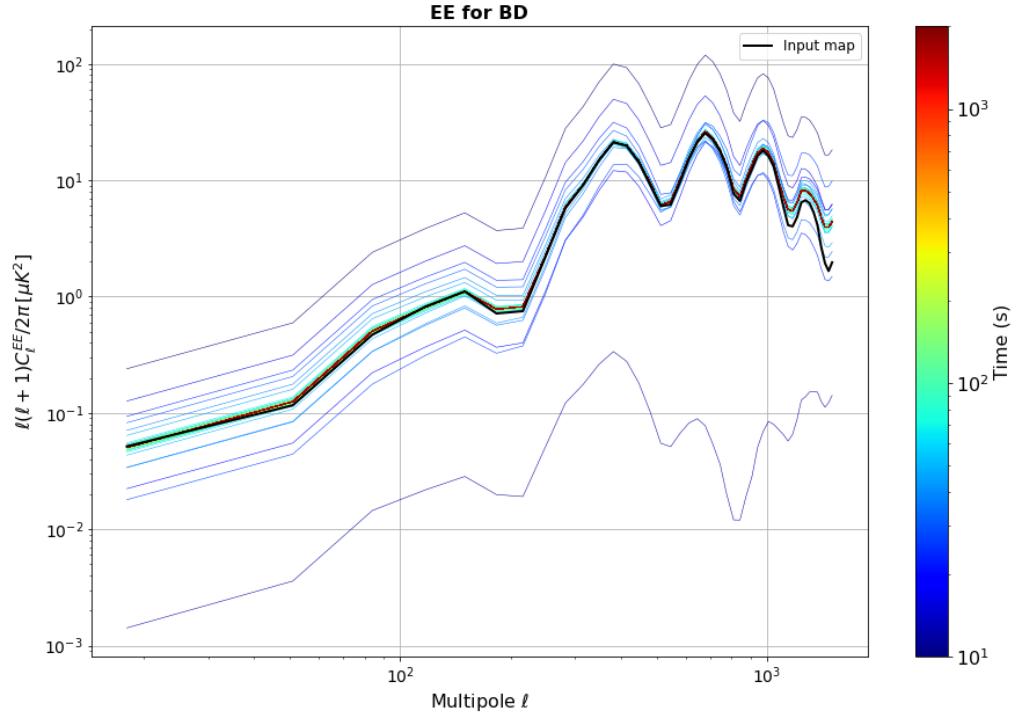


Figure 7.18 – EE angular power spectra of the maps at different stages of convergence of the BD preconditioned CG.

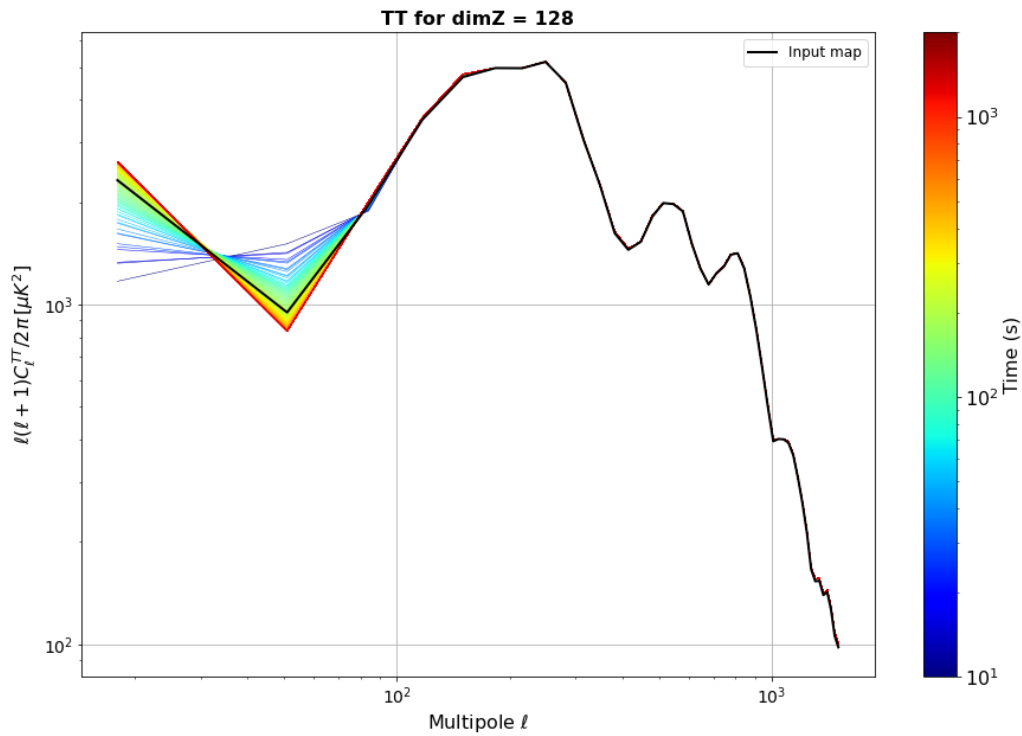


Figure 7.19 – Temperature angular power spectra of the maps at different stages of convergence of the 2lvl-dim $Z = 128$ preconditioned CG.

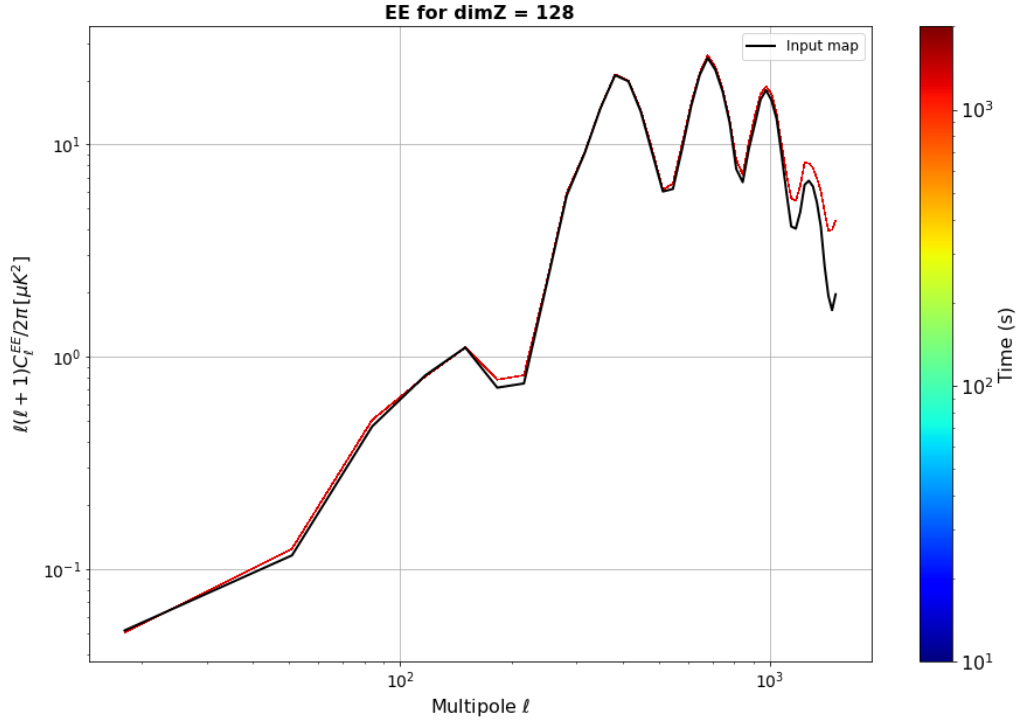


Figure 7.20 – EE angular power spectra of the maps at different stages of convergence of the 2lvl-dim $\mathbf{Z} = 128$ preconditioned CG.

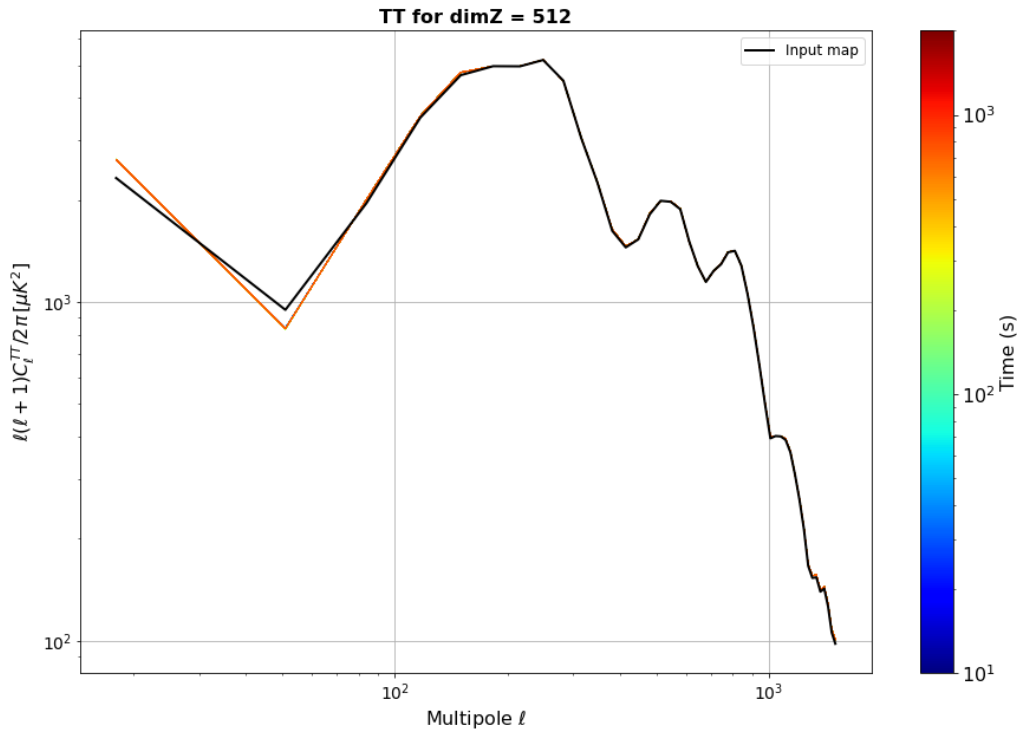


Figure 7.21 – Temperature angular power spectra of the maps at different stages of convergence of the 2lvl-dim $\mathbf{Z} = 512$ preconditioned CG.

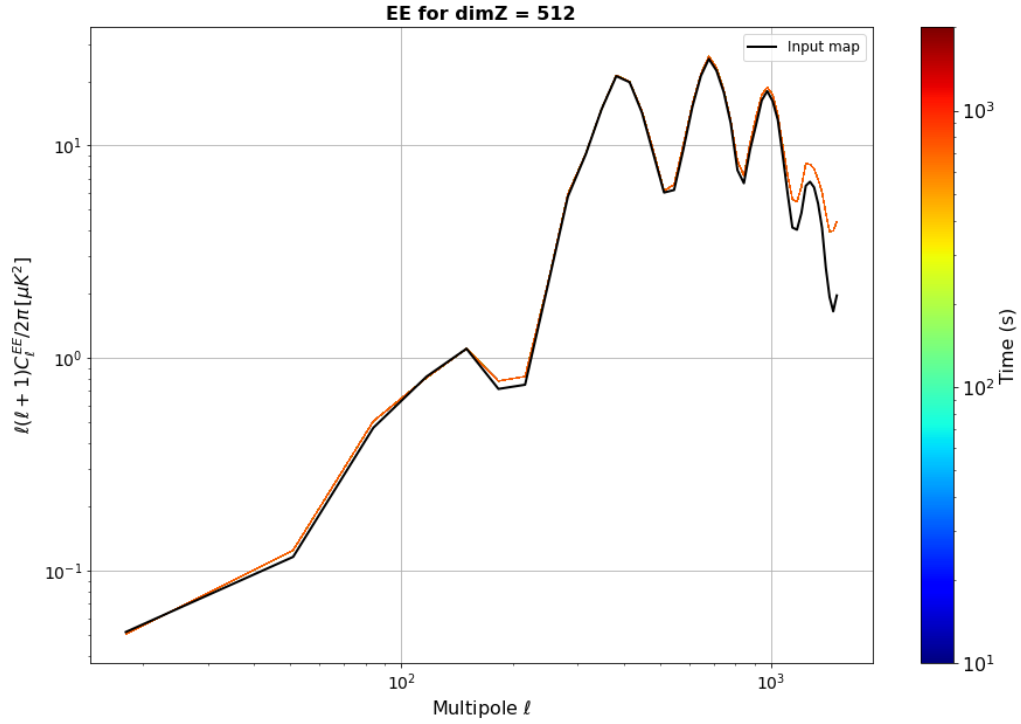


Figure 7.22 – EE angular power spectra of the maps at different stages of convergence of the 2lvl-dim $\mathbf{Z} = 512$ preconditioned CG.

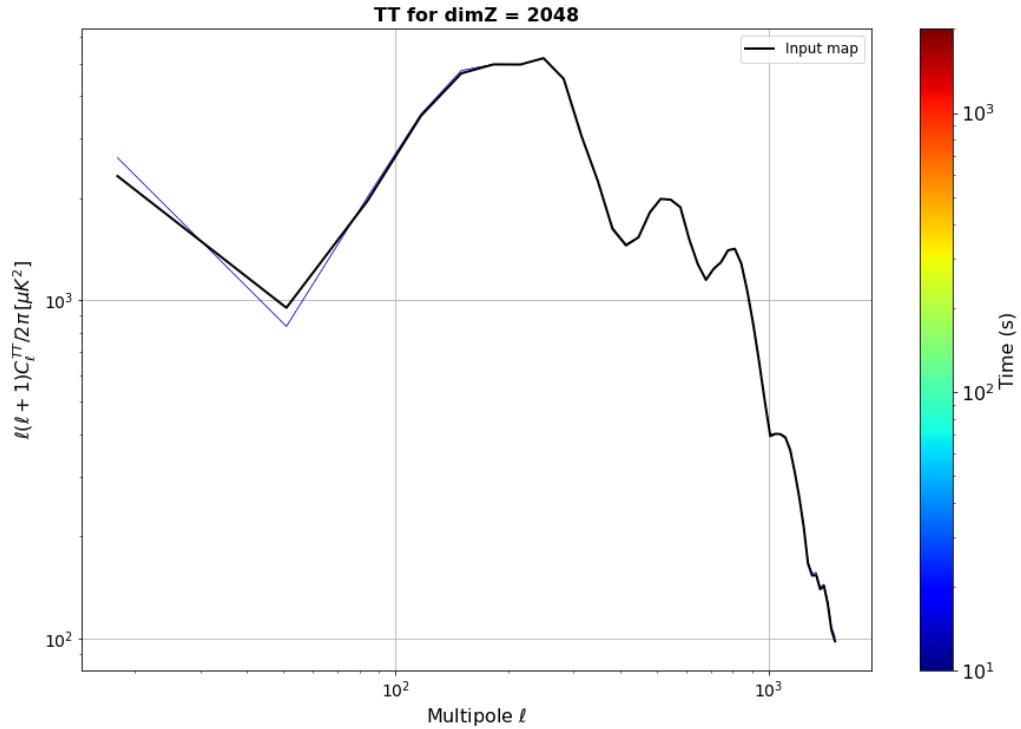


Figure 7.23 – Temperature angular power spectra of the maps at different stages of convergence of the 2lvl-dim $\mathbf{Z} = 2048$ preconditioned CG.

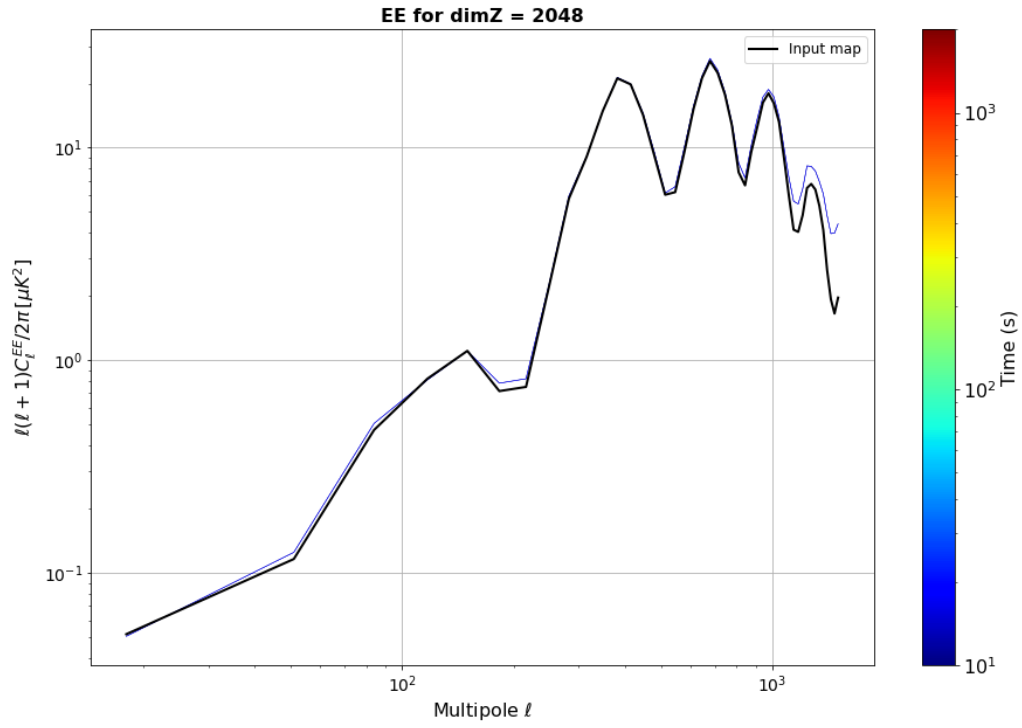


Figure 7.24 – EE angular power spectra of the maps at different stages of convergence of the 2lvl-dim $\mathbf{Z} = 2048$ preconditioned CG.

CHAPTER 8

Map-making applications on satellite and ground-based experiments

Contents

8.1	Satellite experiments: LiteBIRD case	177
8.1.1	Simulation description	177
8.1.2	Full-sky maximum likelihood maps	178
8.1.3	Full-sky templates marginalization maps	180
8.2	Ground-based experiments: Simons Observatory, Small Aperture Telescopes case	182
8.2.1	Simulation	183
8.2.2	Validation: Scan & HWP synchronous signals mitigation	183
8.2.3	Atmospheric effects	187

In this last chapter, I validate and explore the behavior of the map-making methods I developed, using simulated data from both satellite and ground-based experiments. The satellite simulations include only $1/f$ instrumental noise, allowing to compare the different methods in simple and completely controlled settings. The simulations of ground-based observations will first serve to validate the templates marginalization approach for the case of the scan synchronous and HWP synchronous signals by injecting toy models of these systematics. Later, I focus on studying the effects of atmosphere on both the quality of the recovered maps and the convergence of the map-makers. For this purpose, I consider multiple mapping approaches and assess their performance with respect to different parameters used in the map reconstruction. I also investigate the atmospheric intensity-to-polarization leakage induced by calibration errors and its mitigation with the HWP.

8.1 Satellite experiments: LiteBIRD case

In this first section, I demonstrate the two map-making approaches developed in the previous chapters, on satellite observations with data volumes comparable to what is anticipated in the actual data analysis of a satellite mission such as LiteBIRD. I conduct some basic analysis and comparisons to verify that our map-makers behave properly in simple settings without the presence of complex systematic effects.

8.1.1 Simulation description

I simulate full sky observations with TOAST assuming the LiteBIRD instrumental configuration [240]. The simulation features one year of observations at several frequency bands with the LFT, MFT and HFT instruments (Cf. Section 5.1.3). The input sky includes CMB and galactic signal from the *d1s1* foreground model of PySM [449]. The *nside* is set to 256, i.e. the total number of pixels in the maps is 786,432 pixels. We do not include beams in the simulations. The scan strategy from the second Lagrange Point (L2) is summarized in the diagram in figure 8.1. It is described by the rotation of the satellite around the spin axis, with a period of 20 min, and the spin angle—the angle between the spin axis and the boresight direction—is $\beta = 50$ deg. In addition, the spin axis also precesses around the Sun-Earth axis with a period of 192.348 min, with a precession angle $\alpha = 45$ deg. Finally, we also have the yearly revolution of the Sun-Earth axis around the Sun.

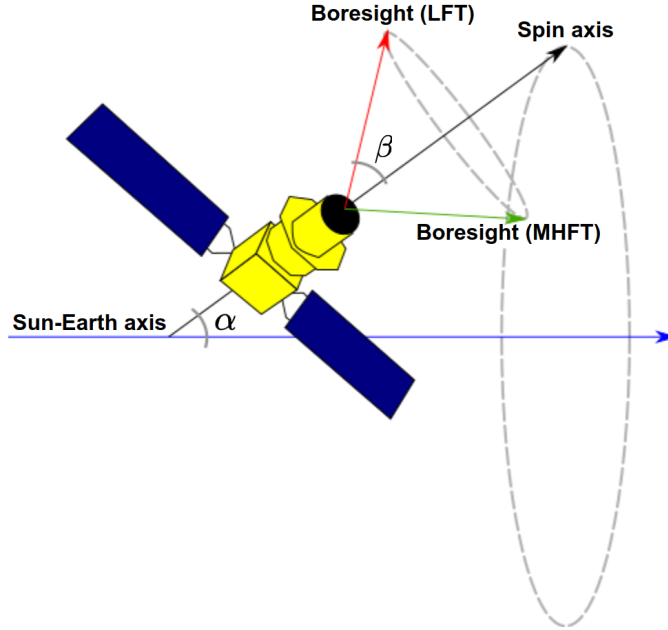


Figure 8.1 – Diagram illustrating the LiteBIRD scan strategy. *Credit: LiteBIRD collaboration^a.*

^a: <https://litebird-sim.readthedocs.io/en/latest/>

The simulation also includes $1/f$ instrumental noise, with characteristic knee frequencies of about 20 mHz (f_k in Eq. (7.1)), and a slope of ~ -1 (α in Eq. (7.1)).

A new realization from the detector noise PSDs is drawn every six hours. The sampling rate is 19 Hz, which means that each detector data amounts to about ~ 600 M time samples. Depending on the instrument and the frequency channel, we need to process data from between 24 and 488 detectors, which means that the size of the data sets generated in the simulations are around $\mathcal{O}(10^{10-11})$ time samples. The half-wave plate rotation frequency is set to 46, 39 and 61 rpm for the LFT, MFT and HFT instruments respectively. Table 8.1 shows the number of detectors and the temperature white noise levels of some selected channels for which results are shown in the following subsections. All runs are performed on the Cori machine at NERSC.

Instrument	Center frequency [GHz]	Ndets	NET array [$\mu K \cdot \sqrt{s}$]
LFT	40	48	18.50
MFT	140	366	3.16
HFT	235	254	5.34

Table 8.1 – Total number of detectors and array noise equivalent temperature (NET) in $\mu K \cdot \sqrt{s}$ for three selected frequency channels of LiteBIRD.

8.1.2 Full-sky maximum likelihood maps

In the simulated data described above, we expect no spatial noise correlations across the focal plane since the instrumental noise is drawn independently for each detector. This is a simplification of the actual case where the minor detector-detector correlation can, and will be, present due to systematic effects, such as read-out induced cross-talk or common thermal modes. In this case, we can get very close to the true maximum likelihood maps by using the banded block-Toeplitz noise model in the map reconstruction. In figure 8.2, I show the Q polarization full-sky maps obtained with, respectively, the LFT instrument at the 40 GHz band, the MFT instrument at the 140 GHz band, and the HFT instrument at the 235 GHz band, processed with a map-maker assuming a Toeplitz-structured noise covariance with a half bandwidth of $\lambda = 8192$ time samples, corresponding to a correlation length of about ~ 430 s. Given the expected proximity of the resulting maps to the optimal solution, I hereafter refer to these as maximum likelihood when comparing them to other maps.

Figure 8.3 shows the noise power spectra of the temperature and polarization maps computed from the difference between the output and input maps of the LFT 40 GHz channel. The underlying maps are obtained with multiple map-making methods: maximum likelihood with the parameters mentioned above, templates marginalization including only polynomial templates with a baseline length fixed to 120 s, and binned map-making. The EE and BB noise spectra are both shown in grey and are consistent with the polarization white noise level expected from detector NETs. They are essentially flat, due to the HWP modulating the signal away from the low frequency $1/f$ noise, hence we do not see any difference in polarization between the different mapping approaches in this case. However, as mentioned previously, the simple configuration considered in

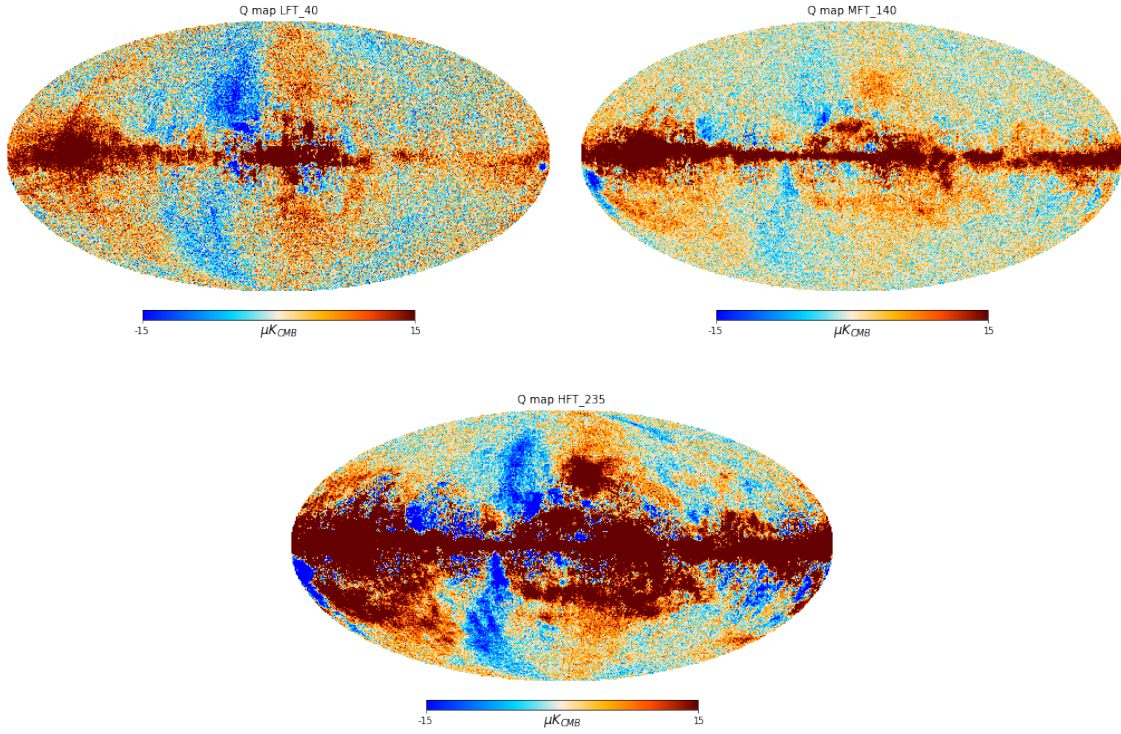


Figure 8.2 – Full-sky maximum likelihood maps from 1 year of simulated observations of the LiteBIRD space mission. The *top left* map corresponds to the LFT 40 GHz channel, the *top right* one to the MFT 140 GHz channel, and the *bottom* map corresponds to the HFT 235 GHz channel.

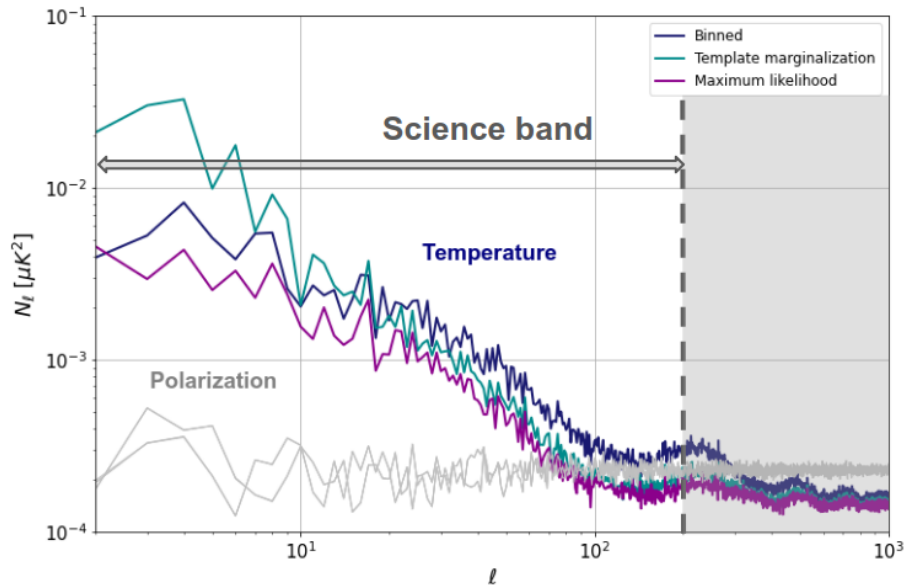


Figure 8.3 – Comparison of the noise power spectra obtained via the different map-making methods available in MAPPRASER. The baseline for the Template marginalization run is set to 120 s. The maps are computed from the simulated data of the LFT 40 GHz channel.

the input simulation is not realistic, and one should expect a more complex noise

structure to arise for the polarization signals when introducing systematics. The TT noise power spectra however show differences between different mapping approaches. In particular, the maximum likelihood estimator yields the least noisy maps as expected. The templates marginalization is closer to the maximum likelihood solution only at small scales, while the noise is boosted at large scales yielding worst estimates than the binned map-maker at $\ell \lesssim 20$. I explore the parametrization and properties of the templates marginalization maps in more details in the next subsection.

It is worth mentioning that an unexpected feature of the TT noise power spectra is the presence of small-scale wiggles possibly indicating the presence of some residual noise correlations at high- ℓ . These are mostly located at sub-beam scales and thus are, for the most part, not really relevant for LiteBIRD science. However, part of the first bump is located at the high- ℓ end of the science band signaling that the temperature noise may not be white at the smallest scales recovered. I have verified that this is not generated by the post-processing of the map, by computing the power spectrum of a map with uncorrelated noise assuming the same white noise level as in the actual LiteBIRD map, and finding no such features. The wiggles do not seem to be present in polarization either. In addition, they are seen with different map-making methods in MAPPRASER (binned, maximum likelihood, and templates marginalization) as well as with the MADAM destriper. A remaining possibility that I was not able to verify, is that these could be caused by issues in the LiteBIRD simulation software itself. This issue is a subject of an ongoing investigation, however, considering that these are real features, they do not seem to be concerning given the LiteBIRD noise level in temperature.

8.1.3 Full-sky templates marginalization maps

In this subsection, I study the behavior of the templates marginalization approach using the specific case of polynomial templates as a way to handle $1/f$ noise. I perform tests using data of the LFT 40 GHz channel, but similar conclusions apply to the rest of the channels. I use the maximum likelihood maps as a reference to which I compare the noise structure in map domain obtained with other mapping methods, focusing on temperature maps, as the noise is white in polarization.

Figure 8.4 shows the difference between the binned noise map and the maximum likelihood noise map, as well as the differences between noise maps obtained with templates marginalization assuming baselines of 20 s or 120 s, and the noise maps obtained with the maximum likelihood estimator. Given that the binned mapper does not correct for the noise correlations, we see stripes in the difference maps. However, these features are not as strong as the usual stripe patterns seen in the binned Planck maps, for example Figure 6 of [277]. This is due to both a lack of sufficient noise power at low frequencies, given the smaller knee frequencies and milder slope of the detector noise expected for LiteBIRD and reflected in the simulation, and to the more optimal cross-linking of the LiteBIRD scanning strategy [240, 315]. Indeed we see in figure 8.3, that the TT noise power spectra of the binned case and the maximum likelihood case at the largest scales are roughly of the same order of magnitude, with the major difference being a highly anisotropic noise properties in the former case manifested by the stripes.

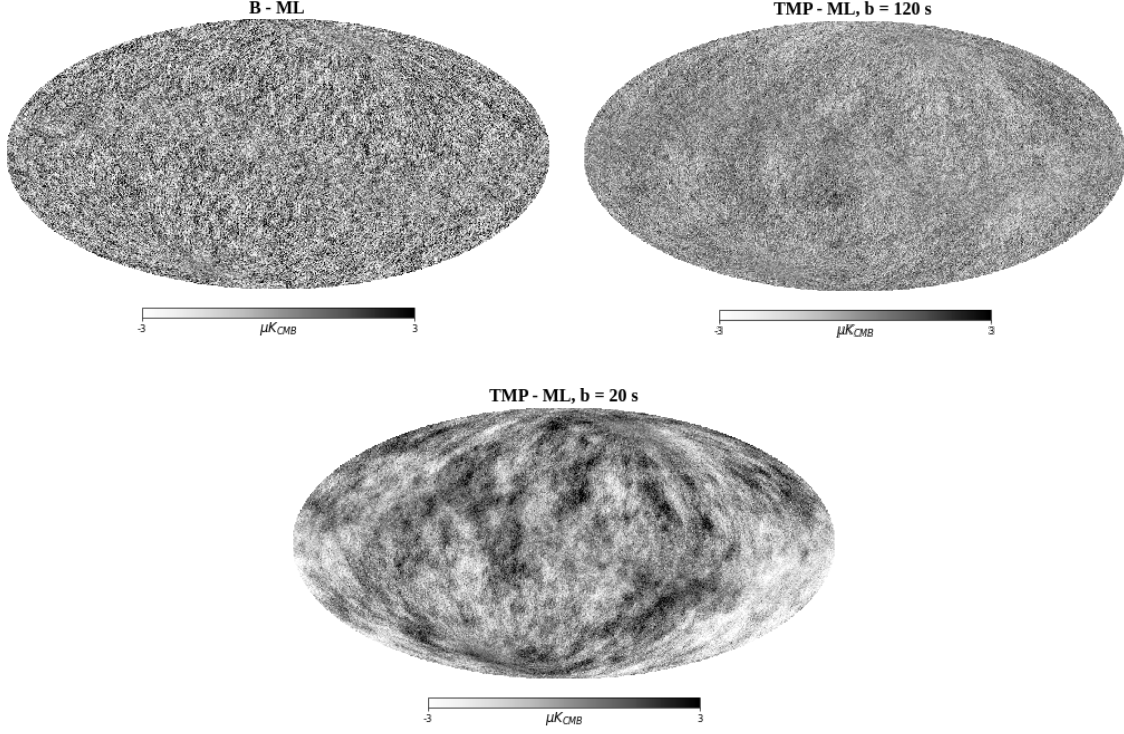


Figure 8.4 – The *top left* map shows the difference between the binned temperature noise map and the maximum likelihood one. The *top right* and *bottom* maps shows the differences between the templates marginalization temperature noise maps, respectively, for a baseline of 120 s and 20 s, and the maximum likelihood one. The maps are computed from the simulated data of the LFT 40 GHz channel.

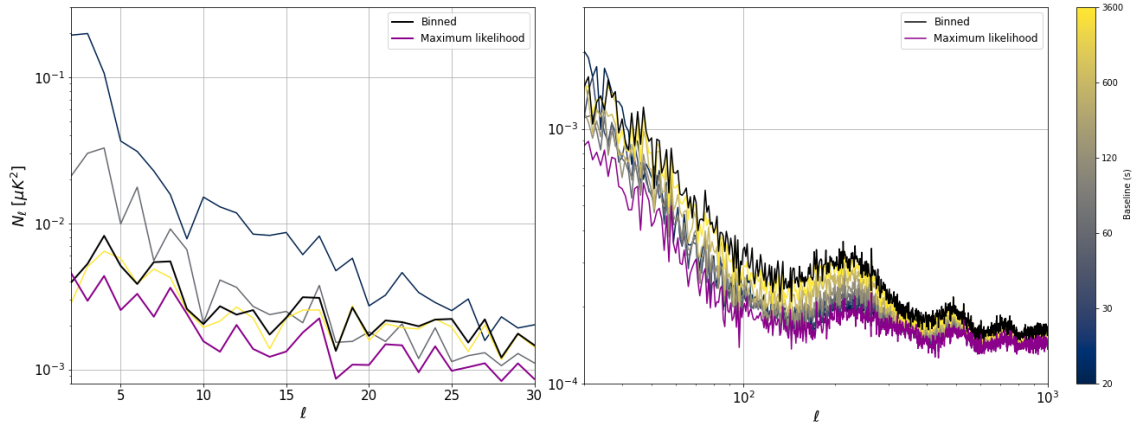


Figure 8.5 – The temperature noise power spectra obtained via the templates marginalization approach for different baselines. The *left* figure shows the $\ell \in [2, 30]$ range, and the *right* figure shows the $\ell \in [30, 10^3]$ range. I also include the noise power spectra from the binned and maximum likelihood maps for comparison. The maps are computed from the simulated data of the LFT 40 GHz channel.

