

IMPERIAL

Forecasted Experimental Sensitivities to the Cosmic Microwave and Stochastic Gravitational Wave Backgrounds

Eleanor Jane Anne Gleave

June 2025

Department of Physics

Imperial College London

under the supervision of Professor Andrew Jaffe

Submitted in part fulfilment of the requirements for the degree of
Doctor of Physics of Imperial College London
and the Diploma of Imperial College London

Abstract

Gravitational wave (GW) observations probe a diffuse, stochastic gravitational wave background (SGWB) in addition to individual cataclysmic events, such as the mergers of two compact objects. The detection and description of the SGWB requires somewhat different techniques to those required for individual events, for both direct and indirect observations.

In this thesis, I first probe the sensitivity of present and future GW experiments to different background sources, including those expected from unresolved compact binaries in their quasi-Newtonian inspiral and eventual mergers, as well as more speculative cosmological sources such as inflation, cosmic strings, and first-order phase transitions, over frequencies ranges where those sources can be described by a single power law. I develop a Fisher matrix formalism to forecast the upcoming sensitivities of individual, and (combinations of) multiple experiments, and novel visualizations taking into account the increase in sensitivity to the SGWB with time.

I subsequently develop this methodology to probe the sensitivities of these experiments to deviations from a single power law, aiming to ultimately evaluate sensitivities to more complicated spectral shapes (which are a more accurate fit to cosmological spectra).

Then, I present the work I have done for the Simons Observatory (SO), a cosmic microwave background (CMB) experiment currently in initial science observations (ISO). Precise measurements of the CMB are a method for indirect observations of the SGWB—measuring the primordial GWB (through effects on the CMB temperature and polarization power spectra) constrains the overall SGWB.

I present the two campaigns of work I carried out, on the optimization of the scanning strategies for the Large and Small Aperture Telescopes, presenting new baseline strategies for each. These strategies are being tested during ISO. There is particular discussion of the implications of the proposed scanning strategies for the sensitivity of SO to the primordial GWB.

Declarations

Statement of Originality

The work contained in this thesis was carried out during my doctoral studies at Imperial College London, under the supervision of Professor Andrew Jaffe. Except where explicitly stated in the text or below, all of the work presented herein is my own.

During my candidacy, I have carried out work as a member of the Simons Observatory (SO) collaboration. As is typical for large scientific collaborations, some of the work presented in this thesis was conducted jointly with other members of SO.

The work done in Chapter 8 was carried out in collaboration with Sara Simon, Reijo Keskitalo, Zack Atkins, Adriaan Duivenvoorden, Frank Qu, and Zack Li. The work in Chapter 9 was done in collaboration with Sara Simon, Ken Ganga, Reijo Keskitalo, Aashrita Mangu, David Alonso, and Akito Kusaka. In both projects I made substantial and original contributions. Details of my specific contributions are clearly indicated in the respective chapters.

Some figures in this thesis have been reproduced or adapted from previously published work, or from unpublished material produced by the SO collaboration. In all such cases, appropriate credit is provided in the figure captions.

London, United Kingdom, June 2025

Eleanor Jane Anne Gleave

Copyright Statement

The copyright of this thesis rests with the author. Unless otherwise indicated, its contents are licensed under a Creative Commons Attribution-Non Commercial 4.0 International Licence

(CC BY-NC). Under this licence, you may copy and redistribute the material in any medium or format. You may also create and distribute modified versions of the work. This is on the condition that you credit the author and do not use it, or any derivative works, for a commercial purpose. When reusing or sharing this work, ensure you make the licence terms clear to others by naming the licence and linking to the licence text. Where a work has been adapted, you should indicate that the work has been changed and describe those changes. Please seek permission from the copyright holder for uses of this work that are not included in this licence or permitted under UK Copyright Law.

Acknowledgements

I would like to begin by thanking my supervisor, Andrew, for his guidance, patience, and steady encouragement throughout my PhD. His insight and advice have been invaluable in shaping both this thesis and my development as a researcher. I am also grateful to Alan, who, alongside Andrew, fought to secure an increase to my funding—a change that made a real difference to my ability to complete this work. I am also grateful to my examiners, Alberto and Carlo, for their careful reading and constructive feedback.