These stripes are suppressed in the case of the maximum likelihood map-making

by down weighting the long temporal modes and in the case of the template marginalization by their explicit filtering. This comes at the price of adding more degrees of freedom, thereby increasing the overall level of the statistical noise. This effect should be particularly noticeable for short baselines, and is clearly visible by comparing the maps of figure 8.4. Processing the maps with short baselines for the polynomial filters boosts the large scale features in the noise map (bottom map), although the map obtained with an intermediate baseline length of 120 s (top right map) shows less stripe patterns than the binned map (at the top left).

These effects can be seen more clearly in the power spectra. Figure 8.5 shows the TT noise power spectra for multiple baselines compared with the binned and maximum likelihood map-makers. The right panel shows the TT noise power spectra at multipoles $\ell \in [30, 10^3]$. As the baseline shortens towards 20 s, the lowest value tested in these runs, the noise level gets closer to the maximum likelihood estimate, hence the increase of the statistical noise at these scales does not dominate over the correction of the systematic errors induced by the noise correlations. Conversely, as we increase the baseline lengths, we get closer to the binned map estimate, which is expected since the offsets span durations close to the length of the stationary intervals defined in the simulation, and the map-maker does the exact same processing as a binned map-maker to any modes shorter than the offsets baselines, thus not properly handling the $1/f$ noise.

The left panel shows the TT noise power spectra at lower multipoles, $\ell \in [2, 30]$. For clarity, only three baselines (20, 60 and 3600 s) are selected. In this regime, the statistical noise dominates as the number of sky modes becomes low and the maps processed with filter baselines shorter than typically $\sim 1/f_k$ become noisier than the binned maps at these scales. These findings are similar to observations made in [303, 277] in applications of destripers to Planck simulations as well as for ground-based observations [434]. In order to be able to process the maps with short baselines ($\lesssim 1/f_k$), one needs to include more information to bring these noise levels down. This could be done by introducing priors on the templates amplitudes, computed from the noise PSDs, as is done for instance in the MADAM or Descart destripers [277, 434], this also typically speeds up the convergence of the map-maker.

8.2 Ground-based experiments: Simons Observatory, Small Aperture Telescopes case

In this section, I study the performance of the map-makers in the context of ground-based observations. In particular, I focus on the case of the Small Aperture Telescopes of the Simons Observatory, studying different time-domain systematics: I first validate the filters of the scan synchronous and HWP synchronous signals on toy models, before focusing on more realistic simulations of atmosphere evaluating how different mapping methods are able to handle its impact on polarization. This work was performed as part of my involvement into the Simons Observatory Pipeline Working Group.

8.2.1 Simulation

All the studies conducted in this section are using the **S2** simulation described in the previous chapter. I summarize the main characteristics of the simulation in table 8.2. Given the large field of view of the SATs [209], the footprint of the sky survey is larger than the $10 \times 10 \text{ deg}^2$ targeted by the boresight pointing. Figure 8.6 shows the resulting hits map.

Parameters		Values
Hardware	Telescope	SAT
	Center frequency	150 GHz
	N_{det}	6020
	Sampling rate	37 Hz
	HWP rotation	2 Hz
Input sky	Nside	512
	Sky coordinates	RA 40..50, DEC -35..-45
	f_{sky}	$\sim 1\%$
Scans	N_{CES}	95
	Δt_{CES}	$\sim 15 \text{ min}$
	Scan speed	1.0 deg.s^{-1}

Table 8.2 – Summary of the **S2** simulation features.

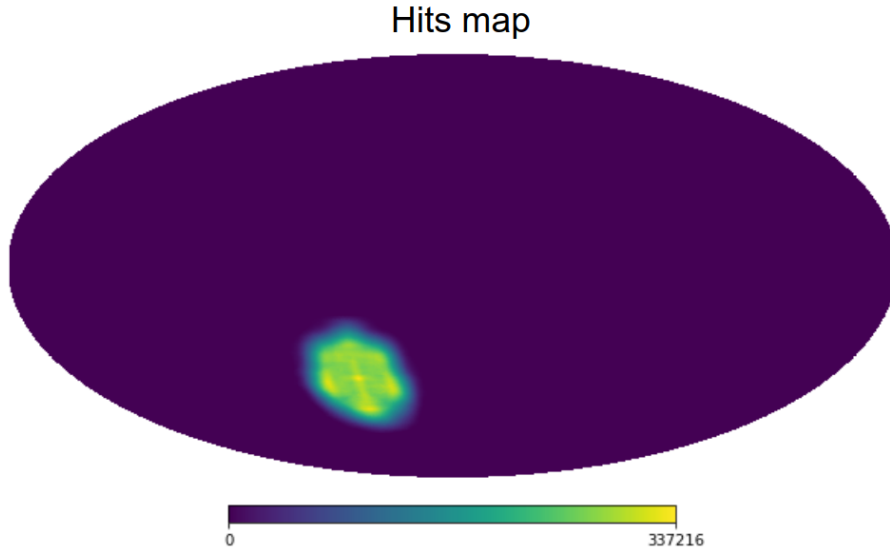


Figure 8.6 – Hits map of the **S2** simulation.

8.2.2 Validation: Scan & HWP synchronous signals mitigation

In this subsection, I validate the templates marginalization procedure at the map level. I select the input signals in the **S2** simulation such that I include only the cosmological CMB signal and, each time, one specific time-domain systematic

signal modeled by a given template for which we know the exact parameters. The sky maps are then reconstructed assuming 3 different noise models:

- A white noise model (Binned mapper).
- An incorrect template model: assumes the correct template class but with the wrong parameters.
- A correct template model: assumes both the correct template class and the correct parameters.

The binned mapper allows us to see the full impact of the systematic signal on the map if not taken into account in the reconstruction. Using an incorrect template model, allows us to qualitatively assess the impact of model errors, and the last configuration using the correct template model allows us to validate the procedure by verifying that we recover a correct estimate of the CMB map.

Scan synchronous signal subtraction

I inject a ground-pickup signal using the scan synchronous signal template, which takes the form of a step function in azimuth $g(\varphi)$; i.e. the ground map is binned in azimuth, with some fixed azimuth step $\Delta\varphi$, and in any given constant elevation scan we have $\forall\varphi \in [\varphi_{min} + i\Delta\varphi, \varphi_{min} + (i+1)\Delta\varphi[, g(\varphi) = g_i = \text{const.}$ Therefore at any given instant t , if the telescope's boresight is pointing towards the i -th bin in azimuth, we can simply write: $SSS(t) = g_i$. I use a linear ground map of the form $g_i = 1 + 0.01i$ (K), with an azimuth bin $\Delta\varphi = 0.33$ deg. For the map reconstruction, given that the azimuth range of the constant elevation scans is of 10 degrees, the correct template model accounts for 30 bins in azimuth for each CES, and only 10 bins in azimuth for each CES are accounted for in the incorrect model run. The reconstructed temperature maps are shown in figure 8.7.

The total offset of the map cannot be recovered, since it is fully degenerate with the ground template. Therefore, it is removed from the difference maps shown in the right panel of the figure. In the binned map, the coherent ground structure completely dominates the sky signal as one would expect, and we essentially see the projection of the offset of the ground map. The middle row maps show an improvement, as the ground signal is partially filtered, and the ground offset removed, however since we are assuming larger azimuth bins, we still get power from the sub-bin fluctuations leaking to our reconstructed sky signal, and given that the corresponding amplitudes fluctuations is of order 10 mK, the ground leakage still dominates over the CMB signal. In the bottom row, we recover the correct CMB map effectively filtering the ground signal. The difference map in the bottom right panel, shows a low amplitude (below $\sim 0.02 \mu\text{K}$) smooth residual structure, which essentially comes from degeneracies between the sky and the ground template. These degenerate modes correspond to the projections of the azimuth bins of the ground template on the sky, which, as the sky rotates, leave stripes which are constant in right ascension (such as spherical harmonics with $m = 0$), with unconstrained relative offsets. Some of these degeneracies can be broken when data from multiple detectors, with shifted azimuth bins, is combined, or with the sky pixels shared between adjacent bins. Nonetheless,

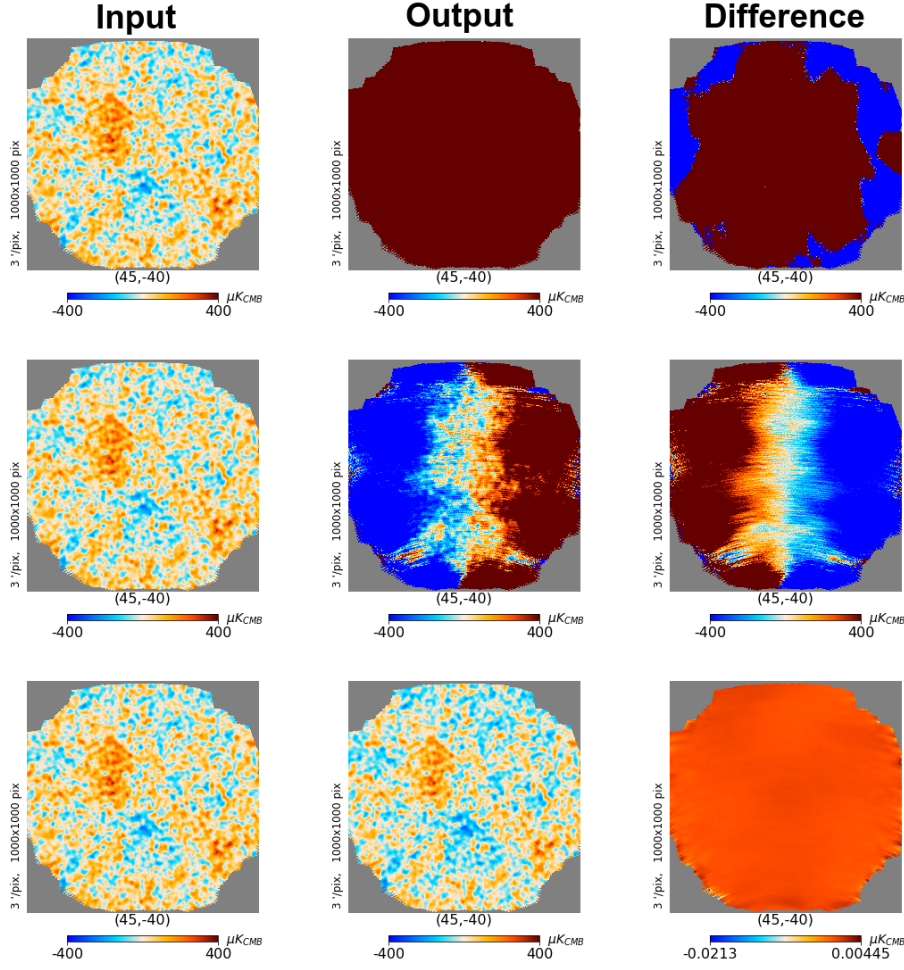


Figure 8.7 – Temperature maps of CMB+SSS mapped with respectively, a binned map-maker, a 10 azimuth bins, and 30 azimuth bins SSS templates map-makers (top to bottom).

these constraints are typically weak, leaving a number of ill-constrained modes in declination which characteristics depend on the pixelization, the size of the azimuthal bins and the scanning strategy [378, 345].

With these results we can validate the SSS filtering operators and the corresponding templates marginalization procedure.

HWP synchronous signal subtraction

I inject a HWP synchronous signal modelled by a set of harmonics of the HWP rotation frequency, coupled with a linear drift term [266, 304, 440]. For each CES and each detector, it takes the following form:

$$\mathcal{H}(\phi_t) = \sum_n A_n(1 + \varepsilon_n t) \cos(n\phi_t) + B_n(1 + \varepsilon_n t) \sin(n\phi_t), \quad (8.1)$$

where ϕ_t is the HWP angle at time t . The harmonics amplitudes A_n and B_n are drawn from a normal distribution $\mathcal{N}(\mu, \sigma)$, with $\mu = 100$ mK, and a dispersion σ/μ of 1%. The linear drift coefficient ε_n is set such that we get a 1% drift every 1000 s (~ 16 min). We inject a total of six harmonics. For the reconstruction,

we consider a HWPSS template model for which the baseline length (i.e. the time interval over which the harmonics amplitudes are considered to be constant) is 10 s. The incorrect model run takes into account only four harmonics of the HWP rotation frequency, f_{HWP} , while the correct model run handles all of the six harmonics. The reconstructed temperature and Q polarization maps are shown, respectively, in figures 8.8 and 8.9.

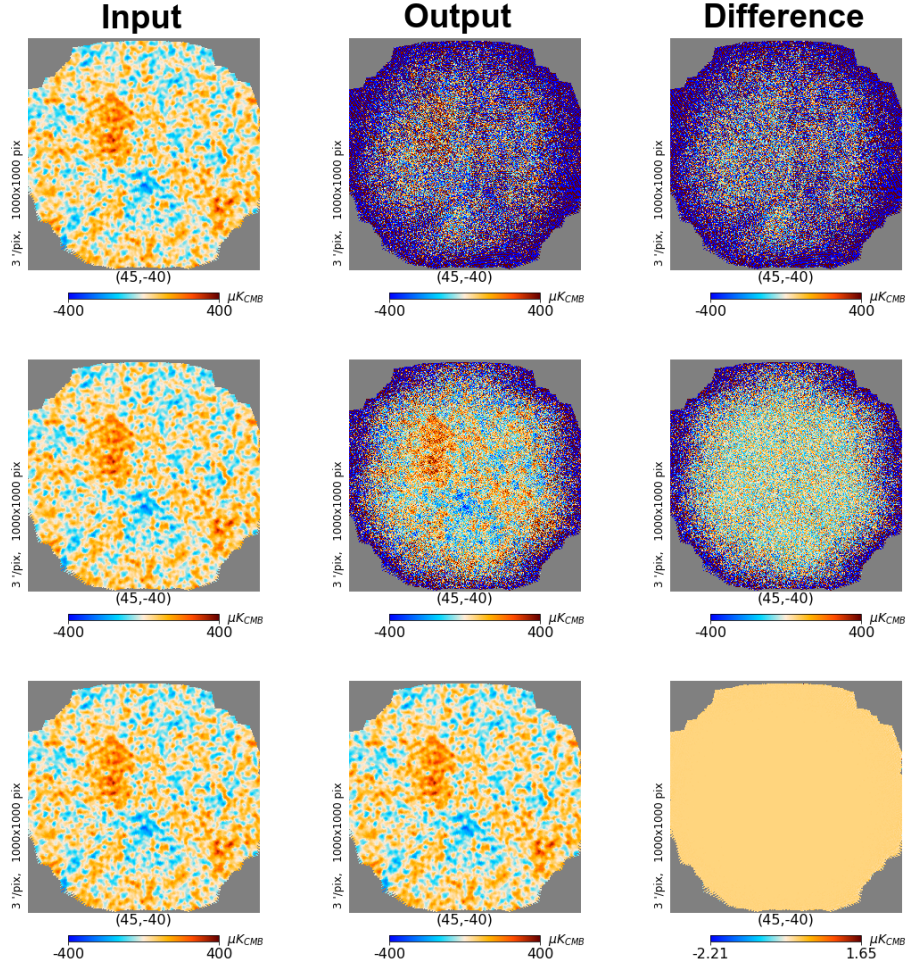


Figure 8.8 – Temperature maps of CMB+HWPSS mapped with respectively, a binned map-maker, a fourth, and sixth order HWPSS templates map-makers (top to bottom).

The binned maps, in the top rows, show the direct projection of the contaminated TOD into the map, and given the amplitude of the HWPSS, the CMB signal is completely dominated by the spurious signal. In the middle row of figure 8.8, the filtering of the first four harmonics reduces the noise level of the temperature map at large scales, as we start to see the sky signal emerging, however the remaining $5f_{\text{HWP}}$ and $6f_{\text{HWP}}$ harmonics still leak as small scale noise given that these components reside in the high frequency tail of the signal. For polarization, the signal is modulated at $4f_{\text{HWP}}$ and the filtering of the first four harmonics removes the large scale stripy structure we see in the top row of figure 8.9, caused by spurious signals varying more slowly than the polarization modulator, but the leakage from the $5f_{\text{HWP}}$ and $6f_{\text{HWP}}$ harmonics is still sufficiently strong for the CMB signal to remain subdominant. This is due to two aspects: (1) the $5f_{\text{HWP}}$ and $6f_{\text{HWP}}$ har-

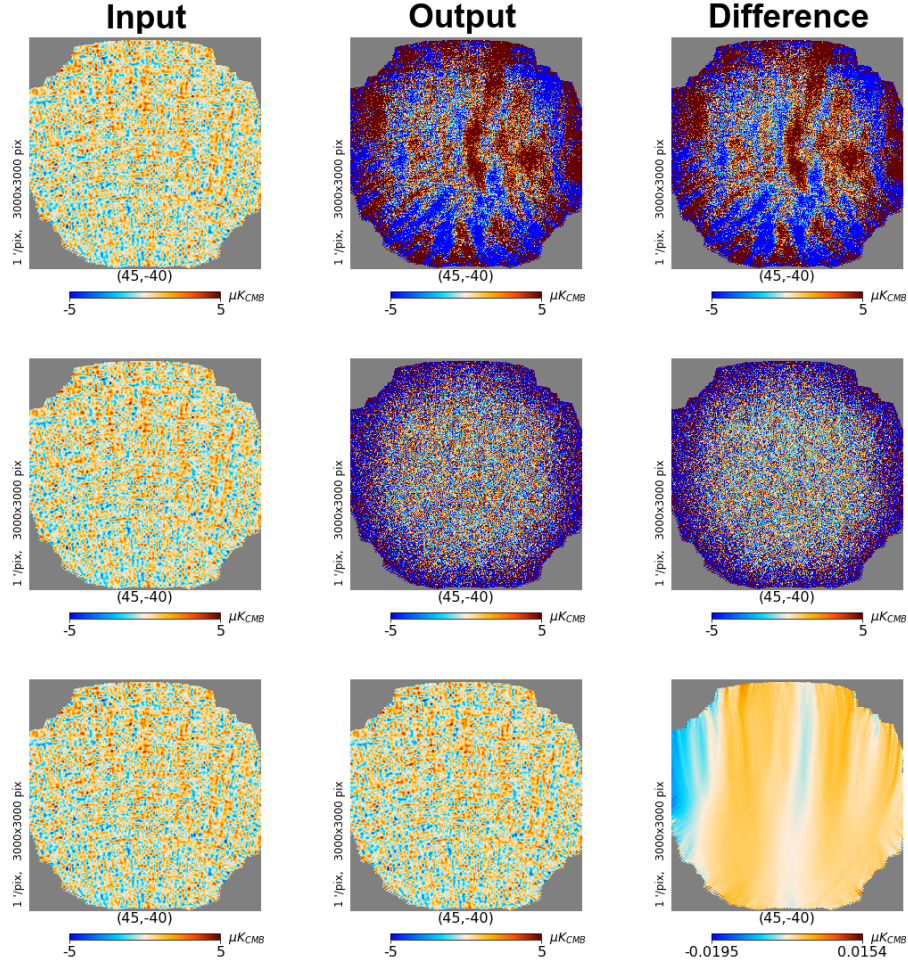


Figure 8.9 – Q polarization maps of CMB+HWPSS mapped with respectively, a binned map-maker, a fourth, and sixth order HWPSS templates map-makers (top to bottom).

monics are closer to the center of the science band in the case of polarization than in temperature because of the HWP modulation, and (2) the polarization signal is fainter in amplitude than temperature. In the bottom rows, we recover the correct input CMB signal, and the HWPSS is effectively filtered thereby validating the HWPSS filtering operators and the full templates marginalization procedure. The residual structure seen in the difference map is below $\sim 0.02 \mu\text{K}$ in amplitude, and is due to the linear drift term that is not perfectly modelled in the map reconstruction, since the harmonics amplitudes are considered to be constant over each 10 s interval.

8.2.3 Atmospheric effects

In this subsection I explore the effects of atmosphere on the polarization maps of the SATs. The simulations are again based on **S2**, featuring a HWP rotating at 2 Hz and using the same TOAST based atmospheric simulation tools. However, to simulate a longer integration time, I scale down the noise by a factor $1/\sqrt{365}$. Given that **S2** represents a total of 24 h of observations, the results will thus be representative of a full effective year of observations. Three approaches will be

considered in this study. The first one will be based on the banded block-Toeplitz noise model, where I assess the impact of the noise bandwidth used in the map reconstruction on the convergence of the solver as well as the quality of the obtained maps. I use template marginalization with polynomial filters as a second approach, and similarly evaluate the impact of different orders on the convergence and quality of the maps. Note that since the actual simulation represents only 24 h of observations, the number of degrees of freedom of the polynomial templates is lower than what is expected in an actual full year of observations, and this test case may be underestimating the statistical noise. The obtained results will need to be confirmed in a full scale simulation. In the last approach, I run the map-maker with Toeplitz-structured noise model on pair differenced data, and investigate the impact of intensity-to-polarization leakage for different levels of gain mismatch between orthogonal pairs. Appendix C presents a detailed analytical calculation to compare the latter approach to the standard case with the full IQU pointing model.

Toeplitz noise covariance: impact of noise bandwidths

As explained in the previous chapter, the atmosphere which is the dominant source of noise in the detector timestreams, follows to a good accuracy a power spectral density of the form

$$P(f) = \sigma_0^2 \left(1 + \left(\frac{f + f_0}{f_k} \right)^\alpha \right), \quad (8.2)$$

behaving as white noise with power σ_0^2 at frequencies above the knee frequency, f_k . Below this latter frequency, it scales as a power law $\propto f^\alpha$. At the lowest frequencies, below f_0 , this scaling breaks again to a white noise behavior, also called *white noise ceiling*. The range of frequencies where the noise behaves as a power law is responsible for the noise correlations. The right panel of figure 8.10 shows the inverse power spectral density of the noise timestream of a single detector for one CES of the S2 simulation, featuring all the above characteristics. The left panel shows the distribution of the knee frequencies for all the detectors in the focal plane and for all the CESs in the simulation.

In the exact maximum likelihood map-making solution, the $N^{-1} \times$ operation is multiplying the timestreams by the correct inverse noise power, thus down weighting the noise in an optimal way and correcting for the effects of correlations. In practice, as explained in Chapter 6, the problem can only be solved if we assume some simple noise structure as in the assumption of a banded block-Toeplitz noise covariance. The size of the bandwidth, characterized by the half bandwidth parameter λ , has effects on both the convergence and the quality of the recovered maps. A larger bandwidth will typically induce a less well conditioned system matrix, hence slowing down the convergence of the solver. However, it will also be better at down weighting the noisiest modes. This latter property can be understood from the fact that the apodization of the time-domain covariance with a Gaussian window of width $\sim \lambda$ for example, corresponds to a convolution operation in frequency-domain with a Gaussian kernel of width $\sim 1/\lambda$, which is effectively a smoothing operation in frequency. This smoothing will result in the flattening of the resulting power spectral density at low frequencies below $\sim 1/\lambda$ (see also Appendix D of [438]). If the bandwidth is not large

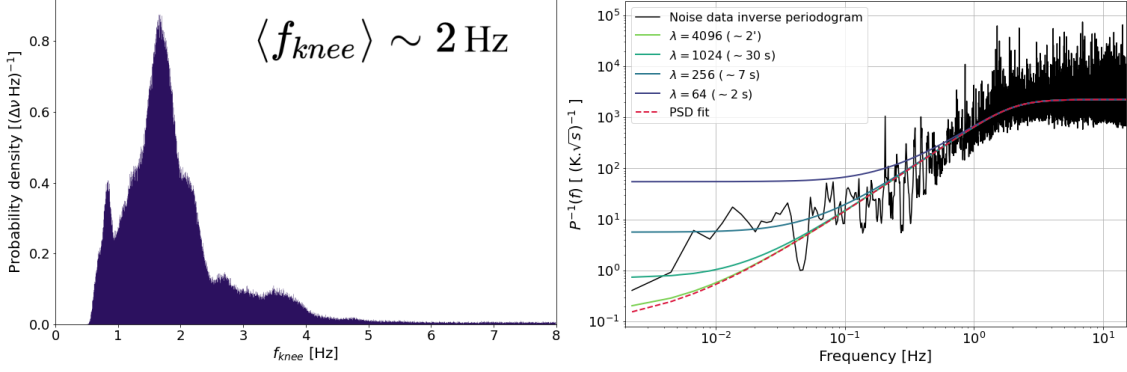


Figure 8.10 – *Left*: Distribution of the knee frequencies induced by atmospheric noise for all detectors in the focal plane and for all the scans of the **S2** simulation. The binning in frequency is $\Delta\nu = 5 \cdot 10^{-3}$ Hz. *Right*: Inverse periodogram computed from the noise timestream of a single detector during one CES, overlaid with the best-fit to our PSD model and the resulting PSDs after apodization with Gaussian windows corresponding to different bandwidths of the Toeplitz blocks.

enough ($< 1/f_0$), the low frequencies will not be sufficiently down weighted leaving some residual correlations in the data. In practice, we do not have to down weight all the modes down to f_0 , but only down to frequencies where the sky signal resides. This flattening effect is shown in the right panel of figure 8.10 for different bandwidths. We should also note that because we ignore the effects of correlations between detectors in the noise model, the resulting maps even for a sufficiently large bandwidth, will be sub-optimal. In the following, I will refer to this mapping method as the down weighting approach.

To demonstrate these effects, I run the **S2** simulation with the scaled down noise, considering different values of the half bandwidth in the noise model used for the map reconstruction, spanning from 1 time sample, corresponding to the limiting case of a binned map-maker with uncorrelated white noise, to 4096 time samples corresponding to a correlation length of roughly 2 minutes given our sampling rate. In this case, as one can see in the right panel of figure 8.10, the bandwidth is sufficiently down weighting the low frequencies, and hence we do not expect a gain in performance beyond this limit. The convergence curves are shown in figure 8.11, with the binned case excluded since it always converges in a single step. Figure 8.12 shows the EE and BB noise power spectra of the different cases compared with the auto-spectra of the input CMB map.

As expected the number of iterations increases monotonically with the bandwidth. We also notice that even though each case represents a factor 4 increase in the bandwidth over the previous one, the resulting number of iterations increases very quickly in the first cases, but quickly seems to become less sensitive to bandwidth changes as we approach the limiting case $\lambda = 4096$. This property could be understood from figure 8.10: as we increase λ , the shape of the PSD is substantially changed in the low frequency part, however as we approach to the limiting case, the solution tends to converge towards the unconvolved PSD and the shape of the resulting PSD becomes stable to any further increase of λ in the frequency range where the sky signal resides. If the spectral properties are less

SAT simulation: Scaling of convergence curves for different bandwidths

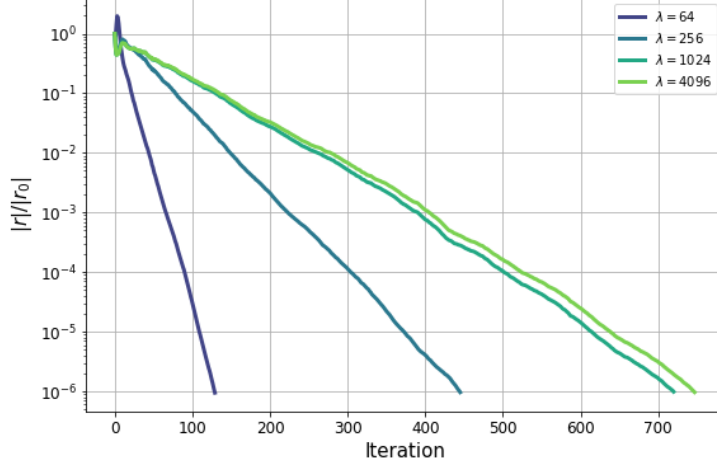


Figure 8.11 – Convergence curves for different values of the half bandwidth in the SAT simulation.

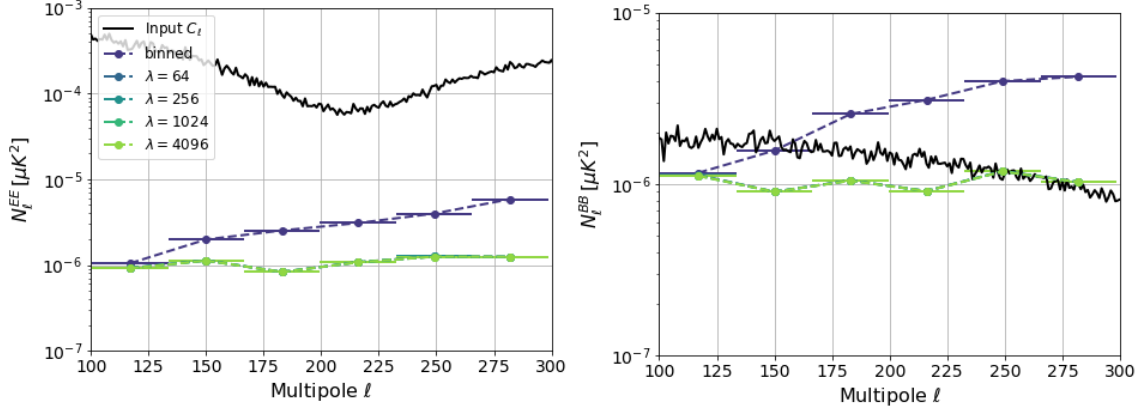


Figure 8.12 – EE and BB noise power spectra for different values of the half-bandwidth, compared with the C_ℓ of the input sky signal (solid black curve), in the SAT simulation.

sensitive to λ , then one should expect similar behavior of the convergence which is fully determined by these spectral properties.

At the power spectrum level, we see a net gain over the binned map in all the other cases. The noise is roughly white in the multipole range of interest, i.e. $\ell \in [100, 300]$, with an amplitude of $10^{-6} \mu\text{K}^2$, largely below the EE signal, and around a factor 2 lower than the BB signal around $\ell \sim 100$ but as the BB signal drops the gap between the two is eventually closed. However, the noise level displays low sensitivity to the bandwidths considered, as all cases seem to perform the same. To quantify the slight changes in the quality of the output between the different cases, I compute the relative differences between the noise power spectra at each case and that of the limiting case $\lambda = 4096$

$$\xi_\ell^{XX}(\lambda) = \frac{N_\ell^{XX}(\lambda) - N_\ell^{XX}(\lambda = 4096)}{N_\ell^{XX}(\lambda = 4096)}, \quad (8.3)$$

where X can be either E or B . Figure 8.13 shows the obtained results. We see a monotonic improvement on the quality of the reconstructed maps as we increase

the bandwidth. The overall improvements are very small however (mostly below 1%), and as expected there is less and less sensitivity as we get closer to the limiting case $\lambda = 4096$.

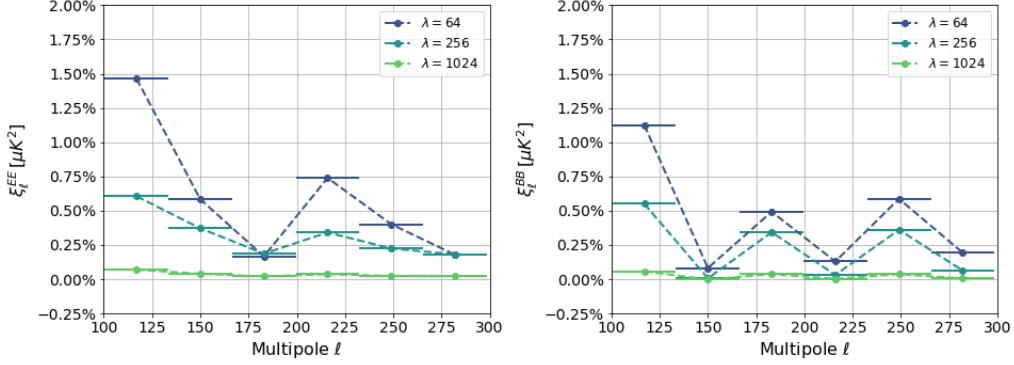


Figure 8.13 – Relative difference of the EE and BB noise power spectra for different values of the half-bandwidth, with the case $\lambda = 4096$, in the SAT simulation.

Templates marginalization: polynomial filtering

Let us now consider the second approach using templates marginalization to reconstruct the sky maps, where we define the templates as a set of Legendre polynomials. The baselines of the polynomial filters are set to 10 s to correspond to each of the left or right sweeps of the telescope. We consider 3 cases: one including only offsets, a second case including 3 orders of polynomials, and a third case including 6 orders of polynomials. Figure 8.14 shows the convergence curves obtained for these three cases overlaid with the previous cases from the down weighting approach.

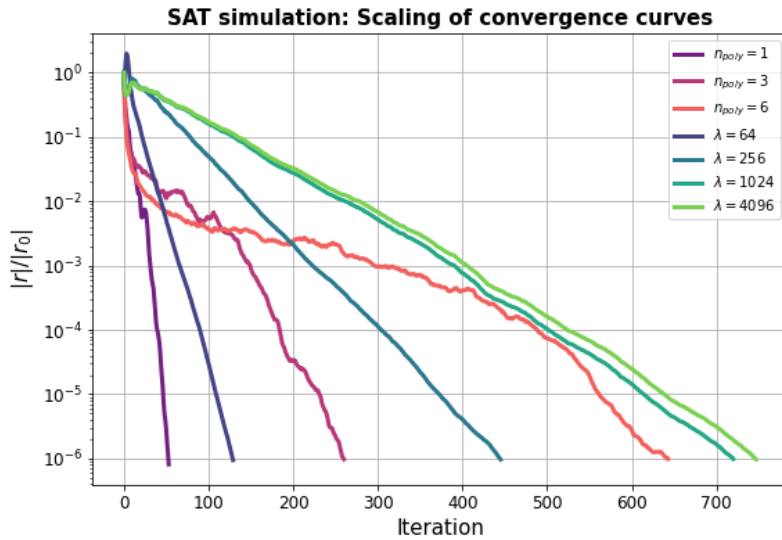


Figure 8.14 – Convergence curves for different orders of the polynomial filters and different values of the half-bandwidth in the SAT simulation.

The map-making solver tends to get slower as we increase the order of polynomials. This behavior is expected because of the degeneracies between polynomial

templates and the sky signal. As we introduce more polynomials, we tend to get more singular and nearly singular modes in the system which then slow down the convergence. For the cases considered however, the number of iterations is comparable to the down weighting case.

Figure 8.15 shows the EE and BB noise power spectra for the different polynomial orders along with the binned and the banded block-Toeplitz case corresponding to $\lambda = 4096$. The noise from polynomial filtering is also flat in the considered multipole range, but has around factor 2 lower amplitude ($\sim 5 \cdot 10^{-7} \mu\text{K}^2$) than the down weighting approach considered in the previous subsection. It is thus orders of magnitude below the EE auto-spectrum, and between a factor 4 and factor 2 lower than the BB auto-spectrum. The better quality of the maps obtained with polynomial filtering can be explained by the fact that polynomial filters remove some of the modes which happen to be correlated across detectors, while these are left uncorrected for in the down weighting approach—where no detector correlations are considered in the noise model—and therefore result in degraded noise properties. In principle, the filtering can remove part of the sky signal together with the noise whenever there is not enough redundancy in the data to recover the former. However, these modes are typically very noisy regardless and do not contribute much to the constraints, hence the corresponding information loss is acceptable.

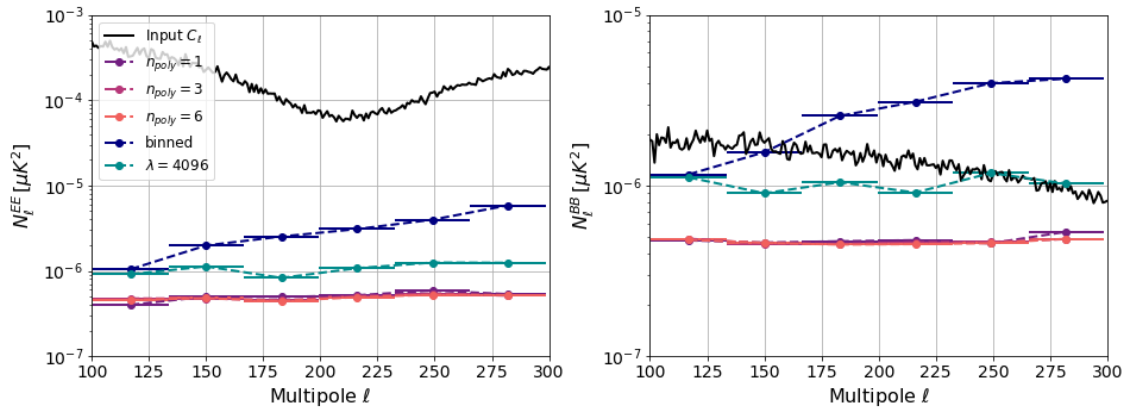


Figure 8.15 – EE and BB noise power spectra for different orders of the polynomial filters, compared with the C_ℓ of the input sky signal (solid black curve), in the SAT simulation.

In order to compare how well the polynomial filtering performs with respect to nearly optimal maps, I run a down weighting map-maker fed with pair-differenced time-ordered data such that the atmospheric contribution to the signal is completely removed, and the only remaining noise contribution is instrumental, which is assumed to be uncorrelated between detectors. In this case, we should obtain a solution that is a reasonably close approximation of the true maximum likelihood map. In Appendix C, I show that the maximum likelihood map-making is in general different than the pair-differencing approach applied here. This is because in this latter case the timestreams of a pair of the detectors are differenced without any weights. Instead, the maximum likelihood approach could be thought of as a pair-differencing scheme where the timestreams are weighted with detector specific weights. This weighting ensures that the maximum likeli-

hood estimates indeed attain the smallest statistical uncertainties. However, if the detector-specific weights in each pair are different, and the atmospheric signal is present then the weighting does not allow for the latter to be perfectly cancelled out in each detector pair. This in turn will result in a residual atmospheric signal in the recovered polarization maps. Both approaches therefore trade between the statistical and systematic errors. In our case however, the noise is dominated by the atmosphere itself and the subdominant detector noise varies on the level of 1%. The weights in each pair are therefore essentially identical and, in the absence of systematic effects, the pair-differencing scheme should be residual free and with nearly the best noise level.

We find the noise power spectra of the maps obtained via pair-differencing to be very close to the maps obtained with polynomial filtering. Figure 8.16 shows the difference of the pair-differencing case relative to the case with $n_{poly} = 6$. It also features the relative differences (also to $n_{poly} = 6$) of the two other polynomial filtering cases, $n_{poly} = 1$ and $n_{poly} = 3$. The near maximum likelihood case (pair-differencing) is overall performing at most 10% better than the $n_{poly} = 6$ case. We also see an overall improvement of a few percents on the maps quality as we increase polynomial orders.

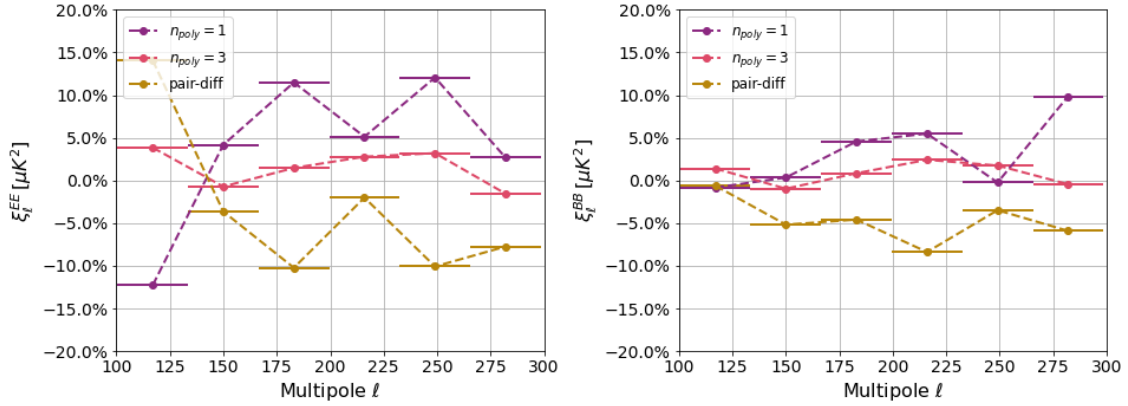


Figure 8.16 – Relative differences of the EE and BB noise power spectra for the pair differencing case along with the different orders of the polynomial filters, relative to $n_{poly} = 6$, in the SAT simulation.

These latter results suggest that a better strategy to handle the effects of atmosphere on the polarization maps may be to run the map-maker on pair differenced data. As discussed above, in the ideal case of no systematics, the noise properties are nearly optimal. In addition, it also provides the advantage of faster convergence since the system matrix is well conditioned in general when intensity is excluded, given that polarization benefits from greater diversity of attack angles in general, due to intrinsic orientations of detectors, sky modulation, in addition to polarization modulators such as the HWP. The case I show in figure 8.16 for example converges in 3 steps. However, in the presence of calibration errors between pairs of detectors, the intensity can leak to the polarization signal and degrade the quality of the maps (as well as convergence). I investigate this in the next subsection for the specific case of gain mismatch, however one should expect similar effects from beam or bandpass mismatch as well [24, 34].

Intensity-to-polarization leakage effects

I consider three levels of gain mismatch in the following: a case with 0.5% gain mismatch in line with the anticipated level of calibration in the Simons Observatory experiment, a more conservative figure of 1%, and a pessimistic case of 5%. The instrumental noise has a very low characteristic knee frequency of 60 mHz. However, even though a large part of the atmosphere is suppressed with the pair differencing, we expect the residual signal to have a sufficiently strong $1/f$ to give rise to an effective knee frequency that is way higher than the instrumental one. In fact, the knee frequency should increase monotonically with the level of gain mismatch, as bigger mismatch leaves more of the atmospheric contribution. This is indeed what is shown in figure 8.17 where we visualize the distribution of the knee frequencies of orthogonal pairs of detectors for different levels of gain mismatch. The increasing knee frequencies should have an effect on both convergence and the noise levels on the resulting maps. However, in the presence of a HWP the polarization modulation should be able to efficiently mitigate these effects. I aim to demonstrate this by running the map-maker, using the down weighting approach, with $\lambda = 4096$, on a case without HWP and another one with the HWP.

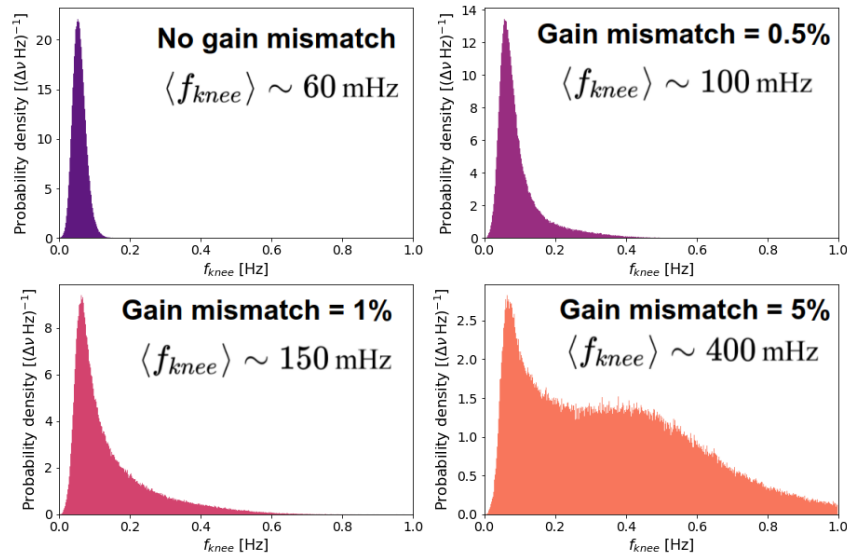


Figure 8.17 – Distribution of the knee frequencies of orthogonal pairs in the **S2** simulation for different levels of gain mismatch. The binning in frequency is $\Delta\nu = 2 \cdot 10^{-3}$ Hz. The residual atmosphere signal causes the right tail of the distribution to get stronger, increasing the effective knee frequency.

Without HWP In the absence of a HWP, the large scale fluctuations of the atmosphere leak directly to the science band of the polarization signal. Even with a percent level suppression, the brightness of the atmosphere is such that it can significantly degrade the quality of the reconstructed maps. Figure 8.18 shows the obtained maps for different levels of gain mismatch. The EE and BB noise power spectra are shown in figure 8.19. The 0.5% and 1% cases boost the noise by up to two orders of magnitude, overwhelming the BB signal. While the pessimistic case of 5% dominates over both EE and BB.

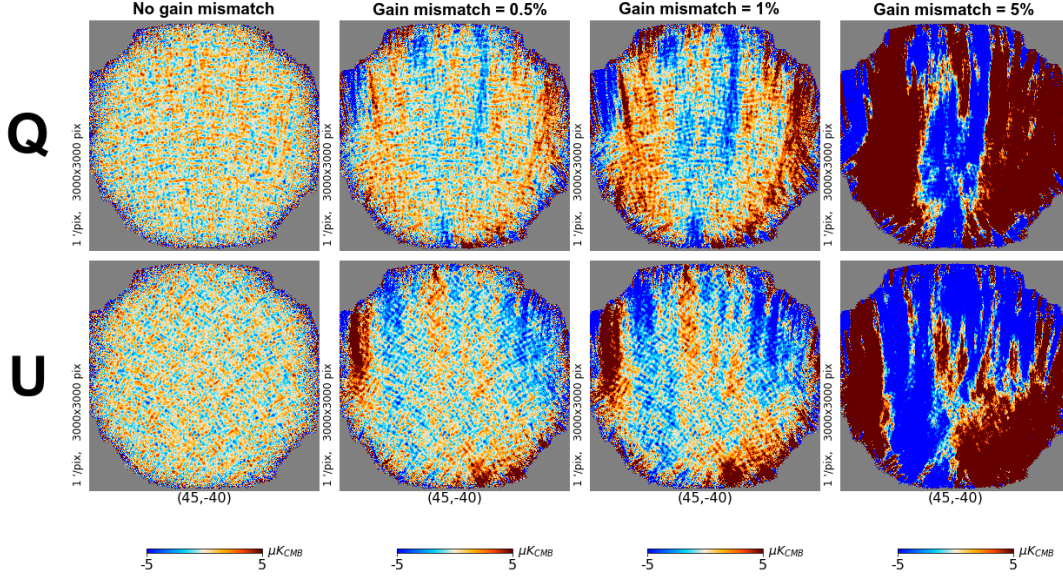


Figure 8.18 – Q and U polarization maps obtained with different levels of gain mismatch and no HWP.

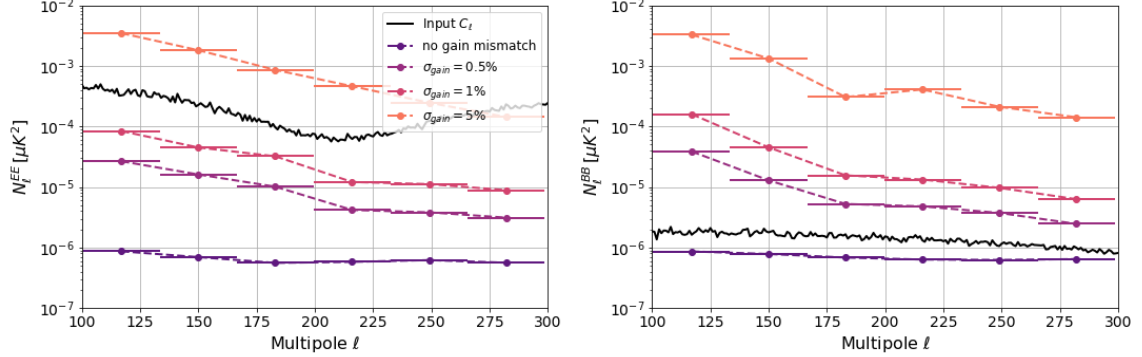


Figure 8.19 – EE and BB noise power spectra for different levels of gain mismatch between orthogonal pairs, in the SAT simulation with no HWP.

With HWP The modulation of the HWP displaces the science band to higher frequencies thereby avoiding the low frequencies where the atmospheric noise is particularly strong. As a result the quality of the reconstructed map is significantly improved, and the mismatch effect is effectively mitigated. Figure 8.20 shows the recovered maps in this case, and figure 8.21 shows the obtained EE and BB noise power spectra where we can see that the noise levels are below the targeted EE and BB signals, and comparable to levels obtained with the optimal case (no mismatch) even for the pessimistic case. The relative differences with respect to the optimal case are found to be respectively below 1%, 2% and 20% for the three cases (from lower to higher gain mismatch). Pair-differencing is thus found to be better than the polynomial filtering approach on a standard IQU model for the more realistic levels of calibration errors (0.5% and 1%), and is in fact significantly better in terms of convergence as it only requires a few steps.

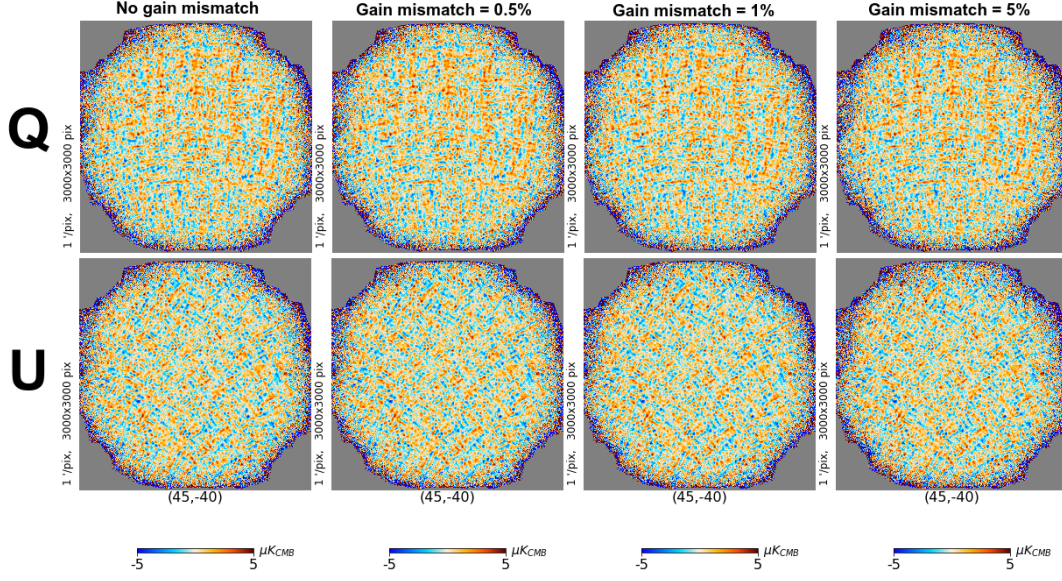


Figure 8.20 – Q and U polarization maps obtained with different levels of gain mismatch in the presence of a HWP rotating at 2 Hz.

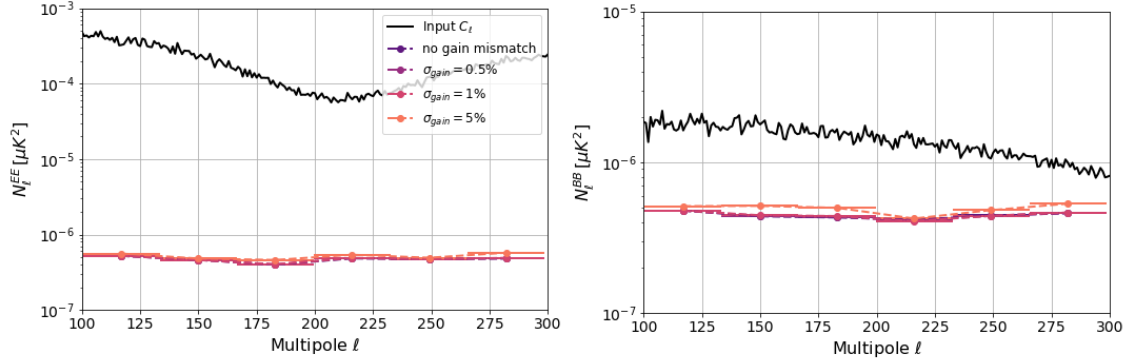


Figure 8.21 – EE and BB noise power spectra for different levels of gain mismatch between orthogonal pairs, in the SAT simulation with HWP.

These results demonstrate the robustness of the HWP mitigation of intensity-to-polarization effects resulting from miscalibration effects. However as I have noted previously, the HWP introduces its own systematics [403, 292, 459, 170]. In particular, one should ensure that the HWP synchronous signals are effectively mitigated. I provided a demonstration for such a procedure in Section 8.2.2 on a toy model. More realistic models have been developed in the context of the Simons Observatory systematics studies [403]. These have only recently been incorporated in the TOAST simulation framework. Therefore one of the natural evolutions of the work conducted in this thesis, is to study the templates marginalization approach on these realistic models.

As we have seen, in the absence of a HWP the gain mismatch effects are quite strong. In particular, this means that the gain mismatch systematics, as well as the bandpass and beam mismatch, could be a more concerning issue for the Large Aperture Telescope which does not feature a polarization modulator. Future work could be oriented towards investigating these effects along with their mitigation strategies for LAT.

Conclusion

With the advent of the era of high precision cosmology, we are able to pursue puzzling questions not only on the origin and structure of the Universe, but also on the subatomic world of particle physics. These questions include the search for primordial B-modes aiming at unveiling the secrets behind the high energy physics governing the very first instants of our Universe's history, the search for light particles, the measurement of neutrino mass and determining their mass hierarchy, as well as investigating the nature of dark matter and dark energy, and many more. Different types of observations bringing complementary insights on these questions are being combined to reach unprecedented sensitivities in our measurements.

In the context of CMB observations in particular, the scientific exploitation of the data collected by the experiments faces a number of challenges brought about by the increasingly more complex design of the instruments and the high levels of sensitivity required to reach the science targets, which in turn calls for more efficient and robust methods and algorithms to process the data, as well as for their cutting-edge implementations capable of fully capitalizing on the power of modern day supercomputers.

I have reviewed all these aspects in the first five chapters of this thesis. I started with a review of the fundamentals of the standard model of cosmology in Chapter 1, before moving in Chapter 2 to an introduction to one of the major scenarios of the physics governing the very early Universe, namely inflation, which search for its B-mode signature is motivating extensive experimental efforts currently being deployed, and more prepared for the near and long term future. In Chapter 3, I presented a review of CMB science, and later discussed the broader picture with multiple probes that are of interest for modern cosmology in Chapter 4. Chapter 5 presents an overview of some of the major CMB experiments currently operating or in preparation. In particular, I presented the science results we have reported in the POLARBEAR collaboration during my PhD. I finish the chapter by a discussion of the data analysis challenges in CMB context, along with an overview of the methods employed in the different stages of the data analysis pipeline.

My involvement in the POLARBEAR-2/Simons Array and Simons Observatory experiments, has been primarily in the context of the preparation of the data analysis pipelines. In particular, my focus was on the development of a massively parallel map-making framework suitable for multi-kilo pixel arrays. This framework implements multiple map-making approaches for reconstructing unbiased sky estimates in the presence of time-domain systematics. The precise mathemat-

ical formalism as well as the technicalities of the different approaches are covered in Chapter 6. Specific examples of systematics are discussed and readily available in the code. The framework also allows for any systematic, which effect is linear in the data model, to be added and accounted for.

In Chapter 7, I give a detailed description of the implementation, evaluate the scaling performance of the code and demonstrate its ability to perform optimally while using a large amount of resources. I also present two novel implementations of linear solvers techniques and study their properties on realistic simulations. In particular, I demonstrate that the two-level preconditioning technique can yield significant advantage in multiple solves, as in the context of null tests or Monte Carlo simulations.

Several applications are subsequently presented in Chapter 8. In particular, I study the different map-making approaches in simple settings for a satellite experiment like LiteBIRD. I have also demonstrated the mitigation of scan and HWP synchronous signals on toy models with the help of simulations of the Simons Observatory Small Aperture Telescope. These simulations are also used to study the map-making methods in the presence of atmosphere. I find the effects of cross-correlations between detectors to be important, degrading the polarization noise level of the down weighting approach by about a factor two. These effects are however mitigated by the filtering approach which yields a noise level close to the optimal case by about 10% or less. Finally, the down weighting approach with pair differenced timestreams produces the best results in the presence of a HWP, which can robustly mitigate gain mismatch induced intensity-to-polarization effects. In this latter case, the noise level gets close by 1-2% to the ideal case with no calibration errors, and the solver converges in only a few steps.

For future work, I plan on continuing the development of the framework, incorporating more sophisticated noise models, e.g. handling detector correlations, which are going to be crucial for high fidelity temperature measurements, but which also have an effect on polarization as we have seen. In addition, I would like to extend the capabilities of the framework to handle non-linear models and perform a time-domain component separation concurrently reducing multi-frequency data and cleaning foregrounds. As a first application, this could be applied to mitigate atmosphere, capitalizing on its frequency scaling. In addition, it would also allow to handle a broader class of frequency dependent systematics [459]. These extensions will require heavy computational optimization of the software. I intend on improving the framework to implement a full hybrid OpenMP/MPI model, as well as implement routines which can capitalize on GPU-based architectures. In parallel, efforts need to be deployed on improving the linear solvers techniques currently implemented. All these developments will be informed by realistic simulations, which I will extend to high resolution observations as well, such as from the Large Aperture Telescope, but crucially I intend to apply all these methods in the actual analysis once the data from the experiments I'm involved with is available.

Block-wise derivation of the unbiased map estimator

In this appendix, we derive the general expressions of unbiased sky estimators given in Chapter 6, Eq. (6.4) and Eq. (6.9). The starting point is the generic data model introduced in Eq. (6.2) which we rewrite here for convenience

$$\mathbf{d} = \mathcal{P} \mathbf{y} + \mathbf{n},$$

where the generalized pointing matrix, \mathcal{P} , and the generalized map, \mathbf{y} , can be written in terms of blocks as (using the same notation as in Chapter 6)

$$\mathcal{P} \equiv \begin{bmatrix} \mathbf{P} & \mathbf{T} \end{bmatrix}, \quad \mathbf{y} \equiv \begin{bmatrix} \mathbf{m} \\ \mathbf{x} \end{bmatrix}. \quad (\text{A.1})$$

We rewrite here, in a block form, the expression of the generalized least square estimator of the generalized map, with a weight matrix \mathbf{W} , given in Eq. (6.3)

$$\begin{bmatrix} \hat{\mathbf{m}} \\ \hat{\mathbf{x}} \end{bmatrix} = \begin{bmatrix} \mathbf{P}^T \mathbf{W}^{-1} \mathbf{P} & \mathbf{P}^T \mathbf{W}^{-1} \mathbf{T} \\ \mathbf{T}^T \mathbf{W}^{-1} \mathbf{P} & \mathbf{T}^T \mathbf{W}^{-1} \mathbf{T} \end{bmatrix}^{-1} \begin{bmatrix} \mathbf{P}^T \\ \mathbf{T}^T \end{bmatrix} \mathbf{W}^{-1} \mathbf{d}. \quad (\text{A.2})$$

We recall that the above expression is only strictly valid if \mathcal{P} is full column rank, which we assume in this derivation (If not, one needs to regularize the inversion). Let us denote the four blocks of the system matrix in the above equation by

$$\mathbf{A} \equiv \mathbf{P}^T \mathbf{W}^{-1} \mathbf{P}, \quad \mathbf{B} \equiv \mathbf{P}^T \mathbf{W}^{-1} \mathbf{T}, \quad \mathbf{C} \equiv \mathbf{T}^T \mathbf{W}^{-1} \mathbf{P}, \quad \mathbf{D} \equiv \mathbf{T}^T \mathbf{W}^{-1} \mathbf{T},$$

we can write its inverse blockwise as follows

$$\begin{bmatrix} \mathbf{A} & \mathbf{B} \\ \mathbf{C} & \mathbf{D} \end{bmatrix}^{-1} = \begin{bmatrix} (\mathbf{A} - \mathbf{B} \mathbf{D}^{-1} \mathbf{C})^{-1} & -(\mathbf{A} - \mathbf{B} \mathbf{D}^{-1} \mathbf{C})^{-1} \mathbf{B} \mathbf{D}^{-1} \\ -(\mathbf{D} - \mathbf{C} \mathbf{A}^{-1} \mathbf{B})^{-1} \mathbf{C} \mathbf{A}^{-1} & (\mathbf{D} - \mathbf{C} \mathbf{A}^{-1} \mathbf{B})^{-1} \end{bmatrix}.$$

Thus, going back to Eq. (A.2), we can write the sky map and template amplitudes estimators as

$$\hat{\mathbf{m}} = (\mathbf{A} - \mathbf{B} \mathbf{D}^{-1} \mathbf{C})^{-1} [\mathbf{P}^T \mathbf{W}^{-1} - \mathbf{B} \mathbf{D}^{-1} \mathbf{T}^T \mathbf{W}^{-1}] \mathbf{d}, \quad (\text{A.3})$$

$$\hat{\mathbf{x}} = (\mathbf{D} - \mathbf{C} \mathbf{A}^{-1} \mathbf{B})^{-1} [\mathbf{T}^T \mathbf{W}^{-1} - \mathbf{C} \mathbf{A}^{-1} \mathbf{P}^T \mathbf{W}^{-1}] \mathbf{d}. \quad (\text{A.4})$$

We can rewrite the inverse kernels that appear at the leftmost side of the above expressions as a function of the original matrices, introducing the operators \mathbf{F}_T and \mathbf{F}_P of Chapter 6

$$\begin{aligned}
 \mathbf{A} - \mathbf{B} \mathbf{D}^{-1} \mathbf{C} &= \mathbf{P}^T \mathbf{W}^{-1} \mathbf{P} - \mathbf{P}^T \mathbf{W}^{-1} \mathbf{T} (\mathbf{T}^T \mathbf{W}^{-1} \mathbf{T})^{-1} \mathbf{T}^T \mathbf{W}^{-1} \mathbf{P} \\
 &= \mathbf{P}^T \left[\mathbf{W}^{-1} - \mathbf{W}^{-1} \mathbf{T} (\mathbf{T}^T \mathbf{W}^{-1} \mathbf{T})^{-1} \mathbf{T}^T \mathbf{W}^{-1} \right] \mathbf{P} \\
 &\equiv \mathbf{P}^T \mathbf{F}_T \mathbf{P}, \\
 \mathbf{D} - \mathbf{C} \mathbf{A}^{-1} \mathbf{B} &= \mathbf{T}^T \mathbf{W}^{-1} \mathbf{T} - \mathbf{T}^T \mathbf{W}^{-1} \mathbf{P} (\mathbf{P}^T \mathbf{W}^{-1} \mathbf{P})^{-1} \mathbf{P}^T \mathbf{W}^{-1} \mathbf{T} \\
 &= \mathbf{T}^T \left[\mathbf{W}^{-1} - \mathbf{W}^{-1} \mathbf{P} (\mathbf{P}^T \mathbf{W}^{-1} \mathbf{P})^{-1} \mathbf{P}^T \mathbf{W}^{-1} \right] \mathbf{T} \\
 &\equiv \mathbf{T}^T \mathbf{F}_P \mathbf{T}
 \end{aligned}$$

Finally, we inject these expressions in equations (A.3) and (A.4) and replace $\mathbf{B} \mathbf{D}^{-1}$ and $\mathbf{C} \mathbf{A}^{-1}$ by their original expressions to obtain the concise forms of Eq. (6.4) and Eq. (6.8)

$$\begin{aligned}
 \hat{\mathbf{m}} &= (\mathbf{P}^T \mathbf{F}_T \mathbf{P})^{-1} \mathbf{P}^T \left[\mathbf{W}^{-1} - \mathbf{W}^{-1} \mathbf{T} (\mathbf{T}^T \mathbf{W}^{-1} \mathbf{T})^{-1} \mathbf{T}^T \mathbf{W}^{-1} \right] \mathbf{d} \\
 \implies \hat{\mathbf{m}} &= (\mathbf{P}^T \mathbf{F}_T \mathbf{P})^{-1} \mathbf{P}^T \mathbf{F}_T \mathbf{d}, \tag{A.5}
 \end{aligned}$$

$$\begin{aligned}
 \hat{\mathbf{x}} &= (\mathbf{T}^T \mathbf{F}_P \mathbf{T})^{-1} \mathbf{T}^T \left[\mathbf{W}^{-1} - \mathbf{W}^{-1} \mathbf{P} (\mathbf{P}^T \mathbf{W}^{-1} \mathbf{P})^{-1} \mathbf{P}^T \mathbf{W}^{-1} \right] \mathbf{d} \\
 \implies \hat{\mathbf{x}} &= (\mathbf{T}^T \mathbf{F}_P \mathbf{T})^{-1} \mathbf{T}^T \mathbf{F}_P \mathbf{d}. \tag{A.6}
 \end{aligned}$$

To derive the expression of the map estimate in the two-step approach, we first start by demonstrating that $\mathbf{A}^{-1} \mathbf{B} (\mathbf{T}^T \mathbf{F}_P \mathbf{T})^{-1} = (\mathbf{P}^T \mathbf{F}_T \mathbf{P})^{-1} \mathbf{B} \mathbf{D}^{-1}$, we start off with the left hand side multiplied by $(\mathbf{P}^T \mathbf{F}_T \mathbf{P})$

$$\begin{aligned}
 (\mathbf{P}^T \mathbf{F}_T \mathbf{P}) \mathbf{A}^{-1} \mathbf{B} (\mathbf{T}^T \mathbf{F}_P \mathbf{T})^{-1} &= (\mathbf{P}^T \mathbf{W}^{-1} \mathbf{P} - \mathbf{P}^T \mathbf{W}^{-1} \mathbf{T} \mathbf{D}^{-1} \mathbf{T}^T \mathbf{W}^{-1} \mathbf{P}) \\
 &\quad \times (\mathbf{P}^T \mathbf{W}^{-1} \mathbf{P})^{-1} \mathbf{P}^T \mathbf{W}^{-1} \mathbf{T} (\mathbf{T}^T \mathbf{F}_P \mathbf{T})^{-1} \\
 &= \mathbf{B} \mathbf{D}^{-1} \left[\underbrace{\mathbf{T}^T (\mathbf{W}^{-1} - \mathbf{W}^{-1} \mathbf{P} \mathbf{A}^{-1} \mathbf{P}^T \mathbf{W}^{-1}) \mathbf{T}}_{= \mathbf{T}^T \mathbf{F}_P \mathbf{T}} \right] \\
 &\quad \times (\mathbf{T}^T \mathbf{F}_P \mathbf{T})^{-1} \\
 &= \mathbf{B} \mathbf{D}^{-1}. \text{(Q.E.D.)}
 \end{aligned}$$

Therefore we rewrite the map estimate in Eq. (A.3) as

$$\hat{\mathbf{m}} = (\mathbf{P}^T \mathbf{F}_T \mathbf{P})^{-1} \mathbf{P}^T \mathbf{W}^{-1} \mathbf{d} - \mathbf{A}^{-1} \mathbf{B} (\mathbf{T}^T \mathbf{F}_P \mathbf{T})^{-1} \mathbf{T}^T \mathbf{W}^{-1} \mathbf{d}, \tag{A.7}$$

By virtue of the Sherman–Morrison–Woodbury formula, we have

$$\begin{aligned} (\mathbf{P}^T \mathbf{F}_T \mathbf{P})^{-1} &= (\mathbf{A} - \mathbf{B} \mathbf{D}^{-1} \mathbf{C})^{-1} = \mathbf{A}^{-1} + \mathbf{A}^{-1} \mathbf{B} (\mathbf{D} - \mathbf{C} \mathbf{A}^{-1} \mathbf{B})^{-1} \mathbf{C} \mathbf{A}^{-1} \\ &= \mathbf{A}^{-1} + \mathbf{A}^{-1} \mathbf{B} (\mathbf{T}^T \mathbf{F}_P \mathbf{T})^{-1} \mathbf{T}^T \mathbf{W}^{-1} \mathbf{P} \mathbf{A}^{-1}, \end{aligned}$$

We inject the above expression in (A.7)

$$\begin{aligned} \hat{\mathbf{m}} &= \mathbf{A}^{-1} \mathbf{P}^T \mathbf{W}^{-1} \mathbf{d} - \mathbf{A}^{-1} \mathbf{B} (\mathbf{T}^T \mathbf{F}_P \mathbf{T})^{-1} \mathbf{T}^T (\mathbf{W}^{-1} - \mathbf{W}^{-1} \mathbf{P} \mathbf{A}^{-1} \mathbf{P}^T \mathbf{W}^{-1}) \mathbf{d} \\ &= \mathbf{A}^{-1} \mathbf{P}^T \mathbf{W}^{-1} \mathbf{d} - \mathbf{A}^{-1} \mathbf{P}^T \mathbf{W}^{-1} \mathbf{T} \underbrace{(\mathbf{T}^T \mathbf{F}_P \mathbf{T})^{-1} \mathbf{T}^T \mathbf{F}_P \mathbf{d}}_{= \hat{\mathbf{x}}}. \end{aligned}$$

Finally we obtain the two-step sky map estimate in the form of Eq. (6.9)

$$\hat{\mathbf{m}} = (\mathbf{P}^T \mathbf{W}^{-1} \mathbf{P})^{-1} \mathbf{P}^T \mathbf{W}^{-1} (\mathbf{d} - \mathbf{T} \hat{\mathbf{x}}). \quad (\text{A.8})$$

APPENDIX B

Customized collective communication operations

We give here a brief description of the basic ideas of the collective communication operations implemented in MIDAPACK and tested in Chapter 7. We refer to these as `MPI_Allreduce`, `Ring` and `Butterfly` schemes. Other schemes are also implemented in MIDAPACK, but they did not perform well in some preliminary cases we tested and hence were not considered in our studies. The content of this appendix is based on written notes from an unpublished paper: Pierre Cargamel, Laura Grigori, Radek Stompor, *Communication algorithms for extreme-scale cosmic microwave background map-making*.

Let us consider that the time-domain data is distributed among n_{proc} processes. Each process i is assigned a segment \mathcal{T}_i from the total time-domain data set. This segment corresponds to a subset of pixels of the observed map, defined as

$$\Pi_i \equiv \{p : \exists t \in \mathcal{T}_i, \mathbf{P}_{t,p} \neq 0\}, \quad (\text{B.1})$$

where $\mathbf{P}_{t,p}$ is the element of the pointing matrix \mathbf{P} at row t and column p . The subsets $\{\Pi_i\}_{i=0, \dots, n_{proc}}$ typically overlap, which means that the pointing operation, $\mathbf{P}^T \mathbf{t}$ (\mathbf{t} stands for the time-ordered data vector), requires collective communication between the processes. We explain the different implemented approaches below.

B.1 Standard `MPI_Allreduce`

The implementation of this first approach is straightforward: each process i first performs a local product of the form $\mathbf{m}|_{\Pi_i} = \mathbf{P}|_{\mathcal{T}_i}^T \mathbf{t}|_{\mathcal{T}_i}$. This step is obviously needed for all approaches. Then it extends its local map to encompass all the observed pixels (possibly more) by allocating a contiguous block of pixels of the

form

$$\Pi_{global} = \llbracket \min\{p \in \Pi_{obs}\}, \max\{p \in \Pi_{obs}\} \rrbracket, \quad \text{with } \Pi_{obs} = \bigcup_{i=0}^{n_{proc}-1} \Pi_i, \quad (\text{B.2})$$

with the extra pixels padded with zeros. A global reduction is then performed by calling the standard `MPI_Allreduce` operation provided by the MPI library. Note that this simple implementation is not necessarily the most economic choice, one can modify it to extend the map only to the subset of shared pixels between at least two distinct processes. In all cases however, the communication traffic in the `MPI_Allreduce` approach is significantly redundant, sending and receiving way more data volumes than what is strictly needed. The next two approaches address this issue, precisely computing the subset of pixels to send and receive to avoid any superfluous data exchanges.

B.2 Ring scheme

In the ring scheme, the collective communication is performed in a sequence of $n_{proc} - 1$ steps. At each step k , each process i sends data to the process $i_{send} = \text{mod}(i + (k + 1), n_{proc})$ and receives data from the process $i_{recv} = \text{mod}(i - (k + 1), n_{proc})$, the received data is subsequently used to update a stored copy of the local map. For any exchange between two processes i and j , the data transfer is limited to the intersection of their local maps, corresponding to pixels in the subset $\Pi_{ij} = \Pi_i \cap \Pi_j$. After $n_{proc} - 1$ steps, all processes have communicated with each other, and the reduce operation is complete. A diagram illustrating the communication steps on a simple example with 4 processes is presented in figure B.1.