I feel fortunate to have been part of such supportive and inspiring communities at Imperial and within the Simons Observatory. I owe particular thanks to Sara, who has been almost like a second supervisor: her generosity with her time, advice, and encouragement has made a lasting difference. I am also thankful to Kimeel for his invaluable computing help, which rescued me on more than one occasion.

Teaching has been one of the most rewarding aspects of this journey. When I doubted whether I was on the right path, it was often the spontaneous questions from students that reminded me why I love what I do. Their genuine curiosity helped me see the subject in new ways and gave me reassurance during moments of imposter syndrome. I am grateful to all the students I have taught for keeping my love of cosmology alive.

To my friends: Zoe, who saw me through the whole process, including me slowly losing my sanity while trying to start a PhD during Covid, and who has remained a steadfast friend — thank you. And to Cathy and Emma, my best friends, whose humour, perspective, and unwavering encouragement have meant so much. Our crossword clubs have become a cherished tradition, giving me the chance to pause, laugh, and feel grounded when I need it most, and reminding me how lucky I am to have you both in my life.

To my parents, thank you for your belief in me and your support, each in your own way. I am especially grateful to my father, whose own time in academia was shorter than he would have liked, but who ensured that I could follow my path without the same limitations.

To my partner, William, thank you for your love, patience, and constant good humour

throughout this journey, especially during the final stages of writing. From making endless cups of tea to proofreading my thesis, and countless other quiet acts of kindness, you enabled me to focus fully on the work—always without complaint. I look forward to returning the favour as you begin writing your own thesis.

And a final thank you to my dog, Mollie—asleep on my lap as I write this—whose patience and companionship have grounded me over the years. It was on long walks together that I found clarity in both my work and in life, and in your quiet presence I am reminded, always, of what matters most.

Contents

I	Cosmology Background	21
1	The Standard Cosmological Model	22
1.1	Introduction	22
1.2	General Relativity	23
1.3	Gravitational Waves	24
1.4	The Friedmann-Lemaître-Robertson-Walker Metric	25
1.5	Expansion of the Universe	26
1.6	Dark Matter	27
1.7	Dark Energy	28
1.8	Thermal History of the Universe	28
1.9	The Stochastic Gravitational Wave Background	33
1.10	Summary and Discussion	33
2	The Cosmic Microwave Background	35
2.1	Introduction	35
2.2	The Surface of Last Scattering	36
2.3	Temperature Anisotropies	36
2.4	Polarization	40
2.5	Analysis of CMB Anisotropies	43
2.6	Observational Techniques and Data Reduction	46
2.7	History of CMB Observations and Results	48
2.8	Summary and Discussion	50
3	The Simons Observatory	51
3.1	Introduction	51
3.2	The Telescopes	52
3.3	Science Goals of SO	55

3.4	Summary and Discussion	60
4	The Stochastic Gravitational Wave Background	62
4.1	Introduction	62
4.2	Astrophysical Foregrounds	63
4.3	Cosmological Sources	65
4.4	Summary and Discussion	72
5	Observing the SGWB	73
5.1	Introduction	73
5.2	Gravitational Wave Experiments	74
5.3	History of SGWB Observations and Results	79
5.4	Summary and Discussion	82
II	Forecasting Experimental Sensitivities to the SGWB	83
6	Modelling the SGWB	84
6.1	Introduction	84
6.2	Instrumental Noise	84
6.3	Astrophysical Foregrounds—Binary Systems	85
6.4	Fisher Matrices and Sensitivities to the SGWB	86
6.5	Power Law Approximations to Cosmological Spectra	89
6.6	Power Law Sensitivities	92
6.7	Combining Multiple Experiments	93
6.8	Adding an Observational Prior	94
6.9	Summary and Discussion	95
7	Forecasting Direct Sensitivities to the SGWB	96
7.1	Introduction (Single Power Laws)	96
7.2	Results (Single Power Laws)	97
7.3	Summary and Discussion (Single Power Laws)	106
7.4	Introduction (More Complex Spectra)	108
7.5	Results (More Complex Spectra)	108
7.6	Summary (More Complex Spectra)	111

III	Forecasting SO Sensitivities to the CMB and SGWB	112
8	Large Aperture Telescope Scan Strategy	113
8.1	Introduction	113
8.2	Proposed Strategies	114
8.3	Technical Constraints and Concerns	120
8.4	Methodology	123
8.5	Scientific Performance	132
8.6	Summary and Discussion	134
9	Small Aperture Telescope Scan Strategy	135
9.1	Introduction and Science Goals	135
9.2	Technical Considerations	136
9.3	Proposed Strategies	137
9.4	Methodology	140
9.5	Results	143
9.6	Summary and Discussion	147
IV	Final Remarks	150
10	Conclusions and Future Directions	151
10.1	Introduction	151
10.2	Key Findings and Contributions	152
10.3	Scientific Implications	153
10.4	Future Directions	153
10.5	Final Remarks	154
	Bibliography	155
A	Foreground Models for the SGWB	177
B	GW Experiment Noise Curves	179
B.1	Ground-Based Interferometers	179
B.2	Space-Based Interferometers	180
B.3	Pulsar Timing Arrays	181
C	SGWB Fiducial Models	182

D Additional LAT Scan Strategy Results **184**

D.1 Initial Simulation Run Results 184

D.2 Improved Simulation Run Results 185

E Additional SAT Scan Strategy Results **186**

E.1 Figures of Merit 187

List of Figures

2.1	<i>B</i> -mode polarization power spectra of the CMB—figure credit: BICEP/Keck Collaboration	42
2.2	Temperature and polarization power spectra of the CMB—figure credit: Planck Collaboration	43
2.3	The effects of varying cosmological parameters on the CMB temperature angular power spectrum—figure credit: Wayne Hu & Scott Dodelson	45
2.4	CMB foregrounds at different frequencies—figure credit: Planck Collaboration and ESA	47
3.1	Renderings of the LAT and LATR—figure credits: SO Collaboration	53
3.2	LAT first light image of Mars—figure credit: SO Collaboration	54
4.1	Example SGWB spectra	66
5.1	Current measurements and eventual sensitivities of GW experiments to the SGWB	75
6.1	Evolving GW experiment sensitivities with time	89
6.2	Power law functions	92
7.1	Sensitivities of individual experiments to simple power laws	98
7.2	Comparison of ‘bowtie’ plots with error bars	100
7.3	Combined sensitivities of overlapping GW experiments	102
7.4	Combined sensitivities of distinct GW experiments	104
7.5	NANOGrav 15 year results plotted as ‘bowties’	105
7.6	NANOGrav 15 year results as an observational prior for LVK-O4	106
7.7	Example results plot for fiducial model with deviation from power law	109
7.8	Sensitivities of experiments to a flat spectrum	110
7.9	Sensitivities of experiments to a curved spectrum	110
8.1	LAT scan Type 1.0 hit map	117

8.2	LAT scan Type 1.1.0 hit map	117
8.3	LAT scan Type 1.1.1 hit map	118
8.4	LAT scan Type 1.2.0 hit map	118
8.5	LAT scan Type 1.2.1 hit map	119
8.6	LAT scan Type 2 hit map	119
8.7	LAT scan Type 3 hit map	120
8.8	Cross-linking maps for LAT scan strategies	126
8.9	Cadence maps for LAT scan strategies	127
8.10	Cross-linking maps for LAT scan optimization	130
8.11	Cadence maps for LAT scan optimization	130
9.1	SAT baseline scan strategy footprints—figure credits: SO Collaboration	138
9.2	SAT SW-SE scan strategy hit map	138
9.3	SAT SW-N-SE scan strategy hit map	139
9.4	SAT SW-NE-NW-SE scan strategy hit map	139
9.5	Example transfer function—figure credit: Ken Ganga, SO Collaboration	141
9.6	New baseline scan strategy for the SAT	148
D.1	Cross-linking maps for improved LAT scan simulations—figure credit: SO Col- laboration	184
D.2	Depth and noise maps for improved LAT scan simulations—figure credit: SO Collaboration	185

List of Tables

5.1	Properties of GW experiments	76
5.2	PTA measurements of the SGWB	80
5.3	LVK upper limits on the SGWB	80
8.1	Technical constraints for the LAT scan strategy	121
8.2	f_{sky} for initial LAT scan simulation run	125
8.3	Summary of LAT scan strategy performances	126
8.4	Figures of merit used for optimization of the LAT Type 1.2.1 scan	129
9.1	f_{sky} values for different SAT scan strategies	144
9.2	Effect of elevation on SAT scan strategy performance	145
9.3	Effect of azimuth throw on SAT scan strategy performance	145
9.4	Effect of scan speed on SAT scan strategy performance	146
9.5	Effect of boresight directions on SAT scan strategy performance	146
9.6	Improvement of new SAT scan strategy on $\sigma(r)$	148
A.1	Coefficients for the foreground model used for SGWB forecasting	178
C.1	Fiducial SGWB models used for forecasting	183
D.1	Coefficients for LAT elevation noise	185
E.1	Observing efficiencies for different SAT scan strategies	187
E.2	Figures of merit for different SAT scan strategies	188