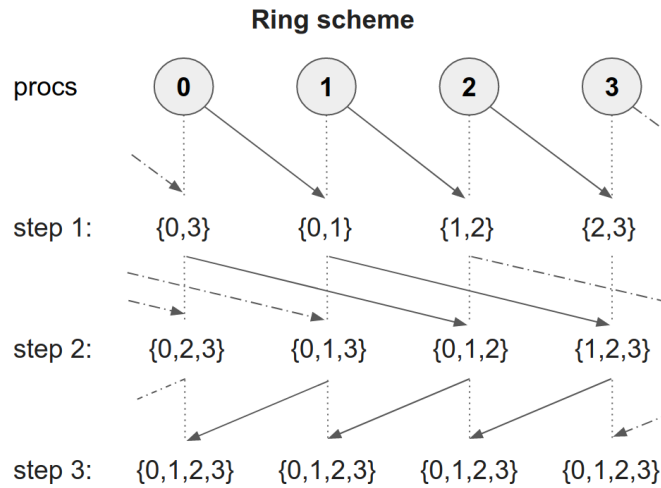


Figure B.1 – Ring communication pattern. *Credit:* Adapted from Pierre Cargemel.

This algorithm requires some precomputation, whereby every process determines the subset of pixels of the local maps intersections, to know exactly what to send and receive at each step. This is also done in $n_{proc} - 1$ steps: every process i sends its subset of pixels to i_{send} and receives a subset from i_{recv} from which it

computes and stores the intersection. The precomputation stage requires heavier communication traffic than the algorithm itself but it is performed only once.

B.3 Butterfly scheme

In this approach, the idea is to further improve on the previous scheme by not only avoiding redundant data transfers, but also by reducing the number of communication steps to avoid scaling performance issues for large jobs. The Butterfly scheme allows to perform the collective communication operation in $\log_2 n_{proc}$ steps instead of n_{proc} . The current implementation is limited to cases where n_{proc} is a power of 2 but can be generalized in future updates.

This algorithm also requires a precomputation stage where for each process i and for each communication step k three sets of pixels are computed: (1) the set of pixels to send, $\mathcal{S}(i, k)$; (2) the set of pixels that the process will receive, $\mathcal{R}(i, k)$; (3) the set of pixels to retain in memory of the process after the communication step is complete, $\mathcal{K}(i, k)$. After precomputation the algorithm performs the following instructions:

```

for all process  $i$  in parallel do
     $\mathbf{m}_i^{(0)} = \mathbf{P}|_{\mathcal{T}_i}^T \mathbf{t}|_{\mathcal{T}_i}$ 
    for  $k = 0$  to  $\log_2 n_{proc} - 1$  do
        send message to process:  $\text{mod } (i + 2^k, n_{proc})$  containing,
             $\mathbf{m}_i^{(k)}|_{\mathcal{S}(i,k)}$ 
        receive message from process:  $r = \text{mod } (i - 2^k + n_{proc}, n_{proc})$ 
            containing,  $\mathbf{m}_r^{(k)}|_{\mathcal{R}(i,k)}$ 
        update  $\mathbf{m}_i^{(k+1)} = \mathbf{m}_i^{(k)} + \mathbf{m}_r^{(k)}$ 
        trim  $\mathbf{m}_i^{(k+1)} = \mathbf{m}_i^{(k+1)}|_{\mathcal{K}(i,k)}$ 
    end
    return  $\mathbf{m}_i^{(\log_2 n_{proc})}$  defined on  $\Pi_i$ 
end

```

Algorithm 4: The Butterfly algorithm.

A simple illustration of the communication pattern for 8 processes is also given in figure B.2.

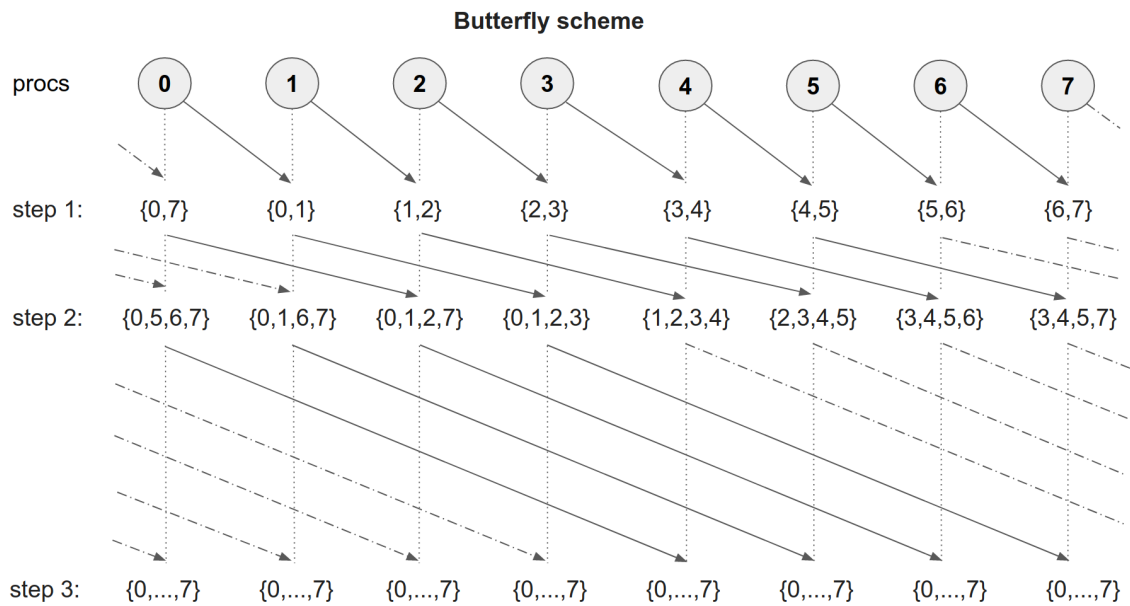


Figure B.2 – Butterfly communication pattern. *Credit:* Adapted from Pierre Cargemel.

APPENDIX C

Analytical calculation of the standard IQU map-making and the pair-differencing case

Consider the following data model

$$\mathbf{d} = \mathbf{P} \mathbf{m} + \mathbf{n}, \quad (\text{C.1})$$

We can further decompose the noise term to

$$\mathbf{n} = \mathbf{n}^{\text{atm.}} + \mathbf{n}^{\text{inst.}}, \quad (\text{C.2})$$

where $\mathbf{n}^{\text{atm.}}$ is the atmospheric contribution to the noise, and $\mathbf{n}^{\text{inst.}}$ is the instrumental noise which is assumed to be uncorrelated between detectors.

Let us denote the even index detectors in the pairs by the subscript \parallel and their orthogonal by \perp . We can construct sums and differences of these pairs as

$$\mathbf{d}_+ = \frac{1}{2}(\mathbf{d}_\parallel + \mathbf{d}_\perp) \quad (\text{C.3})$$

$$\mathbf{d}_- = \frac{1}{2}(\mathbf{d}_\parallel - \mathbf{d}_\perp) \quad (\text{C.4})$$

We introduce a linear operator, \mathcal{D} , which constructs these sums and differences out of the original tod vector \mathbf{d}

$$\mathcal{D} \mathbf{d} = \begin{bmatrix} \mathbf{d}_+ \\ \mathbf{d}_- \end{bmatrix} \quad (\text{C.5})$$

We have $\mathcal{D}^T \mathcal{D} = \mathbf{I}$, and thus we can write a general unbiased map estimator with

a weight, \mathbf{W} , as

$$\hat{\mathbf{m}} = (\mathbf{P}^T \mathbf{W}^{-1} \mathbf{P})^{-1} \mathbf{P}^T \mathbf{W}^{-1} \mathbf{d} \quad (\text{C.6})$$

$$= \left((\mathcal{D} \mathbf{P})^T (\underbrace{\mathcal{D} \mathbf{W}^{-1} \mathcal{D}^T}_{\equiv \mathcal{W}^{-1}}) (\mathcal{D} \mathbf{P}) \right)^{-1} (\mathcal{D} \mathbf{P})^T (\mathcal{D} \mathbf{W}^{-1} \mathcal{D}^T) \mathcal{D} \mathbf{d} \quad (\text{C.7})$$

$$= \left(\begin{bmatrix} \mathbf{P}_+^T & \mathbf{P}_-^T \end{bmatrix} \mathcal{W}^{-1} \begin{bmatrix} \mathbf{P}_+ \\ \mathbf{P}_- \end{bmatrix} \right)^{-1} \begin{bmatrix} \mathbf{P}_+^T & \mathbf{P}_-^T \end{bmatrix} \mathcal{W}^{-1} \begin{bmatrix} \mathbf{d}_+ \\ \mathbf{d}_- \end{bmatrix} \quad (\text{C.8})$$

The transformed weight matrix, \mathcal{W} , reads

$$\mathcal{W} = \begin{bmatrix} \mathcal{D}_+ \mathbf{W} \mathcal{D}_+^T & \mathcal{D}_+ \mathbf{W} \mathcal{D}_-^T \\ \mathcal{D}_- \mathbf{W} \mathcal{D}_+^T & \mathcal{D}_- \mathbf{W} \mathcal{D}_-^T \end{bmatrix} \equiv \begin{bmatrix} \mathbf{W}_{++} & \mathbf{W}_{+-} \\ \mathbf{W}_{-+} & \mathbf{W}_{--} \end{bmatrix} \quad (\text{C.9})$$

Assuming both \mathbf{W}_{++} and \mathbf{W}_{--} are invertible, its block-wise inverse takes the following form

$$\mathcal{W}^{-1} = \begin{bmatrix} \tilde{\mathbf{W}}_{++}^{-1} & \tilde{\mathbf{W}}_{+-}^{-1} \\ \tilde{\mathbf{W}}_{-+}^{-1} & \tilde{\mathbf{W}}_{--}^{-1} \end{bmatrix}$$

with

$$\tilde{\mathbf{W}}_{++}^{-1} \equiv (\mathbf{W}_{++} - \mathbf{W}_{+-} \mathbf{W}_{--}^{-1} \mathbf{W}_{-+})^{-1} \quad (\text{C.10})$$

$$\tilde{\mathbf{W}}_{+-}^{-1} \equiv -\tilde{\mathbf{W}}_{++}^{-1} \mathbf{W}_{+-} \mathbf{W}_{--}^{-1} \quad (\text{C.11})$$

$$\tilde{\mathbf{W}}_{--}^{-1} \equiv (\mathbf{W}_{--} - \mathbf{W}_{-+} \mathbf{W}_{++}^{-1} \mathbf{W}_{+-})^{-1} \quad (\text{C.12})$$

$$\tilde{\mathbf{W}}_{-+}^{-1} \equiv -\tilde{\mathbf{W}}_{--}^{-1} \mathbf{W}_{-+} \mathbf{W}_{++}^{-1} \quad (\text{C.13})$$

Injecting back in the map estimator equation, we get

$$\begin{aligned} \hat{\mathbf{m}} &= \left(\begin{bmatrix} \mathbf{P}_+^T \tilde{\mathbf{W}}_{++}^{-1} + \mathbf{P}_-^T \tilde{\mathbf{W}}_{-+}^{-1} & \mathbf{P}_+^T \tilde{\mathbf{W}}_{+-}^{-1} + \mathbf{P}_-^T \tilde{\mathbf{W}}_{--}^{-1} \end{bmatrix} \begin{bmatrix} \mathbf{P}_+ \\ \mathbf{P}_- \end{bmatrix} \right)^{-1} \\ &\quad \times \begin{bmatrix} \mathbf{P}_+^T \tilde{\mathbf{W}}_{++}^{-1} + \mathbf{P}_-^T \tilde{\mathbf{W}}_{-+}^{-1} & \mathbf{P}_+^T \tilde{\mathbf{W}}_{+-}^{-1} + \mathbf{P}_-^T \tilde{\mathbf{W}}_{--}^{-1} \end{bmatrix} \begin{bmatrix} \mathbf{d}_+ \\ \mathbf{d}_- \end{bmatrix} \end{aligned}$$

Given that \mathbf{P}_+ is a projection on the intensity pixels, and \mathbf{P}_- is a projection on the Q and U pixels. If we restrict these matrices on the domain of pixels where they are not identically zero. We can write the map estimator equation in a block form separating between I and Q/U

$$\begin{bmatrix} \hat{\mathbf{m}}_I \\ \hat{\mathbf{m}}_{Q,U} \end{bmatrix} = \begin{bmatrix} \mathbf{P}_+^T \tilde{\mathbf{W}}_{++}^{-1} \mathbf{P}_+ & \mathbf{P}_+^T \tilde{\mathbf{W}}_{+-}^{-1} \mathbf{P}_- \\ \mathbf{P}_-^T \tilde{\mathbf{W}}_{-+}^{-1} \mathbf{P}_+ & \mathbf{P}_-^T \tilde{\mathbf{W}}_{--}^{-1} \mathbf{P}_- \end{bmatrix}^{-1} \begin{bmatrix} \mathbf{P}_+^T \tilde{\mathbf{W}}_{++}^{-1} \mathbf{d}_+ + \mathbf{P}_+^T \tilde{\mathbf{W}}_{+-}^{-1} \mathbf{d}_- \\ \mathbf{P}_-^T \tilde{\mathbf{W}}_{-+}^{-1} \mathbf{d}_+ + \mathbf{P}_-^T \tilde{\mathbf{W}}_{--}^{-1} \mathbf{d}_- \end{bmatrix}$$

As in Appendix A, let us introduce the quantities, \mathbf{A} , \mathbf{B} , \mathbf{C} and \mathbf{D} such that

$$\mathbf{A} \equiv \mathbf{P}_+^T \tilde{\mathbf{W}}_{++}^{-1} \mathbf{P}_+, \quad \mathbf{B} \equiv \mathbf{P}_+^T \tilde{\mathbf{W}}_{+-}^{-1} \mathbf{P}_-, \quad \mathbf{C} \equiv \mathbf{P}_-^T \tilde{\mathbf{W}}_{-+}^{-1} \mathbf{P}_+, \quad \mathbf{D} \equiv \mathbf{P}_-^T \tilde{\mathbf{W}}_{--}^{-1} \mathbf{P}_-,$$

The block-wise inversion of the leftmost kernel in the map estimator equation yields

$$\begin{aligned}\hat{\mathbf{m}}_I &= (\mathbf{A} - \mathbf{B} \mathbf{D}^{-1} \mathbf{C})^{-1} \left[\mathbf{P}_+^T \left(\tilde{\mathbf{W}}_{++}^{-1} \mathbf{d}_+ + \tilde{\mathbf{W}}_{+-}^{-1} \mathbf{d}_- \right) - \mathbf{B} \mathbf{D}^{-1} \mathbf{P}_-^T \left(\tilde{\mathbf{W}}_{-+}^{-1} \mathbf{d}_+ + \tilde{\mathbf{W}}_{--}^{-1} \mathbf{d}_- \right) \right], \\ \hat{\mathbf{m}}_{Q,U} &= (\mathbf{D} - \mathbf{C} \mathbf{A}^{-1} \mathbf{B})^{-1} \left[\mathbf{P}_-^T \left(\tilde{\mathbf{W}}_{-+}^{-1} \mathbf{d}_+ + \tilde{\mathbf{W}}_{--}^{-1} \mathbf{d}_- \right) - \mathbf{C} \mathbf{A}^{-1} \mathbf{P}_+^T \left(\tilde{\mathbf{W}}_{++}^{-1} \mathbf{d}_+ + \tilde{\mathbf{W}}_{+-}^{-1} \mathbf{d}_- \right) \right],\end{aligned}$$

We can write the left factors as

$$\begin{aligned}\mathbf{A} - \mathbf{B} \mathbf{D}^{-1} \mathbf{C} &= \mathbf{P}_+^T \tilde{\mathbf{W}}_{++}^{-1} \mathbf{P}_+ - \mathbf{P}_+^T \tilde{\mathbf{W}}_{+-}^{-1} \mathbf{P}_- \left(\mathbf{P}_-^T \tilde{\mathbf{W}}_{--}^{-1} \mathbf{P}_- \right)^{-1} \mathbf{P}_-^T \tilde{\mathbf{W}}_{-+}^{-1} \mathbf{P}_+ \\ &= \mathbf{P}_+^T \left(\tilde{\mathbf{W}}_{++}^{-1} - \tilde{\mathbf{W}}_{+-}^{-1} \mathbf{P}_- \left(\mathbf{P}_-^T \tilde{\mathbf{W}}_{--}^{-1} \mathbf{P}_- \right)^{-1} \mathbf{P}_-^T \tilde{\mathbf{W}}_{-+}^{-1} \right) \mathbf{P}_+ \\ &\equiv \mathbf{P}_+^T \mathbf{F}_{\mathbf{P}_1} \mathbf{P}_+ \\ \mathbf{D} - \mathbf{C} \mathbf{A}^{-1} \mathbf{B} &= \mathbf{P}_-^T \tilde{\mathbf{W}}_{--}^{-1} \mathbf{P}_- - \mathbf{P}_-^T \tilde{\mathbf{W}}_{-+}^{-1} \mathbf{P}_+ \left(\mathbf{P}_+^T \tilde{\mathbf{W}}_{++}^{-1} \mathbf{P}_+ \right)^{-1} \mathbf{P}_+^T \tilde{\mathbf{W}}_{+-}^{-1} \mathbf{P}_- \\ &= \mathbf{P}_-^T \left(\tilde{\mathbf{W}}_{--}^{-1} - \tilde{\mathbf{W}}_{-+}^{-1} \mathbf{P}_+ \left(\mathbf{P}_+^T \tilde{\mathbf{W}}_{++}^{-1} \mathbf{P}_+ \right)^{-1} \mathbf{P}_+^T \tilde{\mathbf{W}}_{+-}^{-1} \right) \mathbf{P}_- \\ &\equiv \mathbf{P}_-^T \mathbf{F}_{\mathbf{P}_2} \mathbf{P}_-\end{aligned}$$

Injecting in the map estimators equation we find

$$\begin{aligned}\hat{\mathbf{m}}_I &= (\mathbf{P}_+^T \mathbf{F}_{\mathbf{P}_1} \mathbf{P}_+)^{-1} \left[\mathbf{P}_+^T \left(\tilde{\mathbf{W}}_{++}^{-1} \mathbf{d}_+ + \tilde{\mathbf{W}}_{+-}^{-1} \mathbf{d}_- \right) - \right. \\ &\quad \left. \mathbf{P}_+^T \tilde{\mathbf{W}}_{+-}^{-1} \mathbf{P}_- \left(\mathbf{P}_-^T \tilde{\mathbf{W}}_{--}^{-1} \mathbf{P}_- \right)^{-1} \mathbf{P}_-^T \left(\tilde{\mathbf{W}}_{-+}^{-1} \mathbf{d}_+ + \tilde{\mathbf{W}}_{--}^{-1} \mathbf{d}_- \right) \right] \\ &= (\mathbf{P}_+^T \mathbf{F}_{\mathbf{P}_1} \mathbf{P}_+)^{-1} \mathbf{P}_+^T \left[\mathbf{F}_{\mathbf{P}_1} \mathbf{d}_+ + \underbrace{\left(\tilde{\mathbf{W}}_{+-}^{-1} - \tilde{\mathbf{W}}_{+-}^{-1} \mathbf{P}_- \left(\mathbf{P}_-^T \tilde{\mathbf{W}}_{--}^{-1} \mathbf{P}_- \right)^{-1} \mathbf{P}_-^T \tilde{\mathbf{W}}_{-+}^{-1} \right)}_{\equiv \mathbf{F}_{\mathbf{P}_{12}}} \mathbf{d}_- \right] \\ &= (\mathbf{P}_+^T \mathbf{F}_{\mathbf{P}_1} \mathbf{P}_+)^{-1} \mathbf{P}_+^T [\mathbf{F}_{\mathbf{P}_1} \mathbf{d}_+ + \mathbf{F}_{\mathbf{P}_{12}} \mathbf{d}_-], \\ \hat{\mathbf{m}}_{Q,U} &= (\mathbf{P}_-^T \mathbf{F}_{\mathbf{P}_2} \mathbf{P}_-)^{-1} \left[\mathbf{P}_-^T \left(\tilde{\mathbf{W}}_{-+}^{-1} \mathbf{d}_+ + \tilde{\mathbf{W}}_{--}^{-1} \mathbf{d}_- \right) - \right. \\ &\quad \left. \mathbf{P}_-^T \tilde{\mathbf{W}}_{-+}^{-1} \mathbf{P}_+ \left(\mathbf{P}_+^T \tilde{\mathbf{W}}_{++}^{-1} \mathbf{P}_+ \right)^{-1} \mathbf{P}_+^T \left(\tilde{\mathbf{W}}_{++}^{-1} \mathbf{d}_+ + \tilde{\mathbf{W}}_{+-}^{-1} \mathbf{d}_- \right) \right] \\ &= (\mathbf{P}_-^T \mathbf{F}_{\mathbf{P}_2} \mathbf{P}_-)^{-1} \mathbf{P}_-^T \left[\mathbf{F}_{\mathbf{P}_2} \mathbf{d}_+ + \underbrace{\left(\tilde{\mathbf{W}}_{-+}^{-1} - \tilde{\mathbf{W}}_{-+}^{-1} \mathbf{P}_+ \left(\mathbf{P}_+^T \tilde{\mathbf{W}}_{++}^{-1} \mathbf{P}_+ \right)^{-1} \mathbf{P}_+^T \tilde{\mathbf{W}}_{+-}^{-1} \right)}_{\equiv \mathbf{F}_{\mathbf{P}_{21}}} \mathbf{d}_- \right] \\ &= (\mathbf{P}_-^T \mathbf{F}_{\mathbf{P}_2} \mathbf{P}_-)^{-1} \mathbf{P}_-^T [\mathbf{F}_{\mathbf{P}_2} \mathbf{d}_+ + \mathbf{F}_{\mathbf{P}_{21}} \mathbf{d}_-],\end{aligned}$$

The operators $\mathbf{F}_{\mathbf{P}_{12}}$ and $\mathbf{F}_{\mathbf{P}_{21}}$ respectively filter the polarization and intensity sky signals

$$\begin{aligned}\mathbf{F}_{\mathbf{P}_{12}} \mathbf{P}_- \mathbf{m}_{Q,U} &= \mathbf{0}, \\ \mathbf{F}_{\mathbf{P}_{21}} \mathbf{P}_+ \mathbf{m}_I &= \mathbf{0},\end{aligned}$$

we can decompose the polarization map estimate into sky signal and noise residual

$$\hat{\mathbf{m}}_{Q,U} = \mathbf{m}_{Q,U} + (\mathbf{P}_-^T \mathbf{F}_{\mathbf{P}_2} \mathbf{P}_-)^{-1} \mathbf{P}_-^T [\mathbf{F}_{\mathbf{P}_2} \mathbf{n}_- + \mathbf{F}_{\mathbf{P}_{21}} \mathbf{n}_+], \quad (\text{C.14})$$

Let us now write the different blocks of the actual noise covariance as a function of those of the even and odd index detectors

$$\mathbf{N}_{++} = \frac{1}{4} (\mathbf{N}_{\parallel} + 2 \mathbf{N}_{\parallel \times \perp} + \mathbf{N}_{\perp}) \quad (\text{C.15})$$

$$\mathbf{N}_{--} = \frac{1}{4} (\mathbf{N}_{\parallel} - 2 \mathbf{N}_{\parallel \times \perp} + \mathbf{N}_{\perp}) \quad (\text{C.16})$$

$$\mathbf{N}_{-+} = \mathbf{N}_{+-} = \frac{1}{4} (\mathbf{N}_{\parallel} - \mathbf{N}_{\perp}) \quad (\text{C.17})$$

Given that both detectors in a single pair see the same atmospheric signal we have

$$\mathbf{N}_{\parallel}^{\text{atm}} = \mathbf{N}_{\perp}^{\text{atm}} = \mathbf{N}_{\parallel \times \perp}^{\text{atm}} \quad (\text{C.18})$$

Therefore if we decompose the noise covariance into the instrumental (assumed to be uncorrelated between detectors) and atmospheric parts we get

$$\mathbf{N}_{++} = \mathbf{N}_{\parallel}^{\text{atm}} + \frac{1}{4} (\mathbf{N}_{\parallel}^{\text{inst.}} + \mathbf{N}_{\perp}^{\text{inst.}}) \quad (\text{C.19})$$

$$\mathbf{N}_{--} = \frac{1}{4} (\mathbf{N}_{\parallel}^{\text{inst.}} + \mathbf{N}_{\perp}^{\text{inst.}}) \quad (\text{C.20})$$

$$\mathbf{N}_{-+} = \frac{1}{4} (\mathbf{N}_{\parallel}^{\text{inst.}} - \mathbf{N}_{\perp}^{\text{inst.}}) \quad (\text{C.21})$$

A true maximum likelihood map ($\mathbf{W} = \mathbf{N}$) will be different in general than the pair-differencing case whenever the instrumental noise covariance is not the same for even and odd detectors in the pairs.

Whenever this difference is small, i.e. $\mathbf{N}_{-+} \simeq \mathbf{0}$, then we have

$$\mathbf{F}_{\mathbf{P}1} = \mathbf{N}_{++}^{-1} \quad (\text{C.22})$$

$$\mathbf{F}_{\mathbf{P}2} = \mathbf{N}_{--}^{-1} \quad (\text{C.23})$$

$$\mathbf{F}_{\mathbf{P}12} = \mathbf{F}_{\mathbf{P}21} = \mathbf{0}, \quad (\text{C.24})$$

and the polarization map is exactly the same as that in pair-differencing

$$\hat{\mathbf{m}}_{Q,U} = \mathbf{m}_{Q,U} + (\mathbf{P}_{-}^T \mathbf{N}_{--}^{-1} \mathbf{P}_{-})^{-1} \mathbf{P}_{-}^T \mathbf{N}_{--}^{-1} \mathbf{n}_{-}, \quad (\text{C.25})$$

In the case of our map estimators used in section 8.2.3 however, the weight matrix is block diagonal and ignores all cross-correlations between detectors, in particular, those between detectors in the same pair, which means the atmospheric weights do not cancel out in our approximate of the pair difference noise covariance, as it did in the true noise covariance. We get, assuming we are able to compute each of the detector covariance blocks perfectly

$$\mathbf{W}_{++} = \mathbf{W}_{--} = \frac{1}{2} \mathbf{W}_{\parallel}^{\text{atm.}} + \frac{1}{4} (\mathbf{N}_{\parallel}^{\text{inst.}} + \mathbf{N}_{\perp}^{\text{inst.}}) \quad (\text{C.26})$$

$$\mathbf{W}_{+-} = \mathbf{W}_{-+} = \frac{1}{4} (\mathbf{N}_{\parallel}^{\text{inst.}} - \mathbf{N}_{\perp}^{\text{inst.}}) \quad (\text{C.27})$$

$$\text{with } \mathbf{W}_{\parallel}^{\text{atm.}} = (\delta_{ij} \mathbf{N}_{\parallel,ij}^{\text{atm.}})_{i,j \in \llbracket 0, \mathcal{N}_{\text{pair}}-1 \rrbracket}, \quad (\text{C.28})$$

The weights are thus not optimal, and we do not get the maximum likelihood or pair-differencing estimate even if $\mathbf{N}_{-+} \simeq \mathbf{0}$. Given that the instrumental noise

PSDs are drawn with a 1% dispersion, the effect of N_{+-} should be negligible relative to the effect of $\frac{1}{2}\mathbf{W}_{||}^{\text{atm.}}$ in our noise weights. However, the latter can be easily dealt with by including cross-correlations only between detectors in the same pair, in the noise model. Or somewhat counter intuitively by including only the instrumental noise blocks in our weights, then if the effect of N_{+-} is indeed negligible, we should get exactly the maximum likelihood/pair-differencing polarization map.

The exact same arguments also apply to the templates marginalization approach, however the weights now are diagonal and multiplied by the same filtering operator $\mathbf{F}_{T||} = \mathbf{F}_{T\perp}$. The template map-making should also include both the effects of the atmospheric weight and the non-vanishing difference between noise weights of even and odd detectors in the pairs, corresponding respectively to the first and second terms of the noise residual in Eq. (C.14). However, the effect of the first term should not be as pronounced as in the down weighting approach, since it concerns only the weighting of the non-filtered high frequency noise (approximately white if our templates are modelling the low frequency noise sufficiently well). The second term is also better mitigated in the template approach, as the filtering partially subtracts some of the correlated modes of atmosphere, which are not accounted for in the down weighting approach. This is in accordance with our results of section 8.2.3 where we see that polynomial filters perform better than the down weighting with a factor 2 in the noise power spectra, and $n_{\text{poly}} = 6$ is less than 10% close to the pair-differencing case.

Résumé détaillé de la thèse en français

D.1 Introduction générale

La cosmologie moderne connaît un essor sans précédent avec le développement d'expériences multiples permettant de sonder l'Univers par différentes voies. Grâce à la précision et la complémentarité des observations en cours ou en préparation, les années à venir promettent de nombreuses avancées scientifiques, en particulier dans notre compréhension du modèle cosmologique standard et ses extensions possibles, ainsi que les multiples liens qui existent avec la physique des particules. Dans ce contexte, les observations du fond diffus cosmologique (Cosmic Microwave Background - CMB) jouent un rôle fondamental. Ces dernières ont permis l'établissement du modèle cosmologique standard, avec des mesures des paramètres sous-jacents atteignant des précisions de l'ordre du pour cent. L'objectif actuel de ces observations est non seulement d'améliorer la précision des mesures précédentes mais également de pousser les limites de nos modèles standards, par exemple, à travers la recherche des ondes gravitationnelles primordiales, la recherche des particules reliques non prédites par le modèle standard de la physique des particules, la mesure des masses absolues des neutrinos et leur hiérarchie, ou encore en apportant des éléments d'information sur la nature de l'énergie noire et de la matière noire.

Ces objectifs scientifiques dépendent des mesures de signaux extrêmement faibles, et par conséquent requièrent une grande sensibilité des instruments déployés. Par conséquent l'analyse des données délivrées par ces télescopes devient tout aussi complexe, devant permettre de traiter de gros volumes de données tout en tenant compte de divers effets provenant de sources astrophysiques, ou des instruments eux-même et leur environnement.

Dans cette thèse, on s'intéresse à l'une des premières étapes de l'analyse qui consiste à reconstruire les cartes du ciel à partir des données brutes enregistrées par les détecteurs après les étapes initiales de calibration et de pré-traitement. L'objectif est de pouvoir effectuer une estimation non-biaisée des cartes du ciel

tout en filtrant les contributions dues à des signaux parasites présents dans les données. Pour ce faire, je développe un environnement logiciel qui incorpore plusieurs approches de cartographie du ciel, en s'appuyant sur une implémentation massivement parallèle permettant d'utiliser les capacités des supercalculateurs pour traiter simultanément de gros volumes de données, et qui déploie plusieurs méthodes numériques pour une résolution efficace des systèmes linéaires rencontrés. Plusieurs applications sont considérées ensuite, pour valider les performances des méthodes implémentées, et évaluer l'impact de certains effets systématiques.

Cette thèse se compose de huit chapitres. Les quatre premiers chapitres introduisent le contexte scientifique global, avec une introduction au modèle standard du Big Bang, la théorie de l'inflation, la physique du fond diffus cosmologique et les outils statistiques qui permettent l'interprétation des observations, ainsi que l'approche multi-sonde de la cosmologie. Le chapitre 5 présente un aperçu global des observatoires du fond diffus cosmologique, et une discussion détaillée des méthodes d'analyse de données et des défis majeurs rencontrés par les expériences actuelles et futures. Le chapitre 6 se consacre au problème de cartographie du CMB, et introduit le formalisme générale et la méthodologie adoptée dans ce travail de thèse. Le chapitre 7 se focalise sur les aspects computationnels du problème, avec une description détaillée de l'environnement logiciel développé, et les méthodes numériques utilisées pour la résolution des systèmes linéaires. Le chapitre 8 présente diverses applications des méthodes développées sur des données simulées d'observations au sol et d'observations satellites.

D.2 Contenu par chapitre

Le premier chapitre commence par une description d'un univers homogène en partant du principe cosmologique, puis en introduisant graduellement des notions de base en cosmologie. On discute ensuite la dynamique de l'expansion de l'Univers en la mettant en relation avec son équation d'état. On introduit, par la suite, très brièvement le modèle standard de la physique des particules, suivi d'une discussion plus détaillée des divers contenus énergétiques de l'Univers (photons, baryons, neutrinos, matière noire, et énergie noire). Le chapitre se poursuit avec une discussion de l'histoire thermique de l'Univers, et s'achève avec quelques questions ouvertes du modèle standard.

Le second chapitre se focalise sur l'une des solutions proposées actuellement pour résoudre certaines des questions ouvertes majeures du modèle standard du Big Bang, à savoir, l'inflation. Cette dernière propose l'existence d'une phase d'expansion exponentielle lors des premiers instants du Big Bang. Le chapitre commence par expliquer comment cela résout-il certains des problèmes relevés au chapitre précédent, avant de procéder à une description plus quantitative de cette phase d'expansion rapide. La suite du chapitre se consacre à l'introduction des perturbations cosmologiques à l'origine des structures dans l'Univers, et comment peuvent-elles être générées par l'inflation. On conclut le chapitre par une discussion de différentes observations à l'appui du scénario de l'inflation, et les contraintes qui ont été dérivées sur un certain nombre de modèles inflationnaires.

Le troisième chapitre aborde la description statistique des anisotropies du

fond diffus cosmologique en température et en polarisation, notamment en introduisant la décomposition en modes E/B du champ de polarisation. Il s'ensuit une discussion de l'évolution des perturbations cosmologiques avec les équations de Boltzmann, en dérivant certains des effets majeurs impactant les observations (effet Sachs-Wolfe, oscillations acoustiques, amortissement de Silk). Les anisotropies secondaires (Effet Sunyaev Zel'dovich, Lentillage gravitationnel faible) sont également abordées, en s'intéressant notamment à l'information cosmologique qu'ils peuvent contenir. Dans la suite du chapitre, on s'intéresse à l'état de l'art des contraintes observationnelles sur différents paramètres cosmologiques, ainsi que différents effets physiques associés à ces paramètres. Le chapitre se conclut par une discussion d'autres applications possibles du CMB et leur intérêt scientifique: la recherche des non-Gaussianités, les distortions spectrales, l'étude des amas de galaxies avec l'effet SZ, et la biréfringence cosmique.

Le quatrième chapitre présente un aperçu global d'autres observables d'intérêt en cosmologie et leur complémentarité avec les observations du CMB. Les types d'observations considérées dans ce chapitre sont les relevés de galaxies et l'étude des structures à grande échelle, la forêt de Lyman-alpha, les observations de la raie à 21-cm de l'hydrogène, les supernovae de type Ia avec notamment une brève discussion de la tension actuelle sur la constante de Hubble, et enfin les ondes gravitationnelles.

Dans le cinquième chapitre, je présente d'abord un aperçu de l'historique des expériences du CMB, ainsi que des expériences actuellement en cours en s'arrêtant plus longuement sur les expériences dont je fais partie, soulignant les multiples contributions qui ont été faites et publiées durant ma période de thèse. Deux expériences futures majeures sont ensuite décrites: la mission satellite Lite-BIRD, et l'observatoire au sol CMB-S4. Le chapitre se poursuit avec une description des difficultés principales rencontrées par les expériences CMB, à savoir, les effets systématiques et les avant-plans galactiques ainsi que la contamination due à l'effet de lentille gravitationnelle dans le cadre de la recherche des modes-B primordiaux. Enfin la dernière section du chapitre est consacrée à la présentation du pipeline générique d'analyse de données des expériences du CMB, avec une discussion détaillée des méthodes adoptées dans chacune des étapes de l'analyse. Cela permet d'introduire le travail de recherche de la thèse sur la reconstruction des cartes du ciel, qui est détaillé dans les trois chapitres qui suivent.

Dans le sixième chapitre, je présente d'abord une formulation générale du problème de cartographie du CMB, à partir de laquelle je dérive l'expression générale des estimateurs non-biaisés des cartes du ciel. Un bref développement s'ensuit sur la modélisation de la matrice dite de pointing. Les deux sections qui suivent discutent chacune une approche différente de cartographie du ciel: la première méthode dite de maximum de vraisemblance, repose sur la modélisation des corrélations du bruit via une matrice Toeplitz bande par bloc, et permet de pondérer le bruit dans le domaine de Fourier, en réduisant notamment les contributions dues aux basses fréquences où le bruit en $1/f$ est particulièrement élevé. La deuxième méthode dite de marginalisation des templates, permet de filtrer des signaux parasites présents dans les données en définissant des modèles linéaires de ces signaux. Un certain nombre de templates sont présentées pour modéliser différents effets systématiques: atmosphère et bruit $1/f$, signaux synchrones avec le scan du télescope, signaux synchrones avec la rotation de la lame demi-onde.