List of Acronyms

AGN active galactic nucleus

ALP axion-like particle

BAO baryon acoustic oscillation

BBN Big Bang nucleosynthesis

BH black hole

CIB cosmic infrared background

CMB cosmic microwave background

CNB cosmic neutrino background

DBPL double broken power law

Dec declination

DM dark matter

EMRI extreme mass ratio inspiral

FBOT fast blue optical transient

FLRW Friedmann-Lemaître-Robertson-Walker

FoM figure of merit

FOPT first order phase transition

FoV field of view

GR General Relativity

GRB gamma-ray burst

GW gravitational wave

GWB gravitational wave background

HWP half-wave plate

IGM intergalactic medium

ISO initial science observations

ISW integrated Sachs-Wolfe

KID kinetic inductance detector

kSZ kinematic Sunyaev-Zeldovich effect

LSS large-scale structure

ML maximum likelihood

NS neutron star

nwt ‘noise-with-time’

OT optics tube

PBH primordial black hole

PLISC power law integrated sensitivity curve

PSD power spectral density

PTA pulsar timing array

PTC pulse-tube cooler

PWV precipitable water vapour

QCD quantum chromodynamic

RA right ascension

RS Rees-Sciama

SBPL single broken power law

SFSR single field slow roll

SGWB stochastic gravitational wave background

SLS surface of last scattering

SMBH supermassive black hole

SMBHB supermassive black hole binary

SNR signal-to-noise ratio

SOBH stellar origin black hole

SWE Sachs-Wolfe effect

SZ Sunyaev-Zeldovich effect

TES transition edge sensor

tSZ thermal Sunyaev-Zeldovich

WD white dwarf

List of Collaborations and Experiments

ACT Atacama Cosmology Telescope

AdvACT Advanced ACTPol

BBO Big Bang Observer

BICEP Background Imaging of Cosmic Extragalactic Polarisation

BOOMERanG Balloon Observations Of Millimetric Extragalactic Radiation and Geophysics

CBI Cosmic Background Imager

CCAT Cerro Chajnantor Atacama Telescope

CE Cosmic Explorer

CLASS Cosmology Large Angular Scale Surveyor

CMB-S4 CMB-Stage 4

COBE Cosmic Background Explorer

CPTA Chinese Pulsar Timing Array

DASI Degree Angular Scale Interferometer

DECIGO Deci-hertz Interferometer Gravitational wave Observatory

DES Dark Energy Survey

DESI Dark Energy Spectroscopic Instrument

EPTA European Pulsar Timing Array

ESA European Space Agency

ET Einstein Telescope

FAST Five-hundred-metre Aperture Spherical radio Telescope

FYST Fred Young Submillimeter Telescope

HERA Hydrogen Epoch of Reionization Array

INDIGO Indian Initiative in Gravitational-Wave Observations

InPTA Indian Pulsar Timing Array

IPTA International Pulsar Timing Array

JAXA Japan Aerospace Exploration Agency

KAGRA Kamioka Gravitational Wave Detector

KATRIN Karlsruhe Tritium Neutrino Experiment

LIGO Laser Interferometer Gravitational-Wave Observatory

aLIGO Advanced LIGO

LVK LIGO-Virgo-KAGRA

LISA Laser Interferometer Space Antenna

LiteBIRD Lite (Light) satellite for the studies of B-mode polarization and Inflation from cosmic background Radiation Detection

LSST Large Synoptic Survey Telescope

MAXIMA Millimeter Anisotropy eXperiment IMaging Array

MPTA MeerKAT Pulsar Timing Array

NANOGrav North American Nanohertz Observatory for Gravitational Waves

NASA National Aeronautics and Space Administration

SDSS Sloan Digital Sky Survey

SO Simons Observatory

SKA Square Kilometre Array

SPT South Pole Telescope

POLARBEAR POLARization of the Background Radiation

PPTA Parkes Pulsar Timing Array

Virgo

aVirgo Advanced Virgo

WMAP Wilkinson Microwave Anisotropy Probe

Simons Observatory (SO) Acronyms

ASO Advanced Simons Observatory

SO:JP Simons Observatory: Japan

SO:UK Simons Observatory: UK

AWG analysis working group

TWG technical working group

BB low- ℓ B -modes

CSS calibration, systematics and sensitivity

LAT Large Aperture Telescope

LATR Large Aperture Telescope Receiver

SAT Small Aperture Telescope

Part I

Cosmology Background

Chapter 1

The Standard Cosmological Model

This chapter introduces the standard cosmological model, Λ CDM. Beginning with its foundational principles, we then turn to GR and the FLRW metric. GWs are introduced as a mathematical consequence of GR. We then discuss the hidden energy components of the Universe (dark matter and dark energy), including the evidence supporting them. The thermal history of the Universe is traced from its earliest stages, through to the formation of LSS. The chapter concludes with an introduction to the SGWB as a new observational window into the very early universe.

Portions of the text draw from a set of lecture notes I wrote in 2019 [1], which were in part based on a literature review I wrote the same year [2]. Other sources are cited directly in the text.

1.1 Introduction

Modern cosmology is founded on a simple but profound assumption, that the Universe is homogenous and isotropic at large scales. This ‘Cosmological Principle’ derives from the Copernican principle (i.e. that our position in the Universe is not in some way privileged).

The Cosmological Principle is observationally supported, and can be seen in, e.g., the uniformity of the cosmic microwave background (CMB) (see section 1.8.4 and Chapter 2) and large-scale structure (LSS). Adopting the inherent symmetry of the Cosmological Principle simplifies the maths required to model the Universe, leading us to the Friedmann-Lemaître-Robertson-Walker (FLRW) metric (section 1.4), which underpins our models not only of spacetime, but of an evolving Universe.

To describe the standard model of cosmology in this chapter, it is worth briefly exploring what a ‘standard model’ in science tells us, or more importantly what it does not. A standard

model is a framework of complementary theories which provide a reasonable description of the relevant science. There is no expectation that the model provides a complete or final description, or that there is any essential ‘truth’ in its constituent theories. Indeed, the Standard Model of Particle Physics is known to be incomplete (in that it does not cover dark matter or neutrino masses) and yet remains a useful, and our best, model to explain particle interactions. Rather, a standard model describes the contemporary consensus, and can (and should) be tweaked, updated, or completely re-written as and when it no longer provides the best description of the relevant phenomena.

In cosmology specifically, the standard model—the Λ CDM model—is built from fundamental principles and observations. Along with the Cosmological Principle and the [FLRW](#) metric, the existence of cold dark matter (section [1.6](#)) and the Cosmological Constant Λ are core tenets. It is an imperfect and inconsistent model (see, e.g., discussion of the Hubble tension in section [1.5](#)), but its predictions of the behaviour of the Universe and its contents are excellent.

1.2 General Relativity

General Relativity ([GR](#)), as developed by Einstein in the early 20th Century, remains a fundamental tenet of the modern standard model of cosmology. It proposes that gravitational interactions as we observe them are a consequence of the curvature of spacetime by matter and energy.

The primary equation for [GR](#) is the Einstein field equation:

$$G_{\mu\nu} + \Lambda g_{\mu\nu} = \frac{8\pi G}{c^4} T_{\mu\nu}, \quad (1.1)$$

where Λ is the cosmological constant, G is Newton’s gravitational constant, and c is the speed of light. The three tensors $G_{\mu\nu}$, $g_{\mu\nu}$, and $T_{\mu\nu}$ represent the Einstein, spacetime metric, and stress-energy tensors, respectively.

The inclusion of the Λ term was initially one of mathematical necessity to maintain a ‘steady-state’ universe, as a counterbalance to gravitational attraction. However, its meaning has since evolved to represent dark energy (discussed further in section [1.7](#)). This acts as a salient reminder that even incredible mathematical accuracy of scientific models does not inherently confer physical meaning to the model.

Alternative theories to [GR](#) are continually being developed. There is currently no compelling observational evidence for modified theories of gravity, but cosmological experiments continue

to actively test GR against different theories [3, 4, 5, 6].