Pour clore le chapitre, une brève discussion est menée sur les diverses dégénérescences qui peuvent émerger lors de l'estimation des cartes, selon la nature des données, ainsi que leur résolution numérique.

Le septième chapitre commence par une brève introduction au calcul parallèle, qui est à la base des techniques computationnelles pour résoudre le problème de cartographie. Je donne ensuite un aperçu global de l'environnement logiciel développé dans le cadre de la thèse. Il s'ensuit une description technique détaillée des deux libraries principales qui constituent l'environnement: MIDAPACK (Microwave data analysis package), la librairie bas-niveau qui implémente les opérations matricielles nécessaires lors de la résolution du système linéaire posé par le problème de cartographie, et MAPPRASER (MIDAPACK parallel iterative sky estimator), la librairie haut-niveau qui propose les procédures numériques permettant d'estimer les cartes du ciel. Les deux méthodes de cartographie présentées au chapitre précédent sont implémentées, et peuvent utiliser deux solveurs linéaires possibles: l'algorithme du gradient conjugué préconditionné (PCG) et l'algorithme du gradient conjugué élargi (ECG). La méthode classique du PCG dispose également de divers préconditionneurs: le préconditionneur standard de Block-Jacobi, et deux variantes possibles de préconditionneurs à deux niveaux dites préconditionneur à priori, et préconditionneur à posteriori. Ces derniers s'appuient sur la construction d'un espace de déflation qui doit être idéalement généré par les vecteurs propres associés aux petites valeurs propres du système, l'opérateur sous-jacent permet alors de dé-projeter ce sous-espace du système afin d'accélérer et de stabiliser la convergence. Le premier préconditionneur repose sur une construction à priori de l'espace de déflation sans calcul préalable, exploitant uniquement certaines propriétés connues du problème. Le deuxième préconditionneur repose sur une construction à posteriori où des approximations des vecteurs propres du système sont calculées en utilisant l'algorithme de Lanczos. Dans la dernière partie du chapitre, j'effectue diverses simulations des Small Aperture Telescopes (SATs) de Simons Observatory pour valider les performances de l'implémentation ainsi que les diverses techniques numériques adoptées. Je commence par des tests d'échelle, en vérifiant plusieurs schémas de communication inter-process, pour valider la capacité du logiciel à utiliser simultanément un grand nombre de process MPI. Une démonstration s'ensuit utilisant environ 50,000 process. J'effectue également une première démonstration du solveur ECG, en obtenant une réduction du nombre d'itérations requises avant convergence par rapport à l'approche standard du PCG, tout en soulignant qu'un gain en temps d'exécution nécessitera un travail supplémentaire d'optimisation des opérations implémentées pour ce solveur. Enfin, j'évalue les performances des préconditionneurs à deux niveaux en regardant plusieurs métriques, notamment la convergence aux niveaux des différentes échelles angulaires en calculant le spectre de puissance angulaire à chaque itération, et je conclus que le préconditionneur à posteriori présente un gain significatif en temps d'exécution dans le contexte de résolution de multiples systèmes linéaires similaires, par exemple dans le cas de simulations Monte Carlo.

Le chapitre final se consacre à diverses applications, d'abord sur des simulations du satellite LiteBIRD, en étudiant quelques propriétés statistiques de la reconstruction des cartes de température en présence d'un bruit instrumental simple en $1/f$. La méthode de marginalisation des templates permet une qualité de

reconstruction optimale aux petites échelles angulaires. En revanche, aux larges échelles angulaires la qualité de la reconstruction du ciel se dégrade du fait d'un bruit statistique plus important, ce qui nécessite l'introduction d'information à priori sur l'amplitude des templates. La suite du chapitre est dédiée à des applications sur des simulations d'observation au sol du Simons Observatory SATs. Je m'intéresse d'abord à une démonstration des méthodes de mitigation des signaux synchrones avec le scan du télescope et avec la rotation de la lame demi-onde. La question de l'impact du bruit atmosphérique sur la reconstruction du signal de polarisation est ensuite abordée: la méthode de filtrage polynomial comparée à la pondération du bruit atmosphérique permet une meilleure reconstruction du signal de polarisation, lorsque les corrélations entre détecteurs ne sont pas prises en compte dans le modèle de bruit. Enfin, la différenciation entre les paires orthogonales produit les meilleurs résultats grâce à la suppression de la majeure partie du bruit atmosphérique corrélé entre détecteurs, ainsi qu'à la modulation des signaux de polarisation par la rotation de la lame-demi onde, qui permet de mitiger les effets de fuite d'intensité en polarisation (I-to-P) du bruit atmosphérique, en présence d'erreurs de calibration sur les gains de l'ordre de 0,5% à 1%.

Bibliography

- [1] CMB-S4: Data systems. <https://cmb-s4.org/experiment/data-systems/>. Accessed: 2021-08-11. (cited pp. xiii and 141)
- [2] D. Baumann, “Cosmology Lecture Notes”. <http://cosmology.amsterdam/education/cosmology/>. Accessed: 2021-04-12. (cited pp. viii, 3, 26, and 30)
- [3] J. Lesgourgues, “Master course on “Cosmology” given at the University of Bern, Spring Semester 2015”. https://lesgourg.github.io/courses/Cosmo_JL_Bern.pdf?dl=0. Accessed: 2021-03-01. (cited p. 3)
- [4] Lectures notes on Spectral Distortions of the CMB at the International School of Physics “Enrico Fermi” on Gravitational Waves and Cosmology, July 2017, Varenna, Italy. <http://www.jb.man.ac.uk/~jchluba/Science/Varenna-notes.pdf>. Accessed: 2021-07-02. (cited pp. x and 77)
- [5] LiteBIRD public website. <http://litebird.jp/eng/>. Accessed: 2021-03-27. (cited p. 16)
- [6] Lyman alpha. https://www.youtube.com/watch?v=6Bn7Ka0Tjjw&ab_channel=SSAASecretary. Accessed: 2021-07-13. (cited pp. xi and 90)
- [7] MIDAPACK 1.1b documentation. <https://b3dcmb.github.io/midapack/>. Accessed: 2021-08-11. (cited pp. xiii, 143, and 148)
- [8] PRISTINE. <https://www.ias.u-psud.fr/en/content/pristine>. Accessed: 2021-07-02. (cited p. 79)
- [9] Square Kilometer Array public website. <https://www.skatelescope.org/>. Accessed: 2021-03-27. (cited pp. 16 and 91)
- [10] Voyage 2050 Senior Committee public report. <https://www.cosmos.esa.int/documents/1866264/1866292/Voyage2050-Senior-Committee-report-public.pdf>. Accessed: 2021-07-02. (cited p. 79)
- [11] *Physics Reports*, 427(5-6):257–454, May 2006. URL: <http://dx.doi.org/10.1016/j.physrep.2005.12.006>, doi:10.1016/j.physrep.2005.12.006. (cited p. 13)

- [12] New CAST limit on the axion–photon interaction. *Nature Physics*, 13(6):584–590, May 2017. URL: <http://dx.doi.org/10.1038/nphys4109>, doi:10.1038/nphys4109. (cited p. 17)
- [13] Aad, G. et al. Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC. *Physics Letters B*, 716(1):1–29, Sep 2012. URL: <http://dx.doi.org/10.1016/j.physletb.2012.08.020>, doi:10.1016/j.physletb.2012.08.020. (cited p. 12)
- [14] Abazajian, K. et al. CMB-S4 Science Case, Reference Design, and Project Plan, 2019. [arXiv:1907.04473](https://arxiv.org/abs/1907.04473). (cited p. 110)
- [15] Abazajian, K., Fuller, G.M. and Patel, M. Sterile neutrino hot, warm, and cold dark matter. *Physical Review D*, 64(2), May 2001. URL: <http://dx.doi.org/10.1103/PhysRevD.64.023501>, doi:10.1103/physrevd.64.023501. (cited p. 16)
- [16] Abazajian, K.N. et al. CMB-S4 Science Book, First Edition, 2016. [arXiv:1610.02743](https://arxiv.org/abs/1610.02743). (cited pp. xiii, 81, 110, 118, and 119)
- [17] Abbott, B.P. et al. A Gravitational-wave Measurement of the Hubble Constant Following the Second Observing Run of Advanced LIGO and Virgo. *The Astrophysical Journal*, 909(2):218, Mar 2021. URL: <http://dx.doi.org/10.3847/1538-4357/abdcb7>, doi:10.3847/1538-4357/abdcb7. (cited pp. xi and 97)
- [18] Abbott, B. et al. Observation of Gravitational Waves from a Binary Black Hole Merger. *Physical Review Letters*, 116(6), Feb 2016. URL: <http://dx.doi.org/10.1103/PhysRevLett.116.061102>, doi:10.1103/physrevlett.116.061102. (cited p. 95)
- [19] Abbott, B. et al. Tests of general relativity with the binary black hole signals from the LIGO-Virgo catalog GWTC-1. *Physical Review D*, 100(10), Nov 2019. URL: <http://dx.doi.org/10.1103/PhysRevD.100.104036>, doi:10.1103/physrevd.100.104036. (cited pp. 18 and 95)
- [20] Abbott, B. et al. Tests of General Relativity with GW170817. *Physical Review Letters*, 123(1), Jul 2019. URL: <http://dx.doi.org/10.1103/PhysRevLett.123.011102>, doi:10.1103/physrevlett.123.011102. (cited pp. 18 and 95)
- [21] Abbott, R. et al. GWTC-2: Compact Binary Coalescences Observed by LIGO and Virgo during the First Half of the Third Observing Run. *Physical Review X*, 11(2), Jun 2021. URL: <http://dx.doi.org/10.1103/PhysRevX.11.021053>, doi:10.1103/physrevx.11.021053. (cited p. 96)
- [22] Abbott, R. et al. Tests of general relativity with binary black holes from the second LIGO-Virgo gravitational-wave transient catalog. *Physical Review D*, 103(12), Jun 2021. URL: <http://dx.doi.org/10.1103/PhysRevD.103.121001>, doi:10.1103/physrevd.103.121001. (cited p. 96)

- 103.122002, doi:10.1103/physrevd.103.122002. (cited pp. 18 and 95)
- [23] Abbott, T. et al. Dark Energy Survey year 1 results: Constraints on extended cosmological models from galaxy clustering and weak lensing. *Physical Review D*, 99(12), Jun 2019. URL: <http://dx.doi.org/10.1103/PhysRevD.99.123505>, doi:10.1103/physrevd.99.123505. (cited p. 87)
- [24] Abitbol, M.H. et al. The Simons Observatory: gain, bandpass and polarization-angle calibration requirements for B-mode searches. *Journal of Cosmology and Astroparticle Physics*, 2021(05):032, May 2021. URL: <http://dx.doi.org/10.1088/1475-7516/2021/05/032>, doi:10.1088/1475-7516/2021/05/032. (cited pp. 115 and 193)
- [25] Ackerman, L. et al. Dark matter and dark radiation. *Physical Review D*, 79(2), Jan 2009. URL: <http://dx.doi.org/10.1103/PhysRevD.79.023519>, doi:10.1103/physrevd.79.023519. (cited p. 16)
- [26] Adachi, S. et al. A Measurement of the Degree-scale CMB B-mode Angular Power Spectrum with Polarbear. *The Astrophysical Journal*, 897(1):55, Jul 2020. URL: <http://dx.doi.org/10.3847/1538-4357/ab8f24>, doi:10.3847/1538-4357/ab8f24. (cited p. 105)
- [27] Adachi, S. et al. A Measurement of the CMB E-mode Angular Power Spectrum at Subdegree Scales from 670 Square Degrees of POLARBEAR Data. *The Astrophysical Journal*, 904(1):65, Nov 2020. URL: <http://dx.doi.org/10.3847/1538-4357/abbacd>, doi:10.3847/1538-4357/abbacd. (cited p. 105)
- [28] Adachi, S. et al. Internal Delensing of Cosmic Microwave Background Polarization B-Modes with the POLARBEAR Experiment. *Physical Review Letters*, 124(13), Apr 2020. URL: <http://dx.doi.org/10.1103/PhysRevLett.124.131301>, doi:10.1103/physrevlett.124.131301. (cited pp. 105 and 119)
- [29] Adam, R. et al. Substructure and merger detection in resolved NIKA Sunyaev-Zel'dovich images of distant clusters. *Astronomy Astrophysics*, 614:A118, Jun 2018. URL: <http://dx.doi.org/10.1051/0004-6361/201731950>, doi:10.1051/0004-6361/201731950. (cited p. 79)
- [30] Ade, P.A.R. et al. Planck2013 results. XXII. Constraints on inflation. *Astronomy Astrophysics*, 571:A22, Oct 2014. URL: <http://dx.doi.org/10.1051/0004-6361/201321569>, doi:10.1051/0004-6361/201321569. (cited p. 39)
- [31] Ade, P.A.R. et al. Planck2015 results. *Astronomy Astrophysics*, 594:A19, Sep 2016. URL: <http://dx.doi.org/10.1051/0004-6361/201525821>, doi:10.1051/0004-6361/201525821. (cited p. 82)

-
- [32] Ade, P.A.R. et al. Planck2015 results. *Astronomy Astrophysics*, 594:A27, Sep 2016. URL: <http://dx.doi.org/10.1051/0004-6361/201525823>, doi:10.1051/0004-6361/201525823. (cited p. 79)
- [33] Ade, P.A.R. et al. A Measurement of the Cosmic Microwave Background B-mode Polarization Power Spectrum at Subdegree Scales from Two Years of polarbear Data. *The Astrophysical Journal*, 848(2):121, Oct 2017. URL: <http://dx.doi.org/10.3847/1538-4357/aa8e9f>, doi:10.3847/1538-4357/aa8e9f. (cited p. 105)
- [34] Ade, P.A.R. et al. BICEP2/Keck Array. IV. Optical Characterization and Performance of the BICEP2 and Keck Array Experiments. *The Astrophysical Journal*, 806(2):206, Jun 2015. URL: <http://dx.doi.org/10.1088/0004-637X/806/2/206>, doi:10.1088/0004-637X/806/2/206. (cited pp. 115 and 193)
- [35] Ade, P.A.R. et al. Planck intermediate results. XXII. Frequency dependence of thermal emission from Galactic dust in intensity and polarization. *Astronomy Astrophysics*, 576:A107, Apr 2015. URL: <http://dx.doi.org/10.1051/0004-6361/201424088>, doi:10.1051/0004-6361/201424088. (cited p. 116)
- [36] Ade, P. et al. The Simons Observatory: science goals and forecasts. *Journal of Cosmology and Astroparticle Physics*, 2019(02):056–056, Feb 2019. URL: <http://dx.doi.org/10.1088/1475-7516/2019/02/056>, doi:10.1088/1475-7516/2019/02/056. (cited pp. xii, 81, 108, and 109)
- [37] Ade, P.A. et al. POLARBEAR constraints on cosmic birefringence and primordial magnetic fields. *Physical Review D*, 92(12), Dec 2015. URL: <http://dx.doi.org/10.1103/PhysRevD.92.123509>, doi:10.1103/physrevd.92.123509. (cited pp. 81 and 105)
- [38] Ade, P. et al. BICEP2 / Keck Array IX: New bounds on anisotropies of CMB polarization rotation and implications for axionlike particles and primordial magnetic fields. *Physical Review D*, 96(10), Nov 2017. URL: <http://dx.doi.org/10.1103/PhysRevD.96.102003>, doi:10.1103/physrevd.96.102003. (cited p. 81)
- [39] Ade, P. et al. Constraints on Primordial Gravitational Waves Using Planck, WMAP, and New BICEP2/Keck Observations through the 2015 Season. *Physical Review Letters*, 121(22), Nov 2018. URL: <http://dx.doi.org/10.1103/PhysRevLett.121.221301>, doi:10.1103/physrevlett.121.221301. (cited p. 70)
- [40] Ade, P. et al. Detection of B-Mode Polarization at Degree Angular Scales by BICEP2. *Physical Review Letters*, 112(24), Jun 2014. URL: <http://dx.doi.org/10.1103/PhysRevLett.112.241101>, doi:10.1103/physrevlett.112.241101. (cited p. 121)
- [41] Ade, P. et al. Evidence for Gravitational Lensing of the Cosmic Microwave Background Polarization from Cross-Correlation with the Cosmic Infrared

- Background. *Physical Review Letters*, 112(13), Apr 2014. URL: <http://dx.doi.org/10.1103/PhysRevLett.112.131302>, doi:10.1103/physrevlett.112.131302. (cited p. 104)
- [42] Ade, P. et al. Measurement of the Cosmic Microwave Background Polarization Lensing Power Spectrum with the POLARBEAR Experiment. *Physical Review Letters*, 113(2), Jul 2014. URL: <http://dx.doi.org/10.1103/PhysRevLett.113.021301>, doi:10.1103/physrevlett.113.021301. (cited pp. 70 and 104)
- [43] Aghanim, N. et al. Planck 2018 results. *Astronomy Astrophysics*, 641:A5, Sep 2020. URL: <http://dx.doi.org/10.1051/0004-6361/201936386>, doi:10.1051/0004-6361/201936386. (cited p. 124)
- [44] Agnese, R. et al. First Dark Matter Constraints from a Super-CDMS Single-Charge Sensitive Detector. *Physical Review Letters*, 121(5), Aug 2018. URL: <http://dx.doi.org/10.1103/PhysRevLett.121.051301>, doi:10.1103/physrevlett.121.051301. (not cited)
- [45] Aguilar, A. et al. Evidence for neutrino oscillations from the observation of $\bar{\nu}_e$ appearance in a $\bar{\nu}_\mu$ beam. *Phys. Rev. D*, 64:112007, Nov 2001. URL: <https://link.aps.org/doi/10.1103/PhysRevD.64.112007>, doi:10.1103/PhysRevD.64.112007. (cited p. 18)
- [46] Aguilar-Arevalo, A. et al. Significant Excess of Electronlike Events in the MiniBooNE Short-Baseline Neutrino Experiment. *Physical Review Letters*, 121(22), Nov 2018. URL: <http://dx.doi.org/10.1103/PhysRevLett.121.221801>, doi:10.1103/physrevlett.121.221801. (cited p. 18)
- [47] Ahmad, Q.R. et al. Measurement of the Rate of $\nu_e + d \rightarrow p + p + e^-$ Interactions Produced by ^8B Solar Neutrinos at the Sudbury Neutrino Observatory. *Phys. Rev. Lett.*, 87:071301, Jul 2001. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.87.071301>, doi:10.1103/PhysRevLett.87.071301. (cited p. 13)
- [48] Ahmed, Z. et al. BICEP3: a 95GHz refracting telescope for degree-scale CMB polarization. In Holland, W.S. and Zmuidzinas, J., editors, *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy VII*, volume 9153 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 91531N, August 2014. arXiv:1407.5928, doi:10.1117/12.2057224. (cited p. 102)
- [49] Aker, M. et al. Improved Upper Limit on the Neutrino Mass from a Direct Kinematic Method by KATRIN. *Physical Review Letters*, 123(22), Nov 2019. URL: <http://dx.doi.org/10.1103/PhysRevLett.123.221802>, doi:10.1103/physrevlett.123.221802. (cited p. 15)
- [50] Aker, M. et al. First operation of the KATRIN experiment with tritium. *The European Physical Journal C*, 80(3), Mar 2020. URL: <http://dx.doi.org/10.1140/epjc/s10052-020-7718-z>, doi:10.1140/epjc/s10052-020-7718-z. (cited p. 15)

- [51] Akerib, D.S. et al. Results from a Search for Dark Matter in the Complete LUX Exposure. *Phys. Rev. Lett.*, 118:021303, Jan 2017. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.118.021303>, doi:10.1103/PhysRevLett.118.021303. (cited p. 17)
- [52] Akrami, Y. et al. Planck 2018 results. *Astronomy Astrophysics*, 641:A11, Sep 2020. URL: <http://dx.doi.org/10.1051/0004-6361/201832618>, doi:10.1051/0004-6361/201832618. (cited pp. 116 and 117)
- [53] Alam, S. et al. Completed SDSS-IV extended Baryon Oscillation Spectroscopic Survey: Cosmological implications from two decades of spectroscopic surveys at the Apache Point Observatory. *Physical Review D*, 103(8), Apr 2021. URL: <http://dx.doi.org/10.1103/PhysRevD.103.083533>, doi:10.1103/physrevd.103.083533. (cited pp. xi, xi, 86, 87, 88, and 94)
- [54] Allison, R. et al. Towards a cosmological neutrino mass detection. *Physical Review D*, 92(12), Dec 2015. URL: <http://dx.doi.org/10.1103/PhysRevD.92.123535>, doi:10.1103/physrevd.92.123535. (cited pp. 16 and 111)
- [55] Alonso, D. et al. Homogeneity and isotropy in the Two Micron All Sky Survey Photometric Redshift catalogue. *Monthly Notices of the Royal Astronomical Society*, 449(1):670–684, 03 2015. arXiv:<https://academic.oup.com/mnras/article-pdf/449/1/670/4144355/stv309.pdf>, doi:10.1093/mnras/stv309. (cited pp. viii and 4)
- [56] Alonso, D., Sanchez, J. and Slosar, A. A unified pseudo- C_ℓ framework. *Monthly Notices of the Royal Astronomical Society*, 484(3):4127–4151, January 2019. URL: <http://dx.doi.org/10.1093/mnras/stz093>, doi:10.1093/mnras/stz093. (cited p. 167)
- [57] Alpher, R.A., Bethe, H. and Gamow, G. The Origin of Chemical Elements. *Phys. Rev.*, 73:803–804, Apr 1948. URL: <https://link.aps.org/doi/10.1103/PhysRev.73.803>, doi:10.1103/PhysRev.73.803. (cited p. 2)
- [58] Alpher, R.A. and Herman, R. Evolution of the Universe. , 162(4124):774–775, November 1948. doi:10.1038/162774b0. (cited p. 2)
- [59] Amaro-Seoane, P. et al. Laser Interferometer Space Antenna, 2017. arXiv:1702.00786. (cited p. 97)
- [60] Amole, C. et al. Dark matter search results from the complete exposure of the PICO-60 C3F8 bubble chamber. *Physical Review D*, 100(2), Jul 2019. URL: <http://dx.doi.org/10.1103/PhysRevD.100.022001>, doi:10.1103/physrevd.100.022001. (cited p. 17)
- [61] André, P. et al. PRISM (Polarized Radiation Imaging and Spectroscopy Mission): an extended white paper. *Journal of Cosmology and Astroparticle Physics*, 2014(02):006–006, Feb 2014. URL: <http://dx.doi.org/10.1088/1475-7516/2014/02/006>. (cited p. 17)

- 1088/1475-7516/2014/02/006, doi:10.1088/1475-7516/2014/02/006. (cited p. 79)
- [62] Aprile, E. et al. Dark Matter Search Results from a One Ton-Year Exposure of XENON1T. *Phys. Rev. Lett.*, 121:111302, Sep 2018. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.121.111302>, doi:10.1103/PhysRevLett.121.111302. (cited p. 17)
- [63] Arkani-Hamed, N. et al. Solving the Hierarchy Problem at Reheating with a Large Number of Degrees of Freedom. *Physical Review Letters*, 117(25), Dec 2016. URL: <http://dx.doi.org/10.1103/PhysRevLett.117.251801>, doi:10.1103/physrevlett.117.251801. (cited p. 16)
- [64] Arkani-Hamed, N. et al. Ghost inflation. *Journal of Cosmology and Astroparticle Physics*, 2004(04):001–001, Apr 2004. URL: <http://dx.doi.org/10.1088/1475-7516/2004/04/001>, doi:10.1088/1475-7516/2004/04/001. (cited p. 76)
- [65] Armitage, C. and Wandelt, B. Deconvolution map-making for cosmic microwave background observations. *Physical Review D*, 70(12), Dec 2004. URL: <http://dx.doi.org/10.1103/PhysRevD.70.123007>, doi:10.1103/physrevd.70.123007. (cited p. 133)
- [66] Arvanitaki, A. et al. String axiverse. *Physical Review D*, 81(12), Jun 2010. URL: <http://dx.doi.org/10.1103/PhysRevD.81.123530>, doi:10.1103/physrevd.81.123530. (cited p. 17)
- [67] Arzoumanian, Z. et al. The NANOGrav 12.5 yr Data Set: Search for an Isotropic Stochastic Gravitational-wave Background. *The Astrophysical Journal Letters*, 905(2):L34, Dec 2020. URL: <http://dx.doi.org/10.3847/2041-8213/abd401>, doi:10.3847/2041-8213/abd401. (cited p. 96)
- [68] Audren, B. et al. Neutrino masses and cosmological parameters from a Euclid-like survey: Markov Chain Monte Carlo forecasts including theoretical errors. *Journal of Cosmology and Astroparticle Physics*, 2013(01):026–026, Jan 2013. URL: <http://dx.doi.org/10.1088/1475-7516/2013/01/026>, doi:10.1088/1475-7516/2013/01/026. (cited p. 16)
- [69] Aumont, J. et al. Measurement of the Crab nebula polarization at 90 GHz as a calibrator for CMB experiments. *Astronomy and Astrophysics*, 514:A70, May 2010. URL: <http://dx.doi.org/10.1051/0004-6361/200913834>, doi:10.1051/0004-6361/200913834. (cited p. 119)
- [70] Austermann, J.E. et al. SPTpol: an instrument for CMB polarization measurements with the South Pole Telescope. *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy VI*, Sep 2012. URL: <http://dx.doi.org/10.1117/12.927286>, doi:10.1117/12.927286. (cited p. 102)
- [71] Aybas, D. et al. Search for axion-like dark matter using solid-state nuclear magnetic resonance, 2021. *arXiv:2101.01241*. (cited p. 17)

- [72] Baccigalupi, C. et al. Neural networks and the separation of cosmic microwave background and astrophysical signals in sky maps. *Monthly Notices of the Royal Astronomical Society*, 318(3):769–780, Nov 2000. URL: <http://dx.doi.org/10.1046/j.1365-8711.2000.03751.x>, doi: 10.1046/j.1365-8711.2000.03751.x. (cited p. 122)
- [73] Baluni, V. CP Violating Effects in QCD. *Phys. Rev. D*, 19:2227–2230, 1979. doi:10.1103/PhysRevD.19.2227. (cited p. 17)
- [74] Bardeen, J.M. et al. The Statistics of Peaks of Gaussian Random Fields. , 304:15, May 1986. doi:10.1086/164143. (cited p. 84)
- [75] Bardeen, J.M. Gauge-invariant cosmological perturbations. *Phys. Rev. D*, 22:1882–1905, Oct 1980. URL: <https://link.aps.org/doi/10.1103/PhysRevD.22.1882>, doi:10.1103/PhysRevD.22.1882. (cited p. 40)
- [76] Bardeen, J.M., Steinhardt, P.J. and Turner, M.S. Spontaneous creation of almost scale-free density perturbations in an inflationary universe. , 28(4):679–693, August 1983. doi:10.1103/PhysRevD.28.679. (cited p. 34)
- [77] Barkana, R. Possible interaction between baryons and dark-matter particles revealed by the first stars. *Nature*, 555(7694):71–74, Mar 2018. URL: <http://dx.doi.org/10.1038/nature25791>, doi:10.1038/nature25791. (cited p. 92)
- [78] Bartolo, N. et al. Primordial Black Hole Dark Matter: LISA Serendipity. *Phys. Rev. Lett.*, 122:211301, May 2019. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.122.211301>, doi:10.1103/PhysRevLett.122.211301. (cited p. 18)
- [79] Bartolo, N. et al. Non-Gaussianity from inflation: theory and observations. *Physics Reports*, 402(3-4):103–266, Nov 2004. URL: <http://dx.doi.org/10.1016/j.physrep.2004.08.022>, doi:10.1016/j.physrep.2004.08.022. (cited p. 75)
- [80] Bartolo, N., Matarrese, S. and Riotto, A. Non-Gaussianity from inflation. *Physical Review D*, 65(10), Apr 2002. URL: <http://dx.doi.org/10.1103/PhysRevD.65.103505>, doi:10.1103/physrevd.65.103505. (cited p. 75)
- [81] Bashinsky, S. and Seljak, U. Signatures of relativistic neutrinos in CMB anisotropy and matter clustering. *Physical Review D*, 69(8), Apr 2004. URL: <http://dx.doi.org/10.1103/PhysRevD.69.083002>, doi:10.1103/physrevd.69.083002. (cited p. 72)
- [82] Battistelli, E.S. et al. Intensity and polarization of the atmospheric emission at millimetric wavelengths at Dome Concordia. *Monthly Notices of the Royal Astronomical Society*, 423(2):1293–1299, May 2012. URL: <http://dx.doi.org/10.1111/j.1365-2966.2012.20951.x>, doi:10.1111/j.1365-2966.2012.20951.x. (cited p. 113)

- [83] Baumann, D. TASI Lectures on Inflation, 2012. [arXiv:0907.5424](#). (cited pp. ix, ix, ix, ix, ix, 32, 33, 35, 36, 37, 40, 47, and 48)
- [84] Baumann, D. et al. Phases of new physics in the CMB. *Journal of Cosmology and Astroparticle Physics*, 2016(01):007–007, Jan 2016. URL: <http://dx.doi.org/10.1088/1475-7516/2016/01/007>, doi:10.1088/1475-7516/2016/01/007. (cited p. 72)
- [85] Baumann, D., Green, D. and Wallisch, B. New Target for Cosmic Axion Searches. *Physical Review Letters*, 117(17), Oct 2016. URL: <http://dx.doi.org/10.1103/PhysRevLett.117.171301>, doi:10.1103/physrevlett.117.171301. (cited p. 16)
- [86] Bautista, J.E. et al. The completed SDSS-IV extended Baryon Oscillation Spectroscopic Survey: measurement of the BAO and growth rate of structure of the luminous red galaxy sample from the anisotropic correlation function between redshifts 0.6 and 1. , 500(1):736–762, January 2021. [arXiv:2007.08993](#), doi:10.1093/mnras/staa2800. (cited pp. x and 85)
- [87] Beck, D. *Challenges in CMB Lensing Data Analysis and Scientific Exploitation of Current and Future CMB Polarization Experiments*. PhD thesis, 2019. Thèse de doctorat dirigée par Stompor, Radek Physique de l’Univers.Cosmologie Université de Paris (2019-....) 2019. URL: <http://www.theses.fr/2019UNIP7022>. (cited pp. xii, xii, 99, and 103)
- [88] Becker, R.H. et al. Evidence for Reionization at $z \sim 6$: Detection of a Gunn-Peterson Trough in a $z=6.28$ Quasar. , 122(6):2850–2857, December 2001. [arXiv:astro-ph/0108097](#), doi:10.1086/324231. (cited p. 30)
- [89] Benedict, G.F. et al. Hubble Space Telescope Fine Guidance Sensor Parallaxes of Galactic Cepheid Variable Stars: Period-Luminosity Relations. , 133(4):1810–1827, April 2007. [arXiv:astro-ph/0612465](#), doi:10.1086/511980. (cited p. 94)
- [90] Bennett, C.L. et al. First-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Foreground Emission. , 148(1):97–117, September 2003. [arXiv:astro-ph/0302208](#), doi:10.1086/377252. (cited p. 122)
- [91] Benoît, A. et al. The cosmic microwave background anisotropy power spectrum measured by Archeops. *Astronomy Astrophysics*, 399(3):L19–L23, Feb 2003. URL: <http://dx.doi.org/10.1051/0004-6361:20021850>, doi:10.1051/0004-6361:20021850. (cited p. 100)
- [92] Bernardeau, F. and Uzan, J.P. Non-Gaussianity in multifield inflation. *Physical Review D*, 66(10), Nov 2002. URL: <http://dx.doi.org/10.1103/PhysRevD.66.103506>, doi:10.1103/physrevd.66.103506. (cited p. 75)

-
- [93] Bernus, L. et al. Constraining the Mass of the Graviton with the Planetary Ephemeris INPOP. *Phys. Rev. Lett.*, 123:161103, Oct 2019. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.123.161103>, doi:10.1103/PhysRevLett.123.161103. (cited p. 18)
 - [94] Betoule, M. et al. Improved cosmological constraints from a joint analysis of the SDSS-II and SNLS supernova samples. *Astronomy Astrophysics*, 568:A22, Aug 2014. URL: <http://dx.doi.org/10.1051/0004-6361/201423413>, doi:10.1051/0004-6361/201423413. (cited pp. xi and 93)
 - [95] Betti, M. et al. Neutrino physics with the PTOLEMY project: active neutrino properties and the light sterile case. *Journal of Cosmology and Astroparticle Physics*, 2019(07):047–047, Jul 2019. URL: <http://dx.doi.org/10.1088/1475-7516/2019/07/047>, doi:10.1088/1475-7516/2019/07/047. (cited p. 25)
 - [96] Bezrukov, F. and Shaposhnikov, M. The Standard Model Higgs boson as the inflaton. *Physics Letters B*, 659(3):703–706, Jan 2008. URL: <http://dx.doi.org/10.1016/j.physletb.2007.11.072>, doi:10.1016/j.physletb.2007.11.072. (cited p. 36)
 - [97] Bi, H. and Davidsen, A.F. Evolution of Structure in the Intergalactic Medium and the Nature of the Ly α Forest. *The Astrophysical Journal*, 479(2):523–542, Apr 1997. URL: <http://dx.doi.org/10.1086/303908>, doi:10.1086/303908. (cited p. 90)
 - [98] Birkinshaw, M., Gull, S.F. and Hardebeck, H. The Sunyaev-Zeldovich effect towards three clusters of galaxies. , 309:34–35, May 1984. doi:10.1038/309034a0. (cited p. 79)
 - [99] Blanton, M.R. et al. Sloan Digital Sky Survey IV: Mapping the Milky Way, Nearby Galaxies, and the Distant Universe. , 154(1):28, July 2017. arXiv: 1703.00052, doi:10.3847/1538-3881/aa7567. (cited p. 84)
 - [100] Bleem, L.E. et al. The SPTpol Extended Cluster Survey. *The Astrophysical Journal Supplement Series*, 247(1):25, Mar 2020. URL: <http://dx.doi.org/10.3847/1538-4365/ab6993>, doi:10.3847/1538-4365/ab6993. (cited p. 79)
 - [101] Bond, J.R., Jaffe, A.H. and Knox, L. Estimating the power spectrum of the cosmic microwave background. *Physical Review D*, 57(4):2117–2137, Feb 1998. URL: <http://dx.doi.org/10.1103/PhysRevD.57.2117>, doi:10.1103/physrevd.57.2117. (cited p. 123)
 - [102] Bondi, H. and Gold, T. The Steady-State Theory of the Expanding Universe. , 108:252, January 1948. doi:10.1093/mnras/108.3.252. (cited p. 51)
 - [103] Borrill, J. Power spectrum estimators for large CMB datasets. *Physical Review D*, 59(2), Dec 1998. URL: <http://dx.doi.org/10.1103/PhysRevD.59.027302>, doi:10.1103/physrevd.59.027302. (cited p. 123)

- [104] Bowman, J.D. et al. An absorption profile centred at 78 megahertz in the sky-averaged spectrum. *Nature*, 555(7694):67–70, Mar 2018. URL: <http://dx.doi.org/10.1038/nature25792>, doi:10.1038/nature25792. (cited pp. xi, 30, and 91)
- [105] Boyarsky, A., Ruchayskiy, O. and Shaposhnikov, M. The Role of Sterile Neutrinos in Cosmology and Astrophysics. *Annual Review of Nuclear and Particle Science*, 59(1):191–214, Nov 2009. URL: <http://dx.doi.org/10.1146/annurev.nucl.010909.083654>, doi:10.1146/annurev.nucl.010909.083654. (cited p. 16)
- [106] Braine, T. et al. Extended Search for the Invisible Axion with the Axion Dark Matter Experiment. *Physical Review Letters*, 124(10), Mar 2020. URL: <http://dx.doi.org/10.1103/PhysRevLett.124.101303>, doi:10.1103/physrevlett.124.101303. (cited p. 17)
- [107] Brandt, W.N. et al. Separation of Foreground Radiation from Cosmic Microwave Background Anisotropy Using Multifrequency Measurements. , 424:1, March 1994. doi:10.1086/173867. (cited p. 122)
- [108] Brinckmann, T. et al. The promising future of a robust cosmological neutrino mass measurement. *Journal of Cosmology and Astroparticle Physics*, 2019(01):059–059, Jan 2019. URL: <http://dx.doi.org/10.1088/1475-7516/2019/01/059>, doi:10.1088/1475-7516/2019/01/059. (cited p. 16)
- [109] Brown, L.M. The idea of the neutrino. *Phys. Today*, 31N9:23–28, 1978. doi:10.1063/1.2995181. (cited p. 13)
- [110] Brown, M.L., Castro, P.G. and Taylor, A.N. Cosmic microwave background temperature and polarization pseudo-Estimators and covariances. *Monthly Notices of the Royal Astronomical Society*, 360(4):1262–1280, Jul 2005. URL: <http://dx.doi.org/10.1111/j.1365-2966.2005.09111.x>, doi:10.1111/j.1365-2966.2005.09111.x. (cited p. 124)
- [111] Bucher, M., Goldhaber, A.S. and Turok, N. Open universe from inflation. *Physical Review D*, 52(6):3314–3337, Sep 1995. URL: <http://dx.doi.org/10.1103/PhysRevD.52.3314>, doi:10.1103/physrevd.52.3314. (cited p. 46)
- [112] Bunn, E.F., Hoffman, Y. and Silk, J. The Wiener-filtered COBE DMR Data and Predictions for the Tenerife Experiment. *The Astrophysical Journal*, 464:1, June 1996. arXiv:astro-ph/9509045, doi:10.1086/177294. (cited p. 130)
- [113] Burigana, C. et al. A preliminary study on destriping techniques of PLANCK/LFI measurements versus observational strategy. *arXiv e-prints*, pages astro-ph/9906360, June 1999. arXiv:astro-ph/9906360. (cited pp. 130 and 134)

- [114] Burke-Spolaor, S. et al. The astrophysics of nanohertz gravitational waves. *The Astronomy and Astrophysics Review*, 27(1), Jun 2019. URL: <http://dx.doi.org/10.1007/s00159-019-0115-7>, doi:10.1007/s00159-019-0115-7. (cited p. 96)
- [115] Burles, S., Nollett, K.M. and Turner, M.S. Big-Bang Nucleosynthesis: Linking Inner Space and Outer Space. *arXiv e-prints*, pages astro-ph/9903300, March 1999. arXiv:astro-ph/9903300. (cited pp. viii and 28)
- [116] Calabrese, E., Alonso, D. and Dunkley, J. Complementing the ground-based CMB-S4 experiment on large scales with the PIXIE satellite. *Physical Review D*, 95(6), Mar 2017. URL: <http://dx.doi.org/10.1103/PhysRevD.95.063504>, doi:10.1103/physrevd.95.063504. (cited p. 111)
- [117] Calabrese, E. et al. Precision epoch of reionization studies with next-generation CMB experiments. *Journal of Cosmology and Astroparticle Physics*, 2014(08):010–010, Aug 2014. URL: <http://dx.doi.org/10.1088/1475-7516/2014/08/010>, doi:10.1088/1475-7516/2014/08/010. (cited p. 73)
- [118] Calabrese, M. et al. Multiple lensing of the cosmic microwave background anisotropies. *Journal of Cosmology and Astroparticle Physics*, 2015(03):049–049, mar 2015. doi:10.1088/1475-7516/2015/03/049. (cited p. 67)
- [119] Calderón, R. et al. Negative cosmological constant in the dark sector? *Physical Review D*, 103(2), Jan 2021. URL: <http://dx.doi.org/10.1103/PhysRevD.103.023526>, doi:10.1103/physrevd.103.023526. (cited p. 31)
- [120] Caldwell, R.R., Dave, R. and Steinhardt, P.J. Cosmological Imprint of an Energy Component with General Equation of State. , 80(8):1582–1585, February 1998. arXiv:astro-ph/9708069, doi:10.1103/PhysRevLett.80.1582. (cited p. 18)
- [121] Caldwell, R.R., Gluscevic, V. and Kamionkowski, M. Cross-correlation of cosmological birefringence with CMB temperature. *Physical Review D*, 84(4), Aug 2011. URL: <http://dx.doi.org/10.1103/PhysRevD.84.043504>, doi:10.1103/physrevd.84.043504. (cited p. 81)
- [122] Caldwell, R. A phantom menace? Cosmological consequences of a dark energy component with super-negative equation of state. *Physics Letters B*, 545(1-2):23–29, Oct 2002. URL: [http://dx.doi.org/10.1016/S0370-2693\(02\)02589-3](http://dx.doi.org/10.1016/S0370-2693(02)02589-3), doi:10.1016/S0370-2693(02)02589-3. (cited p. 31)
- [123] Cantalupo, C.M. et al. MADmap: A MASSIVELY PARALLEL MAXIMUM LIKELIHOOD COSMIC MICROWAVE BACKGROUND MAP-MAKER. *The Astrophysical Journal Supplement Series*, 187(1):212–227, Mar 2010. URL: <http://dx.doi.org/10.1088/0067-0049/187/1/212>, doi:10.1088/0067-0049/187/1/212. (cited pp. 152 and 153)

- [124] Carbone, C. et al. Lensed CMB temperature and polarization maps from the Millennium Simulation. *Monthly Notices of the Royal Astronomical Society*, 396(2):668–679, 06 2009. arXiv:<https://academic.oup.com/mnras/article-pdf/396/2/668/3383256/mnras0396-0668.pdf>, doi:10.1111/j.1365-2966.2009.14746.x. (cited p. 67)
- [125] Cardoso, J.F. et al. Component Separation With Flexible Models—Application to Multichannel Astrophysical Observations. *IEEE Journal of Selected Topics in Signal Processing*, 2(5):735–746, November 2008. doi:10.1109/JSTSP.2008.2005346. (cited p. 122)
- [126] Carr, B. and Kühnel, F. Primordial Black Holes as Dark Matter: Recent Developments. *Annual Review of Nuclear and Particle Science*, 70(1):355–394, Oct 2020. URL: <http://dx.doi.org/10.1146/annurev-nucl-050520-125911>, doi:10.1146/annurev-nucl-050520-125911. (cited p. 18)
- [127] Carroll, S.M. Quintessence and the Rest of the World: Suppressing Long-Range Interactions. *Physical Review Letters*, 81(15):3067–3070, Oct 1998. URL: <http://dx.doi.org/10.1103/PhysRevLett.81.3067>, doi:10.1103/physrevlett.81.3067. (cited p. 81)
- [128] Carron, J. Optimal constraints on primordial gravitational waves from the lensed CMB. *Physical Review D*, 99(4), Feb 2019. URL: <http://dx.doi.org/10.1103/PhysRevD.99.043518>, doi:10.1103/physrevd.99.043518. (cited p. 105)
- [129] Carron, J. and Lewis, A. Maximum a posteriori CMB lensing reconstruction. *Physical Review D*, 96(6), Sep 2017. URL: <http://dx.doi.org/10.1103/PhysRevD.96.063510>, doi:10.1103/physrevd.96.063510. (cited pp. 105, 118, and 119)
- [130] Carter, F.W. et al. Year two instrument status of the SPT-3G cosmic microwave background receiver. *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy IX*, Aug 2018. URL: <http://dx.doi.org/10.1117/12.2312426>, doi:10.1117/12.2312426. (cited p. 102)
- [131] Cen, R. and Ostriker, J.P. Where Are the Baryons? *The Astrophysical Journal*, 514(1):1–6, mar 1999. doi:10.1086/306949. (cited p. 13)
- [132] Chabanier, S. *Neutrinos and dark matter cosmology with the Lyman- α forest : the interplay between large-scale evolution and small-scale baryonic physics*. Theses, Université Paris-Saclay, October 2020. URL: <https://tel.archives-ouvertes.fr/tel-03123366>. (cited pp. xi and 90)
- [133] Challinor, A. and Chon, G. Error analysis of quadratic power spectrum estimates for cosmic microwave background polarization: sampling covariance. *Monthly Notices of the Royal Astronomical Society*, 360(2):509–532, Jun 2005. URL: <http://dx.doi.org/10.1111/j.1365-2966.2005.09076.x>, doi:10.1111/j.1365-2966.2005.09076.x. (cited p. 124)

-
- [134] Challinor, A. and Lewis, A. Lensed CMB power spectra from all-sky correlation functions. *Physical Review D*, 71(10), May 2005. URL: <http://dx.doi.org/10.1103/PhysRevD.71.103010>, doi: 10.1103/physrevd.71.103010. (cited p. 69)
- [135] Chang, C.L. et al. SPT-SZ: a Sunyaev-ZePdvovich survey for galaxy clusters. In Young, B., Cabrera, B. and Miller, A., editors, *The Thirteenth International Workshop on Low Temperature Detectors - LTD13*, volume 1185 of *American Institute of Physics Conference Series*, pages 475–477, December 2009. doi: 10.1063/1.3292381. (cited p. 102)
- [136] Chatrchyan, S. et al. Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC. *Physics Letters B*, 716(1):30–61, Sep 2012. URL: <http://dx.doi.org/10.1016/j.physletb.2012.08.021>, doi:10.1016/j.physletb.2012.08.021. (cited p. 12)
- [137] Chen, H.Y., Fishbach, M. and Holz, D.E. A two per cent Hubble constant measurement from standard sirens within five years. *Nature*, 562(7728):545–547, Oct 2018. URL: <http://dx.doi.org/10.1038/s41586-018-0606-0>, doi:10.1038/s41586-018-0606-0. (cited p. 97)
- [138] Chen, X. et al. Observational signatures and non-Gaussianities of general single-field inflation. *Journal of Cosmology and Astroparticle Physics*, 2007(01):002–002, Jan 2007. URL: <http://dx.doi.org/10.1088/1475-7516/2007/01/002>, doi:10.1088/1475-7516/2007/01/002. (cited p. 76)
- [139] Chen, Z.C. and Huang, Q.G. Distinguishing primordial black holes from astrophysical black holes by Einstein Telescope and Cosmic Explorer. *Journal of Cosmology and Astroparticle Physics*, 2020(08):039–039, Aug 2020. URL: <http://dx.doi.org/10.1088/1475-7516/2020/08/039>, doi:10.1088/1475-7516/2020/08/039. (cited p. 18)
- [140] Chen, Z.C., Yuan, C. and Huang, Q.G. Pulsar Timing Array Constraints on Primordial Black Holes with NANOGrav 11-Year Dataset. *Physical Review Letters*, 124(25), Jun 2020. URL: <http://dx.doi.org/10.1103/PhysRevLett.124.251101>, doi:10.1103/physrevlett.124.251101. (cited p. 18)
- [141] Chluba, J. et al. New Horizons in Cosmology with Spectral Distortions of the Cosmic Microwave Background, 2019. arXiv:1909.01593. (cited p. 78)
- [142] Chluba, J. and Sunyaev, R.A. The evolution of CMB spectral distortions in the early Universe. *Monthly Notices of the Royal Astronomical Society*, 419(2):1294–1314, Nov 2011. URL: <http://dx.doi.org/10.1111/j.1365-2966.2011.19786.x>, doi:10.1111/j.1365-2966.2011.19786.x. (cited p. 77)

- [143] Choi, S.K. et al. The Atacama Cosmology Telescope: a measurement of the Cosmic Microwave Background power spectra at 98 and 150 GHz. *Journal of Cosmology and Astroparticle Physics*, 2020(12):045–045, Dec 2020. URL: <http://dx.doi.org/10.1088/1475-7516/2020/12/045>, doi:10.1088/1475-7516/2020/12/045. (cited pp. x and 71)
- [144] Christensen, N. et al. Bayesian methods for cosmological parameter estimation from cosmic microwave background measurements. *Classical and Quantum Gravity*, 18(14):2677–2688, jun 2001. doi:10.1088/0264-9381/18/14/306. (cited p. 125)
- [145] Collaboration, I. Searches for Sterile Neutrinos with the IceCube Detector. *Phys. Rev. Lett.*, 117:071801, Aug 2016. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.117.071801>, doi:10.1103/PhysRevLett.117.071801. (cited p. 18)
- [146] Conley, A. et al. SiFTO: An Empirical Method for Fitting SN Ia Light Curves. *The Astrophysical Journal*, 681(1):482–498, Jul 2008. URL: <http://dx.doi.org/10.1086/588518>, doi:10.1086/588518. (cited p. 93)
- [147] Creminelli, P. and Senatore, L. A smooth bouncing cosmology with scale invariant spectrum. *Journal of Cosmology and Astroparticle Physics*, 2007(11):010–010, Nov 2007. URL: <http://dx.doi.org/10.1088/1475-7516/2007/11/010>, doi:10.1088/1475-7516/2007/11/010. (cited p. 75)
- [148] Crewther, R.J. et al. Chiral Estimate of the Electric Dipole Moment of the Neutron in Quantum Chromodynamics. *Phys. Lett. B*, 88:123, 1979. [Erratum: *Phys.Lett.B* 91, 487 (1980)]. doi:10.1016/0370-2693(79)90128-X. (cited p. 17)
- [149] Cui, X. et al. Dark Matter Results from 54-Ton-Day Exposure of PandaX-II Experiment. *Phys. Rev. Lett.*, 119:181302, Oct 2017. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.119.181302>, doi:10.1103/PhysRevLett.119.181302. (cited p. 17)
- [150] Culverhouse, T. et al. THE QUaD GALACTIC PLANE SURVEY. I. MAPS AND ANALYSIS OF DIFFUSE EMISSION. *The Astrophysical Journal*, 722(2):1057–1077, Sep 2010. URL: <http://dx.doi.org/10.1088/0004-637X/722/2/1057>, doi:10.1088/0004-637x/722/2/1057. (cited p. 121)
- [151] Cyburt, R.H., Fields, B.D. and Olive, K.A. An update on the big bang nucleosynthesis prediction for ${}^7\text{Li}$: the problem worsens. *Journal of Cosmology and Astroparticle Physics*, 2008(11):012, Nov 2008. URL: <http://dx.doi.org/10.1088/1475-7516/2008/11/012>, doi:10.1088/1475-7516/2008/11/012. (cited p. 27)
- [152] Cyr-Racine, F.Y. and Sigurdson, K. Cosmology of atomic dark matter. *Physical Review D*, 87(10), May 2013. URL: <http://dx.doi.org/10.1103/PhysRevD.87.103515>, doi:10.1103/physrevd.87.103515. (cited p. 16)

- [153] Das, S. and Bode, P. A Large Sky Simulation of the Gravitational Lensing of the Cosmic Microwave Background. , 682(1):1–13, July 2008. [arXiv:0711.3793](#), [doi:10.1086/589638](#). (cited p. 67)
- [154] Das, S. et al. Detection of the Power Spectrum of Cosmic Microwave Background Lensing by the Atacama Cosmology Telescope. *Physical Review Letters*, 107(2), Jul 2011. URL: <http://dx.doi.org/10.1103/PhysRevLett.107.021301>, [doi:10.1103/physrevlett.107.021301](#). (cited p. 101)
- [155] Dave, R. et al. Baryons in the Warm-Hot Intergalactic Medium. *The Astrophysical Journal*, 552(2):473–483, May 2001. URL: <http://dx.doi.org/10.1086/320548>, [doi:10.1086/320548](#). (cited p. 13)
- [156] de Bernardis, P. et al. Multiple Peaks in the Angular Power Spectrum of the Cosmic Microwave Background: Significance and Consequences for Cosmology. *The Astrophysical Journal*, 564(2):559–566, Jan 2002. URL: <http://dx.doi.org/10.1086/324298>, [doi:10.1086/324298](#). (cited p. 100)
- [157] de Bernardis, P. et al. A flat Universe from high-resolution maps of the cosmic microwave background radiation. *Nature*, 404(6781):955–959, Apr 2000. URL: <http://dx.doi.org/10.1038/35010035>, [doi:10.1038/35010035](#). (cited p. 100)
- [158] de Graaff, A. et al. Probing the missing baryons with the Sunyaev-Zel’dovich effect from filaments. *Astronomy Astrophysics*, 624:A48, Apr 2019. URL: <http://dx.doi.org/10.1051/0004-6361/201935159>, [doi:10.1051/0004-6361/201935159](#). (cited p. 13)
- [159] de Salas, P.F. and Pastor, S. Relic neutrino decoupling with flavour oscillations revisited. *Journal of Cosmology and Astroparticle Physics*, 2016(07):051–051, Jul 2016. URL: <http://dx.doi.org/10.1088/1475-7516/2016/07/051>, [doi:10.1088/1475-7516/2016/07/051](#). (cited p. 16)
- [160] de Sitter, W. On Einstein’s Theory of Gravitation and its Astronomical Consequences. Third Paper.. *Monthly Notices of the Royal Astronomical Society*, 78(1):3–28, 11 1917. [arXiv:https://academic.oup.com/mnras/article-pdf/78/1/3/4067045/mnras78-0003.pdf](#), [doi:10.1093/mnras/78.1.3](#). (cited p. 2)
- [161] Delabrouille, J. et al. A full sky, low foreground, high resolution CMB map from WMAP. , 493(3):835–857, January 2009. [arXiv:0807.0773](#), [doi:10.1051/0004-6361:200810514](#). (cited p. 122)
- [162] Delabrouille, J. Analysis of the accuracy of a destripping method for future cosmic microwave background mapping with the PLANCK SURVEYOR satellite. *Astron. Astrophys. Suppl. Ser.*, 127(3):555–567, 1998. [doi:10.1051/aas:1998119](#). (cited pp. 130 and 134)

- [163] Desjacques, V., Jeong, D. and Schmidt, F. Large-scale galaxy bias. *Physics Reports*, 733:1–193, Feb 2018. URL: <http://dx.doi.org/10.1016/j.physrep.2017.12.002>, doi:10.1016/j.physrep.2017.12.002. (cited p. 85)
- [164] Di Valentino, E. et al. In the realm of the Hubble tension—a review of solutions *. *Classical and Quantum Gravity*, 38(15):153001, Jul 2021. URL: <http://dx.doi.org/10.1088/1361-6382/ac086d>, doi:10.1088/1361-6382/ac086d. (cited p. 3)
- [165] Dodelson, S. and Schmidt, F. *Modern Cosmology*. Elsevier Science, 2020. URL: <https://books.google.fr/books?id=GGjfywEACAAJ>. (cited pp. ix, 43, and 57)
- [166] Dodelson, S. Coherent Phase Argument for Inflation. *AIP Conference Proceedings*, 2003. URL: <http://dx.doi.org/10.1063/1.1627736>, doi:10.1063/1.1627736. (cited pp. ix, ix, 47, and 48)
- [167] Dodelson, S. and Widrow, L.M. Sterile neutrinos as dark matter. *Physical Review Letters*, 72(1):17–20, Jan 1994. URL: <http://dx.doi.org/10.1103/PhysRevLett.72.17>, doi:10.1103/physrevlett.72.17. (cited p. 14)
- [168] Dolinski, M.J., Poon, A.W. and Rodejohann, W. Neutrinoless Double-Beta Decay: Status and Prospects. *Annual Review of Nuclear and Particle Science*, 69(1):219–251, 2019. arXiv:<https://doi.org/10.1146/annurev-nucl-101918-023407>, doi:10.1146/annurev-nucl-101918-023407. (cited p. 14)
- [169] du Mas des Bourboux, H. et al. The Completed SDSS-IV Extended Baryon Oscillation Spectroscopic Survey: Baryon Acoustic Oscillations with Ly α Forests. *The Astrophysical Journal*, 901(2):153, Oct 2020. URL: <http://dx.doi.org/10.3847/1538-4357/abb085>, doi:10.3847/1538-4357/abb085. (cited p. 90)
- [170] Duivenvoorden, A.J. et al. Probing frequency-dependent half-wave plate systematics for CMB experiments with full-sky beam convolution simulations. *Monthly Notices of the Royal Astronomical Society*, 502(3):4526–4539, Feb 2021. URL: <http://dx.doi.org/10.1093/mnras/stab317>, doi:10.1093/mnras/stab317. (cited p. 196)
- [171] Dvali, G., Gruzinov, A. and Zaldarriaga, M. New mechanism for generating density perturbations from inflation. *Physical Review D*, 69(2), Jan 2004. URL: <http://dx.doi.org/10.1103/PhysRevD.69.023505>, doi:10.1103/physrevd.69.023505. (cited p. 75)
- [172] eBOSS Collaboration et al. The Completed SDSS-IV extended Baryon Oscillation Spectroscopic Survey: Cosmological Implications from two Decades of Spectroscopic Surveys at the Apache Point observatory. *arXiv e-prints*, page arXiv:2007.08991, July 2020. arXiv:2007.08991. (cited p. 3)

- [173] Eddington, A.S. On the Instability of Einstein's Spherical World. *Monthly Notices of the Royal Astronomical Society*, 90(7):668–678, 05 1930. arXiv:<https://academic.oup.com/mnras/article-pdf/90/7/668/2901975/mnras90-0668.pdf>, doi:10.1093/mnras/90.7.668. (cited p. 2)
- [174] Efstathiou, G. A maximum likelihood analysis of the low cosmic microwave background multipoles from the Wilkinson Microwave Anisotropy Probe. *Monthly Notices of the Royal Astronomical Society*, 348(3):885–896, 03 2004. arXiv:<https://academic.oup.com/mnras/article-pdf/348/3/885/4105590/348-3-885.pdf>, doi:10.1111/j.1365-2966.2004.07409.x. (cited p. 123)
- [175] Ehret, K. et al. New ALPS results on hidden-sector lightweights. *Physics Letters B*, 689(4-5):149–155, May 2010. URL: <http://dx.doi.org/10.1016/j.physletb.2010.04.066>, doi:10.1016/j.physletb.2010.04.066. (cited p. 17)
- [176] Einstein, A. Näherungsweise Integration der Feldgleichungen der Gravitation. *Sitzungsberichte der Königlich Preussischen Akademie der Wissenschaften (Berlin)*, pages 688–696, January 1916. (cited p. 95)
- [177] Einstein, A. Kosmologische Betrachtungen zur allgemeinen Relativitätstheorie. *Sitzungsberichte der Königlich Preussischen Akademie der Wissenschaften (Berlin)*, pages 142–152, January 1917. (cited pp. 2 and 18)
- [178] Einstein, A. Über Gravitationswellen. *Sitzungsberichte der Königlich Preussischen Akademie der Wissenschaften (Berlin)*, pages 154–167, January 1918. (cited p. 95)
- [179] Eisenstein, D.J. et al. Detection of the Baryon Acoustic Peak in the Large-Scale Correlation Function of SDSS Luminous Red Galaxies. *The Astrophysical Journal*, 633(2):560–574, Nov 2005. URL: <http://dx.doi.org/10.1086/466512>, doi:10.1086/466512. (cited p. 85)
- [180] El Bouhargani, H. et al. MAPPRaiser: A massively parallel map-making framework for multi-kilo pixel CMB experiments. 2021, in prep. (cited p. 128)
- [181] Eriksen, H.K. et al. Cosmic Microwave Background Component Separation by Parameter Estimation. *The Astrophysical Journal*, 641(2):665–682, Apr 2006. URL: <http://dx.doi.org/10.1086/500499>, doi:10.1086/500499. (cited p. 122)
- [182] Eriksen, H.K. et al. Power Spectrum Estimation from High-Resolution Maps by Gibbs Sampling. , 155(2):227–241, December 2004. arXiv:astro-ph/0407028, doi:10.1086/425219. (cited p. 124)
- [183] Errard, J. et al. MODELING ATMOSPHERIC EMISSION FOR CMB GROUND-BASED OBSERVATIONS. *The Astrophysical Journal*, 809(1):63, Aug 2015. URL: <http://dx.doi.org/10.1088/0004-637X/809/1/63>, doi:10.1088/0004-637X/809/1/63. (cited pp. 113 and 161)

- [184] Essinger-Hileman, T. et al. CLASS: the cosmology large angular scale surveyor. *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy VII*, Jul 2014. URL: <http://dx.doi.org/10.1117/12.2056701>, doi:10.1117/12.2056701. (cited p. 107)
- [185] et al., D.C. The DESI Experiment Part I: Science, Targeting, and Survey Design, 2016. [arXiv:1611.00036](https://arxiv.org/abs/1611.00036). (cited pp. 16 and 106)
- [186] et al., L.S.C. LSST Science Book, Version 2.0, 2009. [arXiv:0912.0201](https://arxiv.org/abs/0912.0201). (cited p. 16)
- [187] et al. (CMB-S4), K.N.A. CMB-S4 Science Book, First Edition, 2016. [arXiv:1610.02743](https://arxiv.org/abs/1610.02743). (cited p. 16)
- [188] et al. (EUCLID), R.L. Euclid Definition Study Report, 2011. [arXiv:1110.3193](https://arxiv.org/abs/1110.3193). (cited p. 16)
- [189] Ewen, H.I. and Purcell, E.M. Observation of a Line in the Galactic Radio Spectrum: Radiation from Galactic Hydrogen at 1,420 Mc./sec. , 168(4270):356, September 1951. doi:10.1038/168356a0. (cited p. 91)
- [190] Fantaye, Y. et al. CMB lensing reconstruction in the presence of diffuse polarized foregrounds. *Journal of Cosmology and Astroparticle Physics*, 2012(12):017–017, Dec 2012. URL: <http://dx.doi.org/10.1088/1475-7516/2012/12/017>, doi:10.1088/1475-7516/2012/12/017. (cited p. 121)
- [191] Faúndez, M.A. et al. Measurement of the Cosmic Microwave Background Polarization Lensing Power Spectrum from Two Years of POLARBEAR Data. *The Astrophysical Journal*, 893(1):85, Apr 2020. URL: <http://dx.doi.org/10.3847/1538-4357/ab7e29>, doi:10.3847/1538-4357/ab7e29. (cited p. 105)
- [192] Faúndez, M.A. et al. Cross-correlation of CMB Polarization Lensing with High-z Submillimeter Herschel-ATLAS Galaxies. *The Astrophysical Journal*, 886(1):38, Nov 2019. URL: <http://dx.doi.org/10.3847/1538-4357/ab4a78>, doi:10.3847/1538-4357/ab4a78. (cited p. 105)
- [193] Fergusson, J.R., Liguori, M. and Shellard, E.P.S. General CMB and primordial bispectrum estimation: Mode expansion, map making, and measures of FNL. *Physical Review D*, 82(2), Jul 2010. URL: <http://dx.doi.org/10.1103/PhysRevD.82.023502>, doi:10.1103/physrevd.82.023502. (cited p. 74)
- [194] Ferreira, E.G.M. Ultra-Light Dark Matter, 2021. [arXiv:2005.03254](https://arxiv.org/abs/2005.03254). (cited p. 81)
- [195] Ferreira, P.G. and Jaffe, A.H. Simultaneous estimation of noise and signal in cosmic microwave background experiments. *Monthly Notices of the Royal Astronomical Society*, 312(1):89–102, Feb 2000. URL: <http://dx.doi.org/10.1046/j.1365-8711.2000.03108.x>, doi:10.1046/j.1365-8711.2000.03108.x. (cited p. 134)

- [196] Field, G.B. The Spin Temperature of Intergalactic Neutral Hydrogen. , 129:536, May 1959. doi:10.1086/146653. (cited p. 92)
- [197] Field, G.B. The Time Relaxation of a Resonance-Line Profile. , 129:551, May 1959. doi:10.1086/146654. (cited p. 92)
- [198] Fixsen, D.J. THE TEMPERATURE OF THE COSMIC MICROWAVE BACKGROUND. *The Astrophysical Journal*, 707(2):916–920, nov 2009. doi:10.1088/0004-637x/707/2/916. (cited pp. 12, 24, and 76)
- [199] Fixsen, D.J. et al. Cosmic Microwave Background Dipole Spectrum Measured by the COBE FIRAS Instrument. , 420:445, January 1994. doi:10.1086/173575. (cited p. 79)
- [200] Fixsen, D.J. et al. The Cosmic Microwave Background Spectrum from the Full COBE FIRAS Data Set. , 473:576, December 1996. arXiv:astro-ph/9605054, doi:10.1086/178173. (cited pp. 2, 19, and 76)
- [201] Fosalba, P. et al. The onion universe: all sky lightcone simulations in spherical shells. *Monthly Notices of the Royal Astronomical Society*, 391(1):435–446, Nov 2008. URL: <http://dx.doi.org/10.1111/j.1365-2966.2008.13910.x>, doi:10.1111/j.1365-2966.2008.13910.x. (cited p. 67)
- [202] Freese, K., Frieman, J.A. and Olinto, A.V. Natural inflation with pseudo Nambu-Goldstone bosons. *Phys. Rev. Lett.*, 65:3233–3236, Dec 1990. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.65.3233>, doi:10.1103/PhysRevLett.65.3233. (cited p. 36)
- [203] Frenk, C. and White, S. Dark matter and cosmic structure. *Annalen der Physik*, 524(9-10):507–534, Sep 2012. URL: <http://dx.doi.org/10.1002/andp.201200212>, doi:10.1002/andp.201200212. (cited p. 17)
- [204] Friedmann, A. Über die Krümmung des Raumes. 10:377–386, January. doi:10.1007/BF01332580. (cited p. 2)
- [205] Friedmann, A. Über die Möglichkeit einer Welt mit konstanter negativer Krümmung des Raumes. *Zeitschrift für Physik*, 21(1):326–332, December 1924. doi:10.1007/BF01328280. (cited p. 2)
- [206] Fukuda, Y. et al. Evidence for Oscillation of Atmospheric Neutrinos. *Phys. Rev. Lett.*, 81:1562–1567, Aug 1998. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.81.1562>, doi:10.1103/PhysRevLett.81.1562. (cited p. 13)
- [207] Fukugita, M. and Yanagida, T. Barygenesis without grand unification. *Physics Letters B*, 174(1):45–47, 1986. URL: <https://www.sciencedirect.com/science/article/pii/0370269386911263>, doi:[https://doi.org/10.1016/0370-2693\(86\)91126-3](https://doi.org/10.1016/0370-2693(86)91126-3). (cited p. 14)

- [208] Furlanetto, S.R., Zaldarriaga, M. and Hernquist, L. The Growth of HiRegions During Reionization. *The Astrophysical Journal*, 613(1):1–15, Sep 2004. URL: <http://dx.doi.org/10.1086/423025>, doi:10.1086/423025. (cited p. 29)
- [209] Galitzki, N. et al. The Simons Observatory: instrument overview. *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy IX*, Jul 2018. URL: <http://dx.doi.org/10.1117/12.2312985>, doi:10.1117/12.2312985. (cited pp. 107, 114, and 183)
- [210] Gando, A. et al. Search for Majorana Neutrinos Near the Inverted Mass Hierarchy Region with KamLAND-Zen. *Physical Review Letters*, 117(8), Aug 2016. URL: <http://dx.doi.org/10.1103/PhysRevLett.117.082503>, doi:10.1103/physrevlett.117.082503. (cited p. 15)
- [211] García-Bellido, J., Rabadán, R. and Zamora, F. Inflationary scenarios from branes at angles. *Journal of High Energy Physics*, 2002(01):036–036, Jan 2002. URL: <http://dx.doi.org/10.1088/1126-6708/2002/01/036>, doi:10.1088/1126-6708/2002/01/036. (cited p. 36)
- [212] Goeckner-Wald, N. et al. Designs for next generation CMB survey strategies from Chile. *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy IX*, Jul 2018. URL: <http://dx.doi.org/10.1117/12.2313898>, doi:10.1117/12.2313898. (cited p. 163)
- [213] Golub, G.H. and van Loan, C.F. *Matrix Computations*. JHU Press, fourth edition, 2013. URL: <http://www.cs.cornell.edu/cv/GVL4/golubandvanloan.htm>. (cited pp. 134, 152, and 157)
- [214] Gorski, K.M. et al. HEALPix: A Framework for High-Resolution Discretization and Fast Analysis of Data Distributed on the Sphere. *The Astrophysical Journal*, 622(2):759–771, Apr 2005. URL: <http://dx.doi.org/10.1086/427976>, doi:10.1086/427976. (cited p. 160)
- [215] Gott, J. R., I. Creation of open universes from de Sitter space. , 295(5847):304–307, January 1982. doi:10.1038/295304a0. (cited p. 46)
- [216] Grigori, L., Sophie, M. and Frédéric, N. Enlarged Krylov Subspace Conjugate Gradient Methods for Reducing Communication. *SIAM Journal on Matrix Analysis and Applications*, 37(2):744–773, 2016. URL: <https://hal.inria.fr/hal-01357899>, doi:10.1137/140989492. (cited p. 153)
- [217] Grigori, L., Stompor, R. and Szydlarski, M. A parallel two-level preconditioner for Cosmic Microwave Background map-making. In *SC '12: Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis*, pages 1–10, 2012. doi:10.1109/SC.2012.10. (cited p. 167)
- [218] Grigori, L. and Tissot, O. Scalable Linear Solvers Based on Enlarged Krylov Subspaces with Dynamic Reduction of Search Directions. *SIAM Journal on Scientific Computing*, 41(5):C522–C547, 2019. arXiv:<https://doi.org/>

- 10.1137/18M1196285, doi:10.1137/18M1196285. (cited pp. 153 and 154)
- [219] Group, P.D. et al. Review of Particle Physics. *Progress of Theoretical and Experimental Physics*, 2020(8), 08 2020. 083C01. arXiv:<https://academic.oup.com/ptep/article-pdf/2020/8/083C01/34673722/ptaa104.pdf>, doi:10.1093/ptep/ptaa104. (cited pp. viii, x, 13, 14, 15, 27, 28, 72, and 73)
- [220] Gunn, J.E. and Peterson, B.A. On the Density of Neutral Hydrogen in Inter-galactic Space. , 142:1633–1636, November 1965. doi:10.1086/148444. (cited p. 30)
- [221] Guth, A.H. Inflationary universe: A possible solution to the horizon and flatness problems. , 23(2):347–356, January 1981. doi:10.1103/PhysRevD.23.347. (cited p. 34)
- [222] Guth, A.H., Kaiser, D.I. and Nomura, Y. Inflationary paradigm after Planck 2013. *Physics Letters B*, 733:112–119, June 2014. arXiv:1312.7619, doi:10.1016/j.physletb.2014.03.020. (cited p. 34)
- [223] Guth, A.H. and Pi, S.Y. Fluctuations in the New Inflationary Universe. *Phys. Rev. Lett.*, 49:1110–1113, Oct 1982. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.49.1110>, doi:10.1103/PhysRevLett.49.1110. (cited p. 34)
- [224] Guy, J. et al. SALT2: using distant supernovae to improve the use of type Ia supernovae as distance indicators. , 466(1):11–21, April 2007. arXiv:astro-ph/0701828, doi:10.1051/0004-6361:20066930. (cited p. 93)
- [225] Halverson, N.W. et al. Degree Angular Scale Interferometer First Results: A Measurement of the Cosmic Microwave Background Angular Power Spectrum. *The Astrophysical Journal*, 568(1):38–45, Mar 2002. URL: <http://dx.doi.org/10.1086/338879>, doi:10.1086/338879. (cited p. 100)
- [226] Hamann, J., Hannestad, S. and Wong, Y.Y. Measuring neutrino masses with a future galaxy survey. *Journal of Cosmology and Astroparticle Physics*, 2012(11):052–052, Nov 2012. URL: <http://dx.doi.org/10.1088/1475-7516/2012/11/052>, doi:10.1088/1475-7516/2012/11/052. (cited p. 16)
- [227] Hanany, S. et al. MAXIMA-1: A Measurement of the Cosmic Microwave Background Anisotropy on Angular Scales of 10[arcmin]–5°. *The Astrophysical Journal*, 545(1):L5–L9, Dec 2000. URL: <http://dx.doi.org/10.1086/317322>, doi:10.1086/317322. (cited p. 100)
- [228] Hand, N. et al. Evidence of Galaxy Cluster Motions with the Kinematic Sunyaev-Zel’dovich Effect. *Physical Review Letters*, 109(4), Jul 2012. URL: <http://dx.doi.org/10.1103/PhysRevLett.109.041101>, doi:10.1103/physrevlett.109.041101. (cited pp. 79 and 101)