1.3 Gravitational Waves

Gravitational waves (GWs) are perturbations in spacetime caused by the accelerations of massive objects, predicted first by Einstein in 1916 as a natural consequence of GR [7]. Following indirect confirmation of their existence through the observation of the decay of a binary system's orbit in 1974 [8, 9], we obtained the first direct detection of gravitational waves by aLIGO in 2015 [10, 11].

GWs exist as perturbations in the metric tensor $g_{\mu\nu}$, against background spacetime. At small scales, we can approximate the metric with the Minkowski metric $\eta_{\mu\nu}$, such that the total metric is

$$g_{\mu\nu} = \eta_{\mu\nu} + h_{\mu\nu}, \quad (1.2)$$

with a small perturbation $|h_{\mu\nu}| \ll 1$. In a transverse-traceless gauge,

$$\partial^\mu h_{\mu\nu} = 0, \quad h^\mu{}_\mu = 0 \quad (1.3)$$

which leads to

$$\partial_\alpha \partial^\alpha h_{\mu\nu} = -16\pi G T_{\mu\nu}. \quad (1.4)$$

In a vacuum, equation 1.4 has plane wave solutions:

$$h_{\mu\nu} = A_{\mu\nu}(k^\lambda) \exp(ik_\lambda x^\lambda), \quad (1.5)$$

with wave vector k_μ , and where $A_{\mu\nu}$ describes the (constant) components of a symmetric tensor [12].

Considering a plane GW propagating along the z -axis, the GW can be described:

$$h_{\mu\nu} = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & h_+ & h_\times & 0 \\ 0 & h_\times & -h_+ & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}, \quad (1.6)$$

where h_+ and $h_\times$ give the two polarization states— h_+ polarization stretches and squeezes spacetime in the x - and y -directions, while $h_\times$ does the same at a 45° offset.

1.3.1 Gravitational Wave Signals

We can categorize gravitational wave signals into general categories [13]:

Coalescing binary systems have a characteristic ‘chirp’ signal as the two objects spiral inwards. A more dramatic signal then follows from the actual merge event. In the aLIGO/aVirgo frequency range ($\sim 10\text{--}1000$ Hz), observable binary systems consist of neutron stars (NSs) and black holes (BHs) (at mass ranges in the orders of 1 and $100 M_{\odot}$, respectively). The ‘chirp’ shape of the signal from a binary can be fitted to a template, to establish the physical properties of the system. Until the recent detection of the stochastic gravitational wave background (SGWB), these formed all of our detections [13, 14].

We also expect to observe GW signals with unknown characteristic shapes—**unmodelled bursts**—as the results of several energetic phenomena. These might include supernova events, or the transition from NS to BH. They would present simply as unexpected excess power in observation runs, and are both hard to predict [15], and hard to detect, since most processing relies on templates [16].

Finally, there is a background of unresolved low-amplitude **stochastic** gravitational wave signals, the incoherent superposition of the signals of a miscellany of energetic events through cosmic history. These include red-shifted signals of the types listed above, but also (very heavily) red-shifted signals from events from the earliest stages of the universe, as far back as inflation [13]. We label this signal the SGWB, which is discussed further in section 1.9 and Chapter 4.

1.4 The Friedmann-Lemaître-Robertson-Walker Metric

The combination of GR and the Cosmological Principle lead to only one mathematically permissible metric, the FLRW metric, which describes a homogenous and isotropic universe which evolves with time.

The FLRW metric is

$$ds^2 = dx^\mu dx^\nu g_{\mu\nu} = -c^2 dt^2 + a(t)^2 \left[\frac{dr^2}{1 - kr^2} + r^2 (d\theta^2 + \sin^2 \theta d\phi^2) \right], \quad (1.7)$$

where ds is the spacetime interval, and t is cosmic time. (r, θ, ϕ) are the comoving spherical co-ordinates—that is, co-ordinates which remain at rest with the expansion of the universe, as opposed to physical co-ordinates which do not. $a(t)$ is the scale factor, which describes the expansion of the universe. Further details of comoving co-ordinates and the scale factor are

given in section 1.5 below. The curvature of space is described by k , such that $k = 0$ describes a flat universe, $k = 1$ a closed universe, and $k = -1$ an open universe.

1.5 Expansion of the Universe

Observational evidence tells us that the Universe is expanding. The first observations of this came in the 1910–20s from Slipher [17] and Hubble [18].

The dynamics of expansion are encoded in the scale factor $a(t)$, which describes the scaling of proper distances between comoving points, such that $r_{\text{proper}} = a