- [229] Hansen, F.K. and Górski, K.M. Fast cosmic microwave background power spectrum estimation of temperature and polarization with Gabor transforms. , 343(2):559–584, August 2003. [arXiv:astro-ph/0207526](#), [doi:10.1046/j.1365-8711.2003.06695.x](#). (cited p. 124)
- [230] Hanson, D. et al. Detection of B-Mode Polarization in the Cosmic Microwave Background with Data from the South Pole Telescope. *Physical Review Letters*, 111(14), Sep 2013. URL: <http://dx.doi.org/10.1103/PhysRevLett.111.141301>, [doi:10.1103/physrevlett.111.141301](#). (cited p. 70)
- [231] Harari, D.D., Hayward, J.D. and Zaldarriaga, M. Depolarization of the cosmic microwave background by a primordial magnetic field and its effect upon temperature anisotropy. *Physical Review D*, 55(4):1841–1850, Feb 1997. URL: <http://dx.doi.org/10.1103/PhysRevD.55.1841>, [doi:10.1103/physrevd.55.1841](#). (cited p. 81)
- [232] Harrington, K. et al. The Cosmology Large Angular Scale Surveyor. In Holland, W.S. and Zmuidzinas, J., editors, *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy VIII*, volume 9914, pages 380 – 400. International Society for Optics and Photonics, SPIE, 2016. URL: <https://doi.org/10.1117/12.2233125>. (cited p. 73)
- [233] Harrison, D.L., van Leeuwen, F. and Ashdown, M.A.J. A deconvolution map-making method for experiments with circular scanning strategies. *Astronomy Astrophysics*, 532:A55, Jul 2011. URL: <http://dx.doi.org/10.1051/0004-6361/201116986>, [doi:10.1051/0004-6361/201116986](#). (cited p. 133)
- [234] Harrison, E.R. Fluctuations at the Threshold of Classical Cosmology. , 1(10):2726–2730, May 1970. [doi:10.1103/PhysRevD.1.2726](#). (cited p. 51)
- [235] Hauser, M.G. and Dwek, E. The Cosmic Infrared Background: Measurements and Implications. *Annual Review of Astronomy and Astrophysics*, 39(1):249–307, Sep 2001. URL: <http://dx.doi.org/10.1146/annurev.astro.39.1.249>, [doi:10.1146/annurev.astro.39.1.249](#). (cited p. 70)
- [236] Hawking, S.W. Black hole explosions? , 248(5443):30–31, March 1974. [doi:10.1038/248030a0](#). (cited p. 31)
- [237] Hawking, S.W. Particle creation by black holes. *Communications in Mathematical Physics*, 43(3):199–220, August 1975. [doi:10.1007/BF02345020](#). (cited p. 31)
- [238] Hawking, S.W. The development of irregularities in a single bubble inflationary universe. *Physics Letters B*, 115(4):295–297, September 1982. [doi:10.1016/0370-2693\(82\)90373-2](#). (cited p. 34)

- [239] Hawkins, E. et al. The 2dF Galaxy Redshift Survey: correlation functions, peculiar velocities and the matter density of the Universe. *Monthly Notices of the Royal Astronomical Society*, 346(1):78–96, 11 2003. arXiv:<https://academic.oup.com/mnras/article-pdf/346/1/78/9375978/346-1-78.pdf>, doi:10.1046/j.1365-2966.2003.07063.x. (cited pp. xi and 87)
- [240] Hazumi, M. et al. LiteBIRD satellite: JAXA’s new strategic L-class mission for all-sky surveys of cosmic microwave background polarization. *Space Telescopes and Instrumentation 2020: Optical, Infrared, and Millimeter Wave*, Dec 2020. URL: <http://dx.doi.org/10.1117/12.2563050>, doi:10.1117/12.2563050. (cited pp. 111, 177, and 180)
- [241] Hensley, B.S. and Bull, P. Mitigating Complex Dust Foregrounds in Future Cosmic Microwave Background Polarization Experiments. *The Astrophysical Journal*, 853(2):127, Jan 2018. URL: <http://dx.doi.org/10.3847/1538-4357/aaa489>, doi:10.3847/1538-4357/aaa489. (cited p. 117)
- [242] Hill, C.A. et al. A cryogenic continuously rotating half-wave plate mechanism for the POLARBEAR-2b cosmic microwave background receiver. *Review of Scientific Instruments*, 91(12):124503, Dec 2020. URL: <http://dx.doi.org/10.1063/5.0029006>, doi:10.1063/5.0029006. (cited p. 106)
- [243] Hill, C.A. et al. Design and development of an ambient-temperature continuously-rotating achromatic half-wave plate for CMB polarization modulation on the POLARBEAR-2 experiment. *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy VIII*, Jul 2016. URL: <http://dx.doi.org/10.1117/12.2232280>, doi:10.1117/12.2232280. (cited p. 106)
- [244] Hilton, M. et al. The Atacama Cosmology Telescope: A Catalog of >4000 Sunyaev–Zel’dovich Galaxy Clusters. *The Astrophysical Journal Supplement Series*, 253(1):3, Feb 2021. URL: <http://dx.doi.org/10.3847/1538-4365/abd023>, doi:10.3847/1538-4365/abd023. (cited p. 79)
- [245] Hivon, E. et al. MASTER of the Cosmic Microwave Background Anisotropy Power Spectrum: A Fast Method for Statistical Analysis of Large and Complex Cosmic Microwave Background Data Sets. *The Astrophysical Journal*, 567(1):2–17, Mar 2002. URL: <http://dx.doi.org/10.1086/338126>, doi:10.1086/338126. (cited pp. 121 and 124)
- [246] Holz, D.E. and Hughes, S.A. Using Gravitational-Wave Standard Sirens. *The Astrophysical Journal*, 629(1):15–22, Aug 2005. URL: <http://dx.doi.org/10.1086/431341>, doi:10.1086/431341. (cited p. 97)
- [247] Hoskin, M.A. The ‘Great Debate’: What Really Happened. *Journal for the History of Astronomy*, 7:169, January 1976. doi:10.1177/002182867600700302. (cited p. 1)

- [248] Howlett, C. et al. Measuring the Growth Rate of Structure with Type IA Supernovae from LSST. *The Astrophysical Journal*, 847(2):128, Sep 2017. URL: <http://dx.doi.org/10.3847/1538-4357/aa88c8>, doi:10.3847/1538-4357/aa88c8. (cited p. 94)
- [249] Hoyle, F. A New Model for the Expanding Universe. , 108:372, January 1948. doi:10.1093/mnras/108.5.372. (cited p. 51)
- [250] Hu, W. Dark Energy Probes in Light of the CMB. In Wolff, S.C. and Lauer, T.R., editors, *Observing Dark Energy*, volume 339 of *Astronomical Society of the Pacific Conference Series*, page 215, August 2005. arXiv:astro-ph/0407158. (cited p. 61)
- [251] Hu, W. Weak lensing of the CMB: A harmonic approach. *Physical Review D*, 62(4), Jul 2000. URL: <http://dx.doi.org/10.1103/PhysRevD.62.043007>, doi:10.1103/physrevd.62.043007. (cited pp. x and 68)
- [252] Hu, W. CMB temperature and polarization anisotropy fundamentals. *Annals of Physics*, 303(1):203–225, Jan 2003. URL: [http://dx.doi.org/10.1016/S0003-4916\(02\)00022-2](http://dx.doi.org/10.1016/S0003-4916(02)00022-2), doi:10.1016/s0003-4916(02)00022-2. (cited p. 61)
- [253] Hu, W. Lecture Notes on CMB Theory: From Nucleosynthesis to Recombination, 2008. arXiv:0802.3688. (cited p. 59)
- [254] Hu, W. and Okamoto, T. Mass Reconstruction with Cosmic Microwave Background Polarization. *The Astrophysical Journal*, 574(2):566–574, Aug 2002. URL: <http://dx.doi.org/10.1086/341110>, doi:10.1086/341110. (cited p. 119)
- [255] Hu, W. and White, M. A CMB polarization primer. *New Astronomy*, 2(4):323–344, Oct 1997. URL: [http://dx.doi.org/10.1016/S1384-1076\(97\)00022-5](http://dx.doi.org/10.1016/S1384-1076(97)00022-5), doi:10.1016/s1384-1076(97)00022-5. (cited p. 59)
- [256] Hubble, E.P. Cepheids in Spiral Nebulae. *Popular Astronomy*, 33:252–255, January 1925. (cited p. 1)
- [257] Hubble, E.P. NGC 6822, a remote stellar system. , 62:409–433, December 1925. doi:10.1086/142943. (cited p. 1)
- [258] Hubble, E. A Relation between Distance and Radial Velocity among Extra-Galactic Nebulae. *Proceedings of the National Academy of Science*, 15(3):168–173, March 1929. doi:10.1073/pnas.15.3.168. (cited pp. 2 and 5)
- [259] Hui, H. et al. BICEP Array: a multi-frequency degree-scale CMB polarimeter. *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy IX*, Jul 2018. URL: <http://dx.doi.org/10.1117/12.2311725>, doi:10.1117/12.2311725. (cited p. 102)
- [260] Humphreys, E.M.L. et al. Toward a New Geometric Distance to the Active Galaxy NGC 4258. III. Final Results and the Hubble Constant. , 775(1):13,

- September 2013. [arXiv:1307.6031](#), [doi:10.1088/0004-637X/775/1/13](#). (cited p. 94)
- [261] Hurier, G., Macías-Pérez, J.F. and Hildebrandt, S. MILCA, a modified internal linear combination algorithm to extract astrophysical emissions from multifrequency sky maps. , 558:A118, October 2013. [arXiv:1007.1149](#), [doi:10.1051/0004-6361/201321891](#). (cited p. 122)
- [262] Ijjas, A., Steinhardt, P.J. and Loeb, A. Inflationary paradigm in trouble after Planck2013. *Physics Letters B*, 723(4-5):261–266, June 2013. [arXiv:1304.2785](#), [doi:10.1016/j.physletb.2013.05.023](#). (cited p. 34)
- [263] Ijjas, A., Steinhardt, P.J. and Loeb, A. Inflationary schism. *Physics Letters B*, 736:142–146, September 2014. [arXiv:1402.6980](#), [doi:10.1016/j.physletb.2014.07.012](#). (cited p. 34)
- [264] Jewell, J., Levin, S. and Anderson, C.H. Application of Monte Carlo Algorithms to the Bayesian Analysis of the Cosmic Microwave Background. , 609(1):1–14, July 2004. [arXiv:astro-ph/0209560](#), [doi:10.1086/383515](#). (cited p. 124)
- [265] Jha, S. *Exploding stars, near and far*. PhD thesis, HARVARD UNIVERSITY, October 2002. (cited pp. viii and 5)
- [266] Johnson, B.R. et al. MAXIPOL: Cosmic Microwave Background Polarimetry Using a Rotating Half-Wave Plate. *The Astrophysical Journal*, 665(1):42–54, Aug 2007. URL: <http://dx.doi.org/10.1086/518105>, [doi:10.1086/518105](#). (cited pp. 114, 137, and 185)
- [267] Kachru, S. et al. Towards inflation in string theory. *Journal of Cosmology and Astroparticle Physics*, 2003(10):013–013, Oct 2003. URL: <http://dx.doi.org/10.1088/1475-7516/2003/10/013>, [doi:10.1088/1475-7516/2003/10/013](#). (cited p. 36)
- [268] KAGRA Collaboration et al. Overview of KAGRA : KAGRA science, 2020. [arXiv:2008.02921](#). (cited p. 96)
- [269] Kaiser, N. On the spatial correlations of Abell clusters. , 284:L9–L12, September 1984. [doi:10.1086/184341](#). (cited p. 84)
- [270] Kaiser, N. Clustering in real space and in redshift space. *Monthly Notices of the Royal Astronomical Society*, 227(1):1–21, 07 1987. [arXiv:https://academic.oup.com/mnras/article-pdf/227/1/1/18522208/mnras227-0001.pdf](#), [doi:10.1093/mnras/227.1.1](#). (cited p. 86)
- [271] Kallosh, R. and Linde, A. Dark energy and the fate of the Universe. *Journal of Cosmology and Astroparticle Physics*, 2003(02):002–002, Feb 2003. URL: <http://dx.doi.org/10.1088/1475-7516/2003/02/002>, [doi:10.1088/1475-7516/2003/02/002](#). (cited p. 31)
- [272] Kamionkowski, M., Kosowsky, A. and Stebbins, A. Statistics of cosmic microwave background polarization. *Physical Review D*, 55(12):7368–7388,

- Jun 1997. URL: <http://dx.doi.org/10.1103/PhysRevD.55.7368>, doi:10.1103/physrevd.55.7368. (cited p. 56)
- [273] Kaneko, D. et al. Deployment of Polarbear-2A. In *18th International Workshop on Low Temperature Detectors*, volume 199, pages 1137–1147, Milano, Italy, July 2019. URL: <https://hal.archives-ouvertes.fr/hal-02614360>, doi:10.1007/s10909-020-02366-w. (cited pp. xii, 105, and 106)
- [274] Kawasaki, M., Yamada, M. and Yanagida, T.T. Observable dark radiation from a cosmologically safe QCD axion. *Phys. Rev. D*, 91:125018, Jun 2015. URL: <https://link.aps.org/doi/10.1103/PhysRevD.91.125018>, doi:10.1103/PhysRevD.91.125018. (cited p. 16)
- [275] Keating, B.G. et al. BICEP: a large angular scale CMB polarimeter. In Fineschi, S., editor, *Polarimetry in Astronomy*, volume 4843 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, pages 284–295, February 2003. doi:10.1117/12.459274. (cited p. 101)
- [276] Keihänen, E. et al. A maximum likelihood approach to the destriping technique. , 428:287–298, December 2004. arXiv:astro-ph/0304411, doi:10.1051/0004-6361:200400060. (cited p. 134)
- [277] Keihänen, E. et al. Making cosmic microwave background temperature and polarization maps with MADAM. *Astronomy and Astrophysics*, 510:A57, Feb 2010. URL: <http://dx.doi.org/10.1051/0004-6361/200912813>, doi:10.1051/0004-6361/200912813. (cited pp. 180 and 182)
- [278] Keihänen, E., Kurki-Suonio, H. and Poutanen, T. Madam- a map-making method for CMB experiments. *Monthly Notices of the Royal Astronomical Society*, 360(1):390–400, Jun 2005. URL: <http://dx.doi.org/10.1111/j.1365-2966.2005.09055.x>, doi:10.1111/j.1365-2966.2005.09055.x. (not cited)
- [279] Keihänen, E. and Reinecke, M. ArtDeco: a beam-deconvolution code for absolute cosmic microwave background measurements. *Astronomy Astrophysics*, 548:A110, Nov 2012. URL: <http://dx.doi.org/10.1051/0004-6361/201220183>, doi:10.1051/0004-6361/201220183. (cited p. 133)
- [280] Kermish, Z.D. et al. The POLARBEAR experiment. *Millimeter, Sub-millimeter, and Far-Infrared Detectors and Instrumentation for Astronomy VI*, Sep 2012. URL: <http://dx.doi.org/10.1117/12.926354>, doi:10.1117/12.926354. (cited pp. 102 and 161)
- [281] Kesden, M., Cooray, A. and Kamionkowski, M. Separation of Gravitational-Wave and Cosmic-Shear Contributions to Cosmic Microwave Background Polarization. *Physical Review Letters*, 89(1), Jun 2002. URL: <http://dx.doi.org/10.1103/PhysRevLett.89.011304>, doi:10.1103/physrevlett.89.011304. (cited p. 118)

- [282] Khosravi, N. et al. H0 tension as a hint for a transition in gravitational theory. *Physical Review D*, 99(10), May 2019. URL: <http://dx.doi.org/10.1103/PhysRevD.99.103526>, doi:10.1103/physrevd.99.103526. (cited p. 94)
- [283] Kirshner, R.P. Hubble's diagram and cosmic expansion. *Proceedings of the National Academy of Sciences*, 101(1):8–13, 2004. URL: <https://www.pnas.org/content/101/1/8>, arXiv: <https://www.pnas.org/content/101/1/8.full.pdf>, doi:10.1073/pnas.2536799100. (cited pp. viii and 5)
- [284] Kitayama, T. et al. The Sunyaev–Zel'dovich effect at 5: RX J1347.51145 imaged by ALMA. *Publications of the Astronomical Society of Japan*, 68(5):88, Sep 2016. URL: <http://dx.doi.org/10.1093/pasj/psw082>, doi:10.1093/pasj/psw082. (cited p. 79)
- [285] Knox, L. and Song, Y.S. Limit on the Detectability of the Energy Scale of Inflation. *Physical Review Letters*, 89(1), Jun 2002. URL: <http://dx.doi.org/10.1103/PhysRevLett.89.011303>, doi:10.1103/physrevlett.89.011303. (cited p. 118)
- [286] Kobakhidze, A. et al. Gravitational waves from a supercooled electroweak phase transition and their detection with pulsar timing arrays. *The European Physical Journal C*, 77(8), Aug 2017. URL: <http://dx.doi.org/10.1140/epjc/s10052-017-5132-y>, doi:10.1140/epjc/s10052-017-5132-y. (cited p. 96)
- [287] Kofman, L. Probing String Theory with Modulated Cosmological Fluctuations. *arXiv e-prints*, pages astro-ph/0303614, March 2003. arXiv: astro-ph/0303614. (cited p. 75)
- [288] Kogut, A. SYNCHROTRON SPECTRAL CURVATURE FROM 22 MHz TO 23 GHz. *The Astrophysical Journal*, 753(2):110, Jun 2012. URL: <http://dx.doi.org/10.1088/0004-637X/753/2/110>, doi:10.1088/0004-637x/753/2/110. (cited p. 118)
- [289] Kogut, A. et al. Three-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Foreground Polarization. *The Astrophysical Journal*, 665(1):355–362, Aug 2007. URL: <http://dx.doi.org/10.1086/519754>, doi:10.1086/519754. (cited p. 118)
- [290] Kogut, A. et al. The Primordial Inflation Explorer (PIXIE): a nulling polarimeter for cosmic microwave background observations. *Journal of Cosmology and Astroparticle Physics*, 2011(07):025–025, Jul 2011. URL: <http://dx.doi.org/10.1088/1475-7516/2011/07/025>, doi:10.1088/1475-7516/2011/07/025. (cited p. 79)
- [291] Komatsu, E. and Spergel, D.N. Acoustic signatures in the primary microwave background bispectrum. *Physical Review D*, 63(6), Feb 2001. URL: <http://dx.doi.org/10.1103/PhysRevD.63.063002>, doi:10.1103/physrevd.63.063002. (cited p. 74)

- [292] Komatsu, K. et al. Demonstration of the broadband half-wave plate using the nine-layer sapphire for the cosmic microwave background polarization experiment. *J. Astron. Telesc. Instrum. Syst.*, 5(4):044008, 2019. [arXiv:1905.13520](#), [doi:10.1117/1.JATIS.5.4.044008](#). (cited p. 196)
- [293] Koopman, B.J. et al. Advanced ACTPol Low-Frequency Array: Read-out and Characterization of Prototype 27 and 39 GHz Transition Edge Sensors. *Journal of Low Temperature Physics*, 193(5-6):1103–1111, May 2018. URL: <http://dx.doi.org/10.1007/s10909-018-1957-5>, [doi:10.1007/s10909-018-1957-5](#). (cited p. 101)
- [294] Korngut, P.M. et al. MUSTANG High Angular Resolution Sunyaev-Zel’dovich Effect Imaging of Substructure in Four Galaxy Clusters. , 734(1):10, June 2011. [arXiv:1010.5494](#), [doi:10.1088/0004-637X/734/1/10](#). (cited p. 79)
- [295] Kosowsky, A. Cosmic Microwave Background Polarization. *Annals of Physics*, 246(1):49–85, Feb 1996. URL: <http://dx.doi.org/10.1006/aphy.1996.0020>, [doi:10.1006/aphy.1996.0020](#). (cited p. 59)
- [296] Kosowsky, A. and Loeb, A. Faraday Rotation of Microwave Background Polarization by a Primordial Magnetic Field. *The Astrophysical Journal*, 469:1, Sep 1996. URL: <http://dx.doi.org/10.1086/177751>, [doi:10.1086/177751](#). (cited p. 81)
- [297] Kovac, J.M. et al. Detection of polarization in the cosmic microwave background using DASI. *Nature*, 420(6917):772–787, Dec 2002. URL: <http://dx.doi.org/10.1038/nature01269>, [doi:10.1038/nature01269](#). (cited p. 100)
- [298] Koyama, K. Cosmological tests of modified gravity. *Reports on Progress in Physics*, 79(4):046902, Mar 2016. URL: <http://dx.doi.org/10.1088/0034-4885/79/4/046902>, [doi:10.1088/0034-4885/79/4/046902](#). (cited p. 18)
- [299] Koyama, K. et al. Non-Gaussianities from ekpyrotic collapse with multiple fields. *Journal of Cosmology and Astroparticle Physics*, 2007(11):024–024, Nov 2007. URL: <http://dx.doi.org/10.1088/1475-7516/2007/11/024>, [doi:10.1088/1475-7516/2007/11/024](#). (cited p. 75)
- [300] Krachmalnicoff, N. et al. S-PASS view of polarized Galactic synchrotron at 2.3 GHz as a contaminant to CMB observations. *Astronomy Astrophysics*, 618:A166, Oct 2018. URL: <http://dx.doi.org/10.1051/0004-6361/201832768>, [doi:10.1051/0004-6361/201832768](#). (cited p. 118)
- [301] Kreisch, C.D., Cyr-Racine, F.Y. and Doré, O. Neutrino puzzle: Anomalies, interactions, and cosmological tensions. *Physical Review D*, 101(12), Jun 2020. URL: <http://dx.doi.org/10.1103/PhysRevD.101.123505>, [doi:10.1103/physrevd.101.123505](#). (cited p. 94)

- [302] Kuo, C.L. et al. High-Resolution Observations of the Cosmic Microwave Background Power Spectrum with ACBAR. *The Astrophysical Journal*, 600(1):32–51, Jan 2004. URL: <http://dx.doi.org/10.1086/379783>, doi:10.1086/379783. (cited p. 100)
- [303] Kurki-Suonio, H. et al. Destriping CMB temperature and polarization maps. *Astronomy Astrophysics*, 506(3):1511–1539, Aug 2009. URL: <http://dx.doi.org/10.1051/0004-6361/200912361>, doi:10.1051/0004-6361/200912361. (cited p. 182)
- [304] Kusaka, A. et al. Modulation of cosmic microwave background polarization with a warm rapidly rotating half-wave plate on the Atacama B-Mode Search instrument. *Review of Scientific Instruments*, 85(2):024501, Feb 2014. URL: <http://dx.doi.org/10.1063/1.4862058>, doi:10.1063/1.4862058. (cited pp. 114, 137, and 185)
- [305] Larsen, P. et al. Demonstration of Cosmic Microwave Background Delensing Using the Cosmic Infrared Background. *Physical Review Letters*, 117(15), Oct 2016. URL: <http://dx.doi.org/10.1103/PhysRevLett.117.151102>, doi:10.1103/physrevlett.117.151102. (cited p. 100)
- [306] Lasky, P.D. et al. Gravitational-Wave Cosmology across 29 Decades in Frequency. *Phys. Rev. X*, 6:011035, Mar 2016. URL: <https://link.aps.org/doi/10.1103/PhysRevX.6.011035>, doi:10.1103/PhysRevX.6.011035. (cited p. 96)
- [307] Leavitt, H.S. 1777 variables in the Magellanic Clouds. *Annals of Harvard College Observatory*, 60:87–108.3, January 1908. (cited p. 94)
- [308] Leavitt, H.S. and Pickering, E.C. Periods of 25 Variable Stars in the Small Magellanic Cloud. *Harvard College Observatory Circular*, 173:1–3, March 1912. (cited p. 94)
- [309] Lehnert, J.L. and Steinhardt, P.J. Non-Gaussian density fluctuations from entropically generated curvature perturbations in ekpyrotic models. *Phys. Rev. D*, 77:063533, Mar 2008. URL: <https://link.aps.org/doi/10.1103/PhysRevD.77.063533>, doi:10.1103/PhysRevD.77.063533. (cited p. 75)
- [310] Lemaître, G. Un Univers homogène de masse constante et de rayon croissant rendant compte de la vitesse radiale des nébuleuses extra-galactiques. *Annales de la Société Scientifique de Bruxelles*, 47:49–59, January 1927. (cited p. 2)
- [311] LEWIS, A. and CHALLINOR, A. Weak gravitational lensing of the CMB. *Physics Reports*, 429(1):1–65, Jun 2006. URL: <http://dx.doi.org/10.1016/j.physrep.2006.03.002>, doi:10.1016/j.physrep.2006.03.002. (cited pp. x, 66, and 67)
- [312] Lewis, A. and Bridle, S. Cosmological parameters from CMB and other data: A Monte Carlo approach. *Physical Review D*, 66(10), Nov

2002. URL: <http://dx.doi.org/10.1103/PhysRevD.66.103511>, doi:10.1103/physrevd.66.103511. (cited p. 125)
- [313] Liddle, A.R. and Leach, S.M. How long before the end of inflation were observable perturbations produced? *Physical Review D*, 68(10), Nov 2003. URL: <http://dx.doi.org/10.1103/PhysRevD.68.103503>, doi:10.1103/physrevd.68.103503. (cited p. 39)
- [314] Linde, A. and Mukhanov, V. Non-Gaussian isocurvature perturbations from inflation. *Physical Review D*, 56(2):R535–R539, Jul 1997. URL: <http://dx.doi.org/10.1103/PhysRevD.56.R535>, doi:10.1103/physrevd.56.r535. (cited p. 75)
- [315] LiteBIRD Collaboration et al. Science from LiteBIRD. (cited p. 180)
- [316] Liu, J., Ortiz-Vazquez, A. and Hill, J.C. Constraining multiplicative bias in CFHTLenS weak lensing shear data. *Physical Review D*, 93(10), May 2016. URL: <http://dx.doi.org/10.1103/PhysRevD.93.103508>, doi:10.1103/physrevd.93.103508. (cited p. 89)
- [317] LSST Science Collaboration et al. LSST Science Book, Version 2.0, 2009. arXiv:0912.0201. (cited pp. 94 and 106)
- [318] Lucca, M. et al. The synergy between CMB spectral distortions and anisotropies. *Journal of Cosmology and Astroparticle Physics*, 2020(02):026–026, Feb 2020. URL: <http://dx.doi.org/10.1088/1475-7516/2020/02/026>, doi:10.1088/1475-7516/2020/02/026. (cited p. 77)
- [319] Lynds, R. The Absorption-Line Spectrum of 4c 05.34. , 164:L73, March 1971. doi:10.1086/180695. (cited p. 90)
- [320] Lyth, D.H. and Stewart, E.D. Inflationary density perturbations with $\Omega < 1$. *Physics Letters B*, 252(3):336–342, December 1990. doi:10.1016/0370-2693(90)90548-K. (cited p. 46)
- [321] Lyth, D.H., Ungarelli, C. and Wands, D. Primordial density perturbation in the curvaton scenario. *Physical Review D*, 67(2), Jan 2003. URL: <http://dx.doi.org/10.1103/PhysRevD.67.023503>, doi:10.1103/physrevd.67.023503. (cited p. 75)
- [322] Maartens, R. et al. Cosmology with the SKA – overview, 2015. arXiv:1501.04076. (cited p. 91)
- [323] Macquart, J.P. et al. A census of baryons in the Universe from localized fast radio bursts. *Nature*, 581(7809):391–395, May 2020. URL: <http://dx.doi.org/10.1038/s41586-020-2300-2>, doi:10.1038/s41586-020-2300-2. (cited p. 13)
- [324] Maino, D. et al. The Planck-LFI instrument: Analysis of the $1/f$ noise and implications for the scanning strategy. *Astronomy and Astrophysics Supplement Series*, 140(3):383–391, Dec 1999. URL: <http://dx.doi.org/>

- 10.1051/aas:1999429, doi:10.1051/aas:1999429. (cited pp. 130 and 134)
- [325] Maino, D. et al. All-sky astrophysical component separation with Fast Independent Component Analysis (FASTICA). *Monthly Notices of the Royal Astronomical Society*, 334(1):53–68, Jul 2002. URL: <http://dx.doi.org/10.1046/j.1365-8711.2002.05425.x>, doi:10.1046/j.1365-8711.2002.05425.x. (cited p. 122)
- [326] Marsh, D.J. Axion cosmology. *Physics Reports*, 643:1–79, Jul 2016. URL: <http://dx.doi.org/10.1016/j.physrep.2016.06.005>, doi:10.1016/j.physrep.2016.06.005. (cited pp. 17 and 81)
- [327] Martin, J. and Ringeval, C. First CMB constraints on the inflationary reheating temperature. *Physical Review D*, 82(2), Jul 2010. URL: <http://dx.doi.org/10.1103/PhysRevD.82.023511>, doi:10.1103/physrevd.82.023511. (cited p. 39)
- [328] Mather, J.C. et al. Measurement of the Cosmic Microwave Background Spectrum by the COBE FIRAS Instrument. , 420:439, January 1994. doi:10.1086/173574. (cited p. 99)
- [329] Mather, J.C. et al. Calibrator Design for the COBE Far-Infrared Absolute Spectrophotometer (FIRAS). , 512(2):511–520, February 1999. arXiv:astro-ph/9810373, doi:10.1086/306805. (cited p. 2)
- [330] Matsumura, T. et al. LiteBIRD: Mission Overview and Focal Plane Layout. *Journal of Low Temperature Physics*, 184(3-4):824–831, August 2016. doi:10.1007/s10909-016-1542-8. (cited p. 73)
- [331] Mention, G. et al. Reactor antineutrino anomaly. *Phys. Rev. D*, 83:073006, Apr 2011. URL: <https://link.aps.org/doi/10.1103/PhysRevD.83.073006>, doi:10.1103/PhysRevD.83.073006. (cited p. 18)
- [332] Millea, M., Anderes, E. and Wandelt, B.D. Bayesian delensing of CMB temperature and polarization. *Physical Review D*, 100(2), Jul 2019. URL: <http://dx.doi.org/10.1103/PhysRevD.100.023509>, doi:10.1103/physrevd.100.023509. (cited p. 118)
- [333] Miller, A.D. et al. A Measurement of the Angular Power Spectrum of the Cosmic Microwave Background from $L = 100$ to 400. , 524(1):L1–L4, October 1999. arXiv:astro-ph/9906421, doi:10.1086/312293. (cited p. 99)
- [334] Minami, Y. and Komatsu, E. New Extraction of the Cosmic Birefringence from the Planck 2018 Polarization Data. *Physical Review Letters*, 125(22), Nov 2020. URL: <http://dx.doi.org/10.1103/PhysRevLett.125.221301>, doi:10.1103/physrevlett.125.221301. (cited p. 82)
- [335] Minami, Y. et al. Simultaneous determination of the cosmic birefringence and miscalibrated polarization angles from CMB experiments. *Progress of Theoretical and Experimental Physics*, 2019(8), 08 2019. 083E02.

- arXiv:<https://academic.oup.com/ptep/article-pdf/2019/8/083E02/29117835/ptz079.pdf>, doi:10.1093/ptep/ptz079. (cited p. 82)
- [336] Miyazaki, S. et al. Hyper Suprime-Cam. In McLean, I.S., Ramsay, S.K. and Takami, H., editors, *Ground-based and Airborne Instrumentation for Astronomy IV*, volume 8446, pages 327 – 335. International Society for Optics and Photonics, SPIE, 2012. URL: <https://doi.org/10.1117/12.926844>. (cited p. 106)
- [337] Moniez, M. Microlensing as a probe of the Galactic structure: 20 years of microlensing optical depth studies. *General Relativity and Gravitation*, 42(9):2047–2074, Feb 2010. URL: <http://dx.doi.org/10.1007/s10714-009-0925-4>, doi:10.1007/s10714-009-0925-4. (cited p. 17)
- [338] Monsalve, R.A. et al. Calibration of the EDGES High-band Receiver to Observe the Global 21 cm Signature from the Epoch of Reionization. , 835(1):49, January 2017. arXiv:1602.08065, doi:10.3847/1538-4357/835/1/49. (cited p. 91)
- [339] Montier, L. et al. Overview of the medium and high frequency telescopes of the LiteBIRD space mission. *Space Telescopes and Instrumentation 2020: Optical, Infrared, and Millimeter Wave*, Dec 2020. URL: <http://dx.doi.org/10.1117/12.2562243>, doi:10.1117/12.2562243. (cited pp. xii and 111)
- [340] Morales, M.F. and Wyithe, J.S.B. Reionization and Cosmology with 21-cm Fluctuations. *Annual Review of Astronomy and Astrophysics*, 48(1):127–171, Aug 2010. URL: <http://dx.doi.org/10.1146/annurev-astro-081309-130936>, doi:10.1146/annurev-astro-081309-130936. (cited p. 30)
- [341] Mroczkowski, T. et al. A MULTI-WAVELENGTH STUDY OF THE SUNYAEV-ZEL'DOVICH EFFECT IN THE TRIPLE-MERGER CLUSTER MACS J0717.5+3745 WITH MUSTANG AND BOLOCAM. *The Astrophysical Journal*, 761(1):47, Nov 2012. URL: <http://dx.doi.org/10.1088/0004-637X/761/1/47>, doi:10.1088/0004-637x/761/1/47. (cited p. 79)
- [342] Mukhanov, V.F. and Chibisov, G.V. Quantum fluctuations and a nonsingular universe. *Soviet Journal of Experimental and Theoretical Physics Letters*, 33:532, May 1981. (cited p. 34)
- [343] Murayama, H. et al. Chaotic inflation and baryogenesis by right-handed sneutrinos. *Phys. Rev. Lett.*, 70:1912–1915, Mar 1993. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.70.1912>, doi:10.1103/PhysRevLett.70.1912. (cited p. 36)
- [344] Mörtzell, E. and Dhawan, S. Does the Hubble constant tension call for new physics? *Journal of Cosmology and Astroparticle Physics*, 2018(09):025–025,

- Sep 2018. URL: <http://dx.doi.org/10.1088/1475-7516/2018/09/025>, doi:10.1088/1475-7516/2018/09/025. (cited p. 94)
- [345] Naess, S. et al. The Atacama Cosmology Telescope: CMB polarization at $200 < \ell < 9000$. *Journal of Cosmology and Astroparticle Physics*, 2014(10):007–007, Oct 2014. URL: <http://dx.doi.org/10.1088/1475-7516/2014/10/007>, doi:10.1088/1475-7516/2014/10/007. (cited p. 185)
- [346] Nakama, T., Carr, B. and Silk, J. Limits on primordial black holes from μ distortions in cosmic microwave background. *Phys. Rev. D*, 97:043525, Feb 2018. URL: <https://link.aps.org/doi/10.1103/PhysRevD.97.043525>, doi:10.1103/PhysRevD.97.043525. (cited p. 18)
- [347] Namikawa, T. et al. Evidence for the Cross-correlation between Cosmic Microwave Background Polarization Lensing from Polarbear and Cosmic Shear from Subaru Hyper Suprime-Cam. *The Astrophysical Journal*, 882(1):62, Sep 2019. URL: <http://dx.doi.org/10.3847/1538-4357/ab3424>, doi:10.3847/1538-4357/ab3424. (cited p. 105)
- [348] Namikawa, T. et al. Atacama Cosmology Telescope: Constraints on cosmic birefringence. *Physical Review D*, 101(8), Apr 2020. URL: <http://dx.doi.org/10.1103/PhysRevD.101.083527>, doi:10.1103/physrevd.101.083527. (cited p. 82)
- [349] Narlikar, J.V. and Wickramasinghe, N.C. Microwave Background in a Steady-state Universe. , 216(5110):43–44, October 1967. doi:10.1038/216043a0. (cited p. 51)
- [350] Natoli, P. et al. A Map-Making algorithm for the Planck Surveyor. *Astronomy Astrophysics*, 372(1):346–356, Jun 2001. URL: <http://dx.doi.org/10.1051/0004-6361:20010393>, doi:10.1051/0004-6361:20010393. (cited p. 153)
- [351] Nicastro, F. et al. Observations of the missing baryons in the warm-hot intergalactic medium. *Nature*, 558(7710):406–409, Jun 2018. URL: <http://dx.doi.org/10.1038/s41586-018-0204-1>, doi:10.1038/s41586-018-0204-1. (cited p. 13)
- [352] O'Dwyer, I.J. et al. Bayesian Power Spectrum Analysis of the First-Year Wilkinson Microwave Anisotropy Probe Data. *The Astrophysical Journal*, 617(2):L99–L102, nov 2004. doi:10.1086/427386. (cited p. 124)
- [353] Ogburn, R. W., I. et al. The BICEP2 CMB polarization experiment. In Holland, W.S. and Zmuidzinas, J., editors, *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy V*, volume 7741 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 77411G, July 2010. doi:10.1117/12.857864. (cited p. 102)
- [354] P. A. R. Ade, T.P.C. et al. A MEASUREMENT OF THE COSMIC MICROWAVE BACKGROUND B-MODE POLARIZATION POWER SPECTRUM AT SUB-DEGREE SCALES WITH POLARBEAR. *The Astrophysical Journal*, 794(2):171, Oct 2014. URL: <http://dx.doi.org/>

- 10.1088/0004-637X/794/2/171, doi:10.1088/0004-637X/794/2/171. (cited pp. 104 and 121)
- [355] Pacheco, P. *An Introduction to Parallel Programming*. Morgan Kaufmann, 2011. (cited p. 142)
- [356] Palanque-Delabrouille, N. et al. Hints, neutrino bounds, and WDM constraints from SDSS DR14 Lyman- α and Planck full-survey data. *Journal of Cosmology and Astroparticle Physics*, 2020(04):038–038, Apr 2020. URL: <http://dx.doi.org/10.1088/1475-7516/2020/04/038>, doi:10.1088/1475-7516/2020/04/038. (cited p. 16)
- [357] Pandey, K.L., Karwal, T. and Das, S. Alleviating the H_0 and σ_8 anomalies with a decaying dark matter model. *Journal of Cosmology and Astroparticle Physics*, 2020(07):026–026, Jul 2020. URL: <http://dx.doi.org/10.1088/1475-7516/2020/07/026>, doi:10.1088/1475-7516/2020/07/026. (cited p. 94)
- [358] Pani, P. and Loeb, A. Constraining primordial black-hole bombs through spectral distortions of the cosmic microwave background. *Phys. Rev. D*, 88:041301, Aug 2013. URL: <https://link.aps.org/doi/10.1103/PhysRevD.88.041301>, doi:10.1103/PhysRevD.88.041301. (cited p. 18)
- [359] Papež, J., Grigori, L. and Stompor, R. Accelerating linear system solvers for time-domain component separation of cosmic microwave background data. *Astronomy Astrophysics*, 638:A73, Jun 2020. URL: <http://dx.doi.org/10.1051/0004-6361/202037687>, doi:10.1051/0004-6361/202037687. (cited p. 153)
- [360] Pearson, R. and Zahn, O. Cosmology from cross correlation of CMB lensing and galaxy surveys. *Physical Review D*, 89(4), Feb 2014. URL: <http://dx.doi.org/10.1103/PhysRevD.89.043516>, doi:10.1103/physrevd.89.043516. (cited p. 16)
- [361] Pearson, T.J. et al. The Anisotropy of the Microwave Background tol=3500: Mosaic Observations with the Cosmic Background Imager. *The Astrophysical Journal*, 591(2):556–574, Jul 2003. URL: <http://dx.doi.org/10.1086/375508>, doi:10.1086/375508. (cited p. 100)
- [362] Peccei, R.D. and Quinn, H.R. Constraints imposed by CP conservation in the presence of pseudoparticles. , 16(6):1791–1797, September 1977. doi:10.1103/PhysRevD.16.1791. (cited p. 17)
- [363] Peccei, R.D. and Quinn, H.R. CP conservation in the presence of pseudoparticles. , 38(25):1440–1443, June 1977. doi:10.1103/PhysRevLett.38.1440. (cited p. 17)
- [364] Peebles, P.J.E. and Yu, J.T. Primeval Adiabatic Perturbation in an Expanding Universe. , 162:815, December 1970. doi:10.1086/150713. (cited p. 51)

-
- [365] Penzias, A.A. and Wilson, R.W. A Measurement of Excess Antenna Temperature at 4080 Mc/s. , 142:419–421, July 1965. doi:10.1086/148307. (cited pp. 2, 51, and 99)
 - [366] Perlmutter, S. et al. Measurements of Ω and Λ from 42 High-Redshift Supernovae. *The Astrophysical Journal*, 517(2):565–586, jun 1999. doi:10.1086/307221. (cited pp. 3 and 18)
 - [367] Perlmutter, S. et al. Measurements of the Cosmological Parameters Ω and Λ from the First Seven Supernovae at $z \geq 0.35$. , 483(2):565–581, July 1997. arXiv:astro-ph/9608192, doi:10.1086/304265. (cited p. 93)
 - [368] Pertoldi, L. Sensitivity to the neutrino mass hierarchy with the JUNO experiment. 2015. (cited pp. viii and 15)
 - [369] Peter, P. and Uzan, J.P. *Primordial Cosmology*. Oxford Graduate Texts. Oxford University Press, 2 2013. (cited pp. ix and 41)
 - [370] Phillips, M.M. The Absolute Magnitudes of Type IA Supernovae. , 413:L105, August 1993. doi:10.1086/186970. (cited pp. 92 and 93)
 - [371] Pietranera, L. et al. Observing cosmic microwave background polarization through ice. , 376(2):645–650, April 2007. arXiv:astro-ph/0611678, doi:10.1111/j.1365-2966.2007.11464.x. (cited p. 113)
 - [372] Pietrzyński, G. et al. A distance to the Large Magellanic Cloud that is precise to one per cent. *Nature*, 567(7747):200–203, Mar 2019. URL: <http://dx.doi.org/10.1038/s41586-019-0999-4>, doi:10.1038/s41586-019-0999-4. (cited p. 94)
 - [373] Planck Collaboration et al. Planck 2018 results. VI. Cosmological parameters. , 641:A6, September 2020. arXiv:1807.06209, doi:10.1051/0004-6361/201833910. (cited pp. xi, 3, 13, 18, 19, 32, 62, 70, 73, 94, and 97)
 - [374] Planck Collaboration et al. Planck intermediate results. L. Evidence of spatial variation of the polarized thermal dust spectral energy distribution and implications for CMB B-mode analysis. , 599:A51, March 2017. arXiv:1606.07335, doi:10.1051/0004-6361/201629164. (cited p. 117)
 - [375] Planck Collaboration et al. Planck 2018 results. IX. Constraints on primordial non-Gaussianity. , 641:A9, September 2020. arXiv:1905.05697, doi:10.1051/0004-6361/201935891. (cited pp. 48, 75, and 76)
 - [376] Planck Collaboration et al. Planck 2018 results. IV. Diffuse component separation. , 641:A4, September 2020. arXiv:1807.06208, doi:10.1051/0004-6361/201833881. (cited pp. xiii and 116)
 - [377] Planck Collaboration et al. Planck 2018 results - X. Constraints on inflation. *A&A*, 641:A10, 2020. doi:10.1051/0004-6361/201833887. (cited pp. ix, ix, 46, 48, 49, and 50)

- [378] Poletti, D. et al. Making maps of cosmic microwave background polarization for B-mode studies: the POLARBEAR example. *Astronomy Astrophysics*, 600:A60, Mar 2017. URL: <http://dx.doi.org/10.1051/0004-6361/201629467>, doi:10.1051/0004-6361/201629467. (cited pp. 129, 130, and 185)
- [379] Polnarev, A.G. Polarization and Anisotropy Induced in the Microwave Background by Cosmological Gravitational Waves. , 29:607–613, December 1985. (cited p. 59)
- [380] Pontecorvo, B. Inverse beta processes and nonconservation of lepton charge. *Zh. Eksp. Teor. Fiz.*, 34:247, 1957. (cited p. 14)
- [381] Pospelov, M., Ritz, A. and Skordis, C. Pseudoscalar Perturbations and Polarization of the Cosmic Microwave Background. *Physical Review Letters*, 103(5), Jul 2009. URL: <http://dx.doi.org/10.1103/PhysRevLett.103.051302>, doi:10.1103/physrevlett.103.051302. (cited p. 81)
- [382] Preskill, J.P. Cosmological Production of Superheavy Magnetic Monopoles. *Phys. Rev. Lett.*, 43:1365–1368, Nov 1979. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.43.1365>, doi:10.1103/PhysRevLett.43.1365. (cited p. 34)
- [383] Press, W.H. et al. *Numerical Recipes 3rd Edition: The Art of Scientific Computing*. Cambridge University Press, USA, 3 edition, 2007. (cited p. 148)
- [384] Primakoff, H. Photo-Production of Neutral Mesons in Nuclear Electric Fields and the Mean Life of the Neutral Meson. *Phys. Rev.*, 81:899–899, Mar 1951. URL: <https://link.aps.org/doi/10.1103/PhysRev.81.899>, doi:10.1103/PhysRev.81.899. (cited p. 17)
- [385] Pritchard, J.R. and Loeb, A. 21 cm cosmology in the 21st century. *Reports on Progress in Physics*, 75(8):086901, Jul 2012. URL: <http://dx.doi.org/10.1088/0034-4885/75/8/086901>, doi:10.1088/0034-4885/75/8/086901. (cited p. 90)
- [386] Prunet, S. et al. Noise estimation in CMB time-streams and fast map-making. Application to the BOOMERanG98 data. *arXiv e-prints*, pages astro-ph/0101073, January 2001. arXiv:astro-ph/0101073. (cited p. 151)
- [387] Puglisi, G. et al. Iterative map-making with two-level preconditioning for polarized cosmic microwave background data sets. *Astronomy Astrophysics*, 618:A62, Oct 2018. URL: <http://dx.doi.org/10.1051/0004-6361/201832710>, doi:10.1051/0004-6361/201832710. (cited pp. 157 and 158)
- [388] Ratra, B. and Peebles, P.J.E. Cosmological consequences of a rolling homogeneous scalar field. , 37(12):3406–3427, June 1988. doi:10.1103/PhysRevD.37.3406. (cited p. 18)

- [389] Ratra, B. and Peebles, P.J.E. Inflation in an open universe. *Phys. Rev. D*, 52:1837–1894, Aug 1995. URL: <https://link.aps.org/doi/10.1103/PhysRevD.52.1837>, doi:10.1103/PhysRevD.52.1837. (cited p. 46)
- [390] Rees, M.J. and Sciama, D.W. Large-scale Density Inhomogeneities in the Universe. , 217(5128):511–516, February 1968. doi:10.1038/217511a0. (cited p. 64)
- [391] Reichardt, C.L. Observing the Epoch of Reionization with the Cosmic Microwave Background. *Astrophysics and Space Science Library*, page 227–245, 2016. URL: http://dx.doi.org/10.1007/978-3-319-21957-8_8, doi:10.1007/978-3-319-21957-8_8. (cited pp. ix, 62, and 63)
- [392] Renk, J. et al. Galileon gravity in light of ISW, CMB, BAO and H_0 data. *Journal of Cosmology and Astroparticle Physics*, 2017(10):020–020, Oct 2017. URL: <http://dx.doi.org/10.1088/1475-7516/2017/10/020>, doi:10.1088/1475-7516/2017/10/020. (cited p. 94)
- [393] Revenu, B. et al. Destriping of polarized data in a CMB mission with a circular scanning strategy. *Astron. Astrophys. Suppl. Ser.*, 142(3):499–509, 2000. doi:10.1051/aas:2000308. (cited p. 134)
- [394] Riess, A.G. et al. Cosmic Distances Calibrated to 1% Precision with Gaia EDR3 Parallaxes and Hubble Space Telescope Photometry of 75 Milky Way Cepheids Confirm Tension with Λ CDM. *The Astrophysical Journal*, 908(1):L6, Feb 2021. URL: <http://dx.doi.org/10.3847/2041-8213/abdbaf>, doi:10.3847/2041-8213/abdbaf. (cited p. 94)
- [395] Riess, A.G. et al. Large Magellanic Cloud Cepheid Standards Provide a 1% Foundation for the Determination of the Hubble Constant and Stronger Evidence for Physics beyond Λ CDM. *The Astrophysical Journal*, 876(1):85, May 2019. URL: <http://dx.doi.org/10.3847/1538-4357/ab1422>, doi:10.3847/1538-4357/ab1422. (cited pp. xi, 94, and 97)
- [396] Riess, A.G. et al. Observational Evidence from Supernovae for an Accelerating Universe and a Cosmological Constant. *The Astronomical Journal*, 116(3):1009–1038, sep 1998. doi:10.1086/300499. (cited pp. 3 and 18)
- [397] Riess, A.G., Press, W.H. and Kirshner, R.P. A Precise Distance Indicator: Type IA Supernova Multicolor Light-Curve Shapes. , 473:88, December 1996. arXiv:astro-ph/9604143, doi:10.1086/178129. (cited p. 93)
- [398] Rogers, K.K. et al. Spin-SILC: CMB polarization component separation with spin wavelets. , 463(3):2310–2322, December 2016. arXiv:1605.01417, doi:10.1093/mnras/stw2128. (cited p. 122)
- [399] Rubin, V.C., Ford, W. K., J. and Thonnard, N. Rotational properties of 21 SC galaxies with a large range of luminosities and radii, from NGC 4605 ($R=4\text{kpc}$) to UGC 2885 ($R=122\text{kpc}$). , 238:471–487, June 1980. doi:10.1086/158003. (cited pp. 3 and 17)

- [400] Saad, Y. et al. A Deflated Version of the Conjugate Gradient Algorithm. *SIAM Journal on Scientific Computing*, 21(5):1909–1926, 2000. arXiv:<https://doi.org/10.1137/S1064829598339761>, doi:10.1137/S1064829598339761. (cited p. 154)
- [401] Saad, Y. *Iterative Methods for Sparse Linear Systems*. Other Titles in Applied Mathematics. SIAM, second edition, 2003. URL: http://www-users.cs.umn.edu/~saad/IterMethBook_2ndEd.pdf, doi:10.1137/1.9780898718003. (cited pp. 153, 157, and 158)
- [402] Sachs, R.K. and Wolfe, A.M. Perturbations of a Cosmological Model and Angular Variations of the Microwave Background. , 147:73, January 1967. doi:10.1086/148982. (cited pp. 51 and 64)
- [403] Salatino, M. et al. Studies of systematic uncertainties for Simons Observatory: polarization modulator related effects. In Zmuidzinas, J. and Gao, J.R., editors, *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy IX*, volume 10708, pages 769 – 790. International Society for Optics and Photonics, SPIE, 2018. URL: <https://doi.org/10.1117/12.2312993>. (cited pp. xiii, 114, 115, and 196)
- [404] Saleem, M. et al. The Science Case for LIGO-India, 2021. arXiv:2105.01716. (cited p. 97)
- [405] Sazonov, S.Y. and Sunyaev, R.A. Microwave polarization in the direction of galaxy clusters induced by the CMB quadrupole anisotropy. *Monthly Notices of the Royal Astronomical Society*, 310(3):765–772, 12 1999. arXiv:<https://academic.oup.com/mnras/article-pdf/310/3/765/3134788/310-3-765.pdf>, doi:10.1046/j.1365-8711.1999.02981.x. (cited p. 63)
- [406] Schaffer, K.K. et al. THE FIRST PUBLIC RELEASE OF SOUTH POLE TELESCOPE DATA: MAPS OF A 95 deg² FIELD FROM 2008 OBSERVATIONS. *The Astrophysical Journal*, 743(1):90, Nov 2011. URL: <http://dx.doi.org/10.1088/0004-637X/743/1/90>, doi:10.1088/0004-637x/743/1/90. (cited p. 121)
- [407] Schutz, B.F. Determining the Hubble constant from gravitational wave observations. , 323(6086):310–311, September 1986. doi:10.1038/323310a0. (cited p. 97)
- [408] Scott, P.F. et al. First results from the Very Small Array – III. The cosmic microwave background power spectrum. *Monthly Notices of the Royal Astronomical Society*, 341(4):1076–1083, Jun 2003. URL: <http://dx.doi.org/10.1046/j.1365-8711.2003.06354.x>, doi:10.1046/j.1365-8711.2003.06354.x. (cited p. 100)
- [409] Seljak, U. and Zaldarriaga, M. A Line-of-Sight Integration Approach to Cosmic Microwave Background Anisotropies. *The Astrophysical Journal*, 469:437, Oct 1996. URL: <http://dx.doi.org/10.1086/177793>, doi:10.1086/177793. (cited p. 59)

-
- [410] Seljak, U. and Hirata, C.M. Gravitational lensing as a contaminant of the gravity wave signal in the CMB. *Physical Review D*, 69(4), Feb 2004. URL: <http://dx.doi.org/10.1103/PhysRevD.69.043005>, doi:10.1103/physrevd.69.043005. (cited pp. 118 and 119)
- [411] Senatore, L., Smith, K.M. and Zaldarriaga, M. Non-Gaussianities in single field inflation and their optimal limits from the WMAP 5-year data. *Journal of Cosmology and Astroparticle Physics*, 2010(01):028–028, Jan 2010. URL: <http://dx.doi.org/10.1088/1475-7516/2010/01/028>, doi:10.1088/1475-7516/2010/01/028. (cited p. 75)
- [412] Shapley, H. and Curtis, H.D. The Scale of the Universe. *Bulletin of the National Research Council*, 2(11):171–217, May 1921. (cited p. 1)
- [413] Sheehy, C. and Slosar, A. No evidence for dust B-mode decorrelation in Planck data. *Physical Review D*, 97(4), Feb 2018. URL: <http://dx.doi.org/10.1103/PhysRevD.97.043522>, doi:10.1103/physrevd.97.043522. (cited p. 117)
- [414] Sherwin, B.D. and Schmittfull, M. Delensing the CMB with the cosmic infrared background. *Physical Review D*, 92(4), Aug 2015. URL: <http://dx.doi.org/10.1103/PhysRevD.92.043005>, doi:10.1103/physrevd.92.043005. (cited p. 118)
- [415] Sikivie, P. Experimental Tests of the "Invisible" Axion. *Phys. Rev. Lett.*, 51:1415–1417, Oct 1983. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.51.1415>, doi:10.1103/PhysRevLett.51.1415. (cited p. 17)
- [416] Silk, J. Fluctuations in the Primordial Fireball. , 215(5106):1155–1156, September 1967. doi:10.1038/2151155a0. (cited p. 51)
- [417] Silk, J. Cosmic Black-Body Radiation and Galaxy Formation. , 151:459, February 1968. doi:10.1086/149449. (cited p. 51)
- [418] Silverstein, E. and Tong, D. Scalar speed limits and cosmology: Acceleration from D-cceleration. *Physical Review D*, 70(10), Nov 2004. URL: <http://dx.doi.org/10.1103/PhysRevD.70.103505>, doi:10.1103/physrevd.70.103505. (cited p. 76)
- [419] Slosar, A. et al. Dark Energy and Modified Gravity, 2019. arXiv:1903.12016. (cited p. 19)
- [420] Smith, K.M. et al. Delensing CMB polarization with external datasets. *Journal of Cosmology and Astroparticle Physics*, 2012(06):014–014, Jun 2012. URL: <http://dx.doi.org/10.1088/1475-7516/2012/06/014>, doi:10.1088/1475-7516/2012/06/014. (cited p. 118)
- [421] Smith, K.M., Zahn, O. and Doré, O. Detection of gravitational lensing in the cosmic microwave background. *Physical Review D*, 76(4), Aug 2007. URL: <http://dx.doi.org/10.1103/PhysRevD.76.043510>, doi:10.1103/physrevd.76.043510. (cited p. 100)

- [422] Smith, R.E. et al. Stable clustering, the halo model and non-linear cosmological power spectra. *Monthly Notices of the Royal Astronomical Society*, 341(4):1311–1332, 06 2003. arXiv:<https://academic.oup.com/mnras/article-pdf/341/4/1311/3500043/341-4-1311.pdf>, doi:10.1046/j.1365-8711.2003.06503.x. (cited pp. x and 67)
- [423] Smoot, G.F. et al. Structure in the COBE Differential Microwave Radiometer First-Year Maps. , 396:L1, September 1992. doi:10.1086/186504. (cited pp. 52 and 99)
- [424] Stairs, I.H. et al. Measurement of Relativistic Orbital Decay in the PSR B153412 Binary System. *The Astrophysical Journal*, 505(1):352–357, sep 1998. doi:10.1086/306151. (cited p. 95)
- [425] Staniszewski, Z. et al. The Keck Array: A Multi Camera CMB Polarimeter at the South Pole. *Journal of Low Temperature Physics*, 167(5-6):827–833, June 2012. doi:10.1007/s10909-012-0510-1. (cited p. 102)
- [426] Starobinsky, A.A. Dynamics of phase transition in the new inflationary universe scenario and generation of perturbations. *Physics Letters B*, 117(3-4):175–178, November 1982. doi:10.1016/0370-2693(82)90541-X. (cited p. 34)
- [427] Stebor, N. et al. The Simons Array CMB polarization experiment. In Holland, W.S. and Zmuidzinas, J., editors, *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy VIII*, volume 9914 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 99141H, July 2016. doi:10.1117/12.2233103. (cited p. 106)
- [428] Steinhardt, P.J. and Turok, N. The cyclic model simplified. *New Astronomy Reviews*, 49(2-6):43–57, May 2005. URL: <http://dx.doi.org/10.1016/j.newar.2005.01.003>, doi:10.1016/j.newar.2005.01.003. (cited p. 31)
- [429] Stompor, R. et al. Making maps of the cosmic microwave background: The MAXIMA example. *Physical Review D*, 65(2), Dec 2001. URL: <http://dx.doi.org/10.1103/PhysRevD.65.022003>, doi:10.1103/physrevd.65.022003. (cited pp. 129, 134, and 151)
- [430] Stompor, R. et al. Maximum likelihood algorithm for parametric component separation in cosmic microwave background experiments. *Monthly Notices of the Royal Astronomical Society*, 392(1):216–232, Jan 2009. URL: <http://dx.doi.org/10.1111/j.1365-2966.2008.14023.x>, doi:10.1111/j.1365-2966.2008.14023.x. (cited p. 122)
- [431] Strigari, L.E. Galactic searches for dark matter. *Physics Reports*, 531(1):1–88, Oct 2013. URL: <http://dx.doi.org/10.1016/j.physrep.2013.05.004>, doi:10.1016/j.physrep.2013.05.004. (cited p. 17)
- [432] Sullivan, M. et al. The dependence of Type Ia Supernovae luminosities on their host galaxies. *Monthly Notices of the Royal Astronomical Society*, page no–no, May 2010. URL: <http://dx.doi.org/10.1111/>

- j.1365-2966.2010.16731.x, doi:10.1111/j.1365-2966.2010.16731.x. (cited p. 93)
- [433] Sunyaev, R.A. and Zeldovich, Y.B. The Observations of Relic Radiation as a Test of the Nature of X-Ray Radiation from the Clusters of Galaxies. *Comments on Astrophysics and Space Physics*, 4:173, November 1972. (cited p. 63)
- [434] Sutton, D. et al. Map making in small field modulated CMB polarization experiments: approximating the maximum likelihood method. *Monthly Notices of the Royal Astronomical Society*, 393(3):894–910, 02 2009. arXiv:<https://academic.oup.com/mnras/article-pdf/393/3/894/18569414/mnras0393-0894.pdf>, doi:10.1111/j.1365-2966.2008.14195.x. (cited pp. 130 and 182)
- [435] Sutton, D. et al. Fast and precise map-making for massively multi-detector CMB experiments. *Monthly Notices of the Royal Astronomical Society*, 407(3):1387–1402, Aug 2010. URL: <http://dx.doi.org/10.1111/j.1365-2966.2010.16954.x>, doi:10.1111/j.1365-2966.2010.16954.x. (cited p. 130)
- [436] Suzuki, A. et al. The Polarbear-2 and the Simons Array Experiments. *Journal of Low Temperature Physics*, 184(3-4):805–810, Jan 2016. URL: <http://dx.doi.org/10.1007/s10909-015-1425-4>, doi:10.1007/s10909-015-1425-4. (cited p. 106)
- [437] Svrcek, P. and Witten, E. Axions in string theory. *Journal of High Energy Physics*, 2006(06):051–051, Jun 2006. URL: <http://dx.doi.org/10.1088/1126-6708/2006/06/051>, doi:10.1088/1126-6708/2006/06/051. (cited p. 17)
- [438] Szydlarski, M., Grigori, L. and Stompor, R. Accelerating the cosmic microwave background map-making procedure through preconditioning. *Astronomy Astrophysics*, 572:A39, Nov 2014. URL: <http://dx.doi.org/10.1051/0004-6361/201323210>, doi:10.1051/0004-6361/201323210. (cited pp. 133, 153, 155, 157, 158, 167, and 188)
- [439] Takakura, S. et al. Measurements of Tropospheric Ice Clouds with a Ground-based CMB Polarization Experiment, POLARBEAR. *The Astrophysical Journal*, 870(2):102, Jan 2019. URL: <http://dx.doi.org/10.3847/1538-4357/aaf381>, doi:10.3847/1538-4357/aaf381. (cited p. 113)
- [440] Takakura, S. et al. Performance of a continuously rotating half-wave plate on the POLARBEAR telescope. *Journal of Cosmology and Astroparticle Physics*, 2017(05):008–008, May 2017. URL: <http://dx.doi.org/10.1088/1475-7516/2017/05/008>, doi:10.1088/1475-7516/2017/05/008. (cited pp. 103 and 185)

- [441] Tang, J. et al. Comparison of Two-Level Preconditioners Derived from Deflation, Domain Decomposition and Multigrid Methods. *Journal of Scientific Computing*, 39 (3), 2009, 39, 06 2009. doi:10.1007/s10915-009-9272-6. (cited p. 155)
- [442] Tanimura, H. et al. A search for warm/hot gas filaments between pairs of SDSS Luminous Red Galaxies. *Monthly Notices of the Royal Astronomical Society*, 483(1):223–234, Nov 2018. URL: <http://dx.doi.org/10.1093/mnras/sty3118>, doi:10.1093/mnras/sty3118. (cited p. 13)
- [443] Tashiro, H. CMB spectral distortions and energy release in the early universe. *Progress of Theoretical and Experimental Physics*, 2014(6), 06 2014. 06B107. arXiv:<https://academic.oup.com/ptep/article-pdf/2014/6/06B107/4448175/ptu066.pdf>, doi:10.1093/ptep/ptu066. (cited p. 77)
- [444] Taylor, J.H. and Weisberg, J.M. A new test of general relativity - Gravitational radiation and the binary pulsar PSR 1913+16. , 253:908–920, February 1982. doi:10.1086/159690. (cited p. 95)
- [445] Tegmark, M. How to Make Maps from Cosmic Microwave Background Data without Losing Information. *The Astrophysical Journal*, 480(2):L87–L90, May 1997. URL: <http://dx.doi.org/10.1086/310631>, doi:10.1086/310631. (cited p. 120)
- [446] Tegmark, M. and de Oliveira-Costa, A. How to measure CMB polarization power spectra without losing information. *Physical Review D*, 64(6), Aug 2001. URL: <http://dx.doi.org/10.1103/PhysRevD.64.063001>, doi:10.1103/physrevd.64.063001. (cited p. 123)
- [447] Tegmark, M., de Oliveira-Costa, A. and Hamilton, A.J. High resolution foreground cleaned CMB map from WMAP. , 68(12):123523, December 2003. arXiv:astro-ph/0302496, doi:10.1103/PhysRevD.68.123523. (cited p. 122)
- [448] The Simons Observatory Collaboration et al. The Simons Observatory: Astro2020 Decadal Project Whitepaper, 2019. arXiv:1907.08284. (cited p. 109)
- [449] Thorne, B. et al. PySM: Python Sky Model, April 2017. arXiv:1704.007. (cited p. 177)
- [450] Thornton, R.J. et al. THE ATACAMA COSMOLOGY TELESCOPE: THE POLARIZATION-SENSITIVE ACTPol INSTRUMENT. *The Astrophysical Journal Supplement Series*, 227(2):21, Dec 2016. URL: <http://dx.doi.org/10.3847/1538-4365/227/2/21>, doi:10.3847/1538-4365/227/2/21. (cited pp. 101 and 161)
- [451] Tiburzi, C. Pulsars Probe the Low-Frequency Gravitational Sky: Pulsar Timing Arrays Basics and Recent Results. *Publications of the Astronomical Society of Australia*, 35, 2018. URL: <http://dx.doi.org/10.1017/pasa.2018.7>, doi:10.1017/pasa.2018.7. (cited p. 96)

-
- [452] Tran, H. et al. Comparison of the crossed and the Gregorian Mizuguchi-Dragone for wide-field millimeter-wave astronomy. , 47(2):103–109, January 2008. doi:10.1364/AO.47.000103. (cited p. 103)
- [453] Tripp, R. A two-parameter luminosity correction for Type IA supernovae. , 331:815–820, March 1998. (cited p. 93)
- [454] Tristram, M. et al. Planck constraints on the tensor-to-scalar ratio. *Astronomy Astrophysics*, 647:A128, Mar 2021. URL: <http://dx.doi.org/10.1051/0004-6361/202039585>, doi:10.1051/0004-6361/202039585. (cited pp. 48, 70, 74, and 102)
- [455] Troitsky, A.V. and Osharin, A.M. Polarization of thermal microwave radiation of the cloudy atmosphere. *Radiophysics and Quantum Electronics*, 43(5):356–365, May 2000. doi:10.1007/BF02677151. (cited p. 113)
- [456] Troxel, M. and Ishak, M. The intrinsic alignment of galaxies and its impact on weak gravitational lensing in an era of precision cosmology. *Physics Reports*, 558:1–59, Feb 2015. URL: <http://dx.doi.org/10.1016/j.physrep.2014.11.001>, doi:10.1016/j.physrep.2014.11.001. (cited p. 88)
- [457] Ueda, S. et al. A Cool Core Disturbed: Observational Evidence for the Coexistence of Subsonic Sloshing Gas and Stripped Shock-heated Gas around the Core of RX J1347.5–1145. *The Astrophysical Journal*, 866(1):48, Oct 2018. URL: <http://dx.doi.org/10.3847/1538-4357/aadd9d>, doi:10.3847/1538-4357/aadd9d. (cited pp. x and 80)
- [458] Velten, H.E.S., vom Marttens, R.F. and Zimdahl, W. Aspects of the cosmological “coincidence problem”. *The European Physical Journal C*, 74(11), Nov 2014. URL: <http://dx.doi.org/10.1140/epjc/s10052-014-3160-4>, doi:10.1140/epjc/s10052-014-3160-4. (cited p. 30)
- [459] Vergès, C., Errard, J. and Stompor, R. Framework for analysis of next generation, polarized CMB data sets in the presence of Galactic foregrounds and systematic effects. *Physical Review D*, 103(6), Mar 2021. URL: <http://dx.doi.org/10.1103/PhysRevD.103.063507>, doi:10.1103/physrevd.103.063507. (cited pp. 132, 196, and 198)
- [460] Vernizzi, F. and Wands, D. Non-Gaussianities in two-field inflation. *Journal of Cosmology and Astroparticle Physics*, 2006(05):019–019, May 2006. URL: <http://dx.doi.org/10.1088/1475-7516/2006/05/019>, doi:10.1088/1475-7516/2006/05/019. (cited p. 75)
- [461] Walker, E.S. et al. Supernova Legacy Survey: using spectral signatures to improve Type Ia supernovae as distance indicators. , 410(2):1262–1282, January 2011. arXiv:1008.2308, doi:10.1111/j.1365-2966.2010.17519.x. (cited p. 93)

- [462] Wandelt, B.D., Hivon, E. and Górski, K.M. Cosmic microwave background anisotropy power spectrum statistics for high precision cosmology. *Physical Review D*, 64(8), Sep 2001. URL: <http://dx.doi.org/10.1103/PhysRevD.64.083003>, doi:10.1103/physrevd.64.083003. (cited p. 124)
- [463] Wandelt, B.D., Larson, D.L. and Lakshminarayanan, A. Global, exact cosmic microwave background data analysis using Gibbs sampling. *Physical Review D*, 70(8), Oct 2004. URL: <http://dx.doi.org/10.1103/PhysRevD.70.083511>, doi:10.1103/physrevd.70.083511. (cited p. 124)
- [464] Waterhouse, T.P. and Zibin, J.P. The cosmic variance of Omega, 2008. [arXiv:0804.1771](https://arxiv.org/abs/0804.1771). (cited p. 46)
- [465] Wehus, I.K., Naess, S.K. and Eriksen, H.K. BAYESIAN NOISE ESTIMATION FOR NON-IDEAL COSMIC MICROWAVE BACKGROUND EXPERIMENTS. *The Astrophysical Journal Supplement Series*, 199(1):15, Feb 2012. URL: <http://dx.doi.org/10.1088/0067-0049/199/1/15>, doi:10.1088/0067-0049/199/1/15. (cited pp. 134 and 151)
- [466] Weinberg, D.H. The Lyman- α Forest as a Cosmological Tool. *AIP Conference Proceedings*, 2003. URL: <http://dx.doi.org/10.1063/1.1581786>, doi:10.1063/1.1581786. (cited p. 90)
- [467] Weinberg, D.H. et al. Observational probes of cosmic acceleration. *Physics Reports*, 530(2):87–255, Sep 2013. URL: <http://dx.doi.org/10.1016/j.physrep.2013.05.001>, doi:10.1016/j.physrep.2013.05.001. (cited p. 88)
- [468] Weinberg, S. The cosmological constant problem. *Reviews of Modern Physics*, 61(1):1–23, January 1989. doi:10.1103/RevModPhys.61.1. (cited p. 12)
- [469] Weinberg, S. Goldstone Bosons as Fractional Cosmic Neutrinos. *Physical Review Letters*, 110(24), Jun 2013. URL: <http://dx.doi.org/10.1103/PhysRevLett.110.241301>, doi:10.1103/physrevlett.110.241301. (cited p. 16)
- [470] Weisberg, J.M., Taylor, J.H. and Fowler, L.A. Gravitational waves from an orbiting pulsar. *Scientific American*, 245:74–82, October 1981. doi:10.1038/scientificamerican1081-74. (cited p. 95)
- [471] Westbrook, B. et al. Detector fabrication development for the LiteBIRD satellite mission, 2021. [arXiv:2101.05306](https://arxiv.org/abs/2101.05306). (cited pp. xii and 111)
- [472] Will, C.M. Solar system versus gravitational-wave bounds on the graviton mass. *Classical and Quantum Gravity*, 35(17):17LT01, Jul 2018. URL: <http://dx.doi.org/10.1088/1361-6382/aad13c>, doi:10.1088/1361-6382/aad13c. (cited p. 18)

- [473] Wong, K.C. et al. H0LiCOW – XIII. A 2.4 per cent measurement of H0 from lensed quasars: 5.3σ tension between early- and late-Universe probes. *Monthly Notices of the Royal Astronomical Society*, 498(1):1420–1439, 09 2019. arXiv:<https://academic.oup.com/mnras/article-pdf/498/1/1420/33755111/stz3094.pdf>, doi:10.1093/mnras/stz3094. (cited pp. 3 and 94)
- [474] Wouthuysen, S.A. On the excitation mechanism of the 21-cm (radio-frequency) interstellar hydrogen emission line. , 57:31–32, January 1952. doi:10.1086/106661. (cited p. 92)
- [475] Yamamoto, K., Sasaki, M. and Tanaka, T. Large-Angle Cosmic Microwave Background Anisotropy in an Open Universe in the One-Bubble Inflationary Scenario. *The Astrophysical Journal*, 455:412, Dec 1995. URL: <http://dx.doi.org/10.1086/176588>, doi:10.1086/176588. (cited p. 46)
- [476] Young, S. and Byrnes, C.T. Signatures of non-gaussianity in the isocurvature modes of primordial black hole dark matter. *Journal of Cosmology and Astroparticle Physics*, 2015(04):034–034, Apr 2015. URL: <http://dx.doi.org/10.1088/1475-7516/2015/04/034>, doi:10.1088/1475-7516/2015/04/034. (cited p. 18)
- [477] Yvon, D. and Mayet, F. Mirage: A new iterative map-making code for CMB experiments. *Astronomy Astrophysics*, 436(2):729–739, May 2005. URL: <http://dx.doi.org/10.1051/0004-6361:20035920>, doi:10.1051/0004-6361:20035920. (cited p. 153)
- [478] Zaldarriaga, M. and Seljak, U. All-sky analysis of polarization in the microwave background. *Physical Review D*, 55(4):1830–1840, Feb 1997. URL: <http://dx.doi.org/10.1103/PhysRevD.55.1830>, doi:10.1103/physrevd.55.1830. (cited pp. 55, 56, and 59)
- [479] Zeldovich, Y.B. A hypothesis, unifying the structure and the entropy of the Universe. , 160:1P, January 1972. doi:10.1093/mnras/160.1.1P. (cited p. 51)
- [480] Zeldovich, Y. and Khlopov, M. On the concentration of relic magnetic monopoles in the universe. *Physics Letters B*, 79(3):239–241, 1978. URL: <https://www.sciencedirect.com/science/article/pii/0370269378902320>, doi:[https://doi.org/10.1016/0370-2693\(78\)90232-0](https://doi.org/10.1016/0370-2693(78)90232-0). (cited p. 34)
- [481] Zwicky, F. Die Rotverschiebung von extragalaktischen Nebeln. *Helvetica Physica Acta*, 6:110–127, January 1933. (cited pp. 3 and 17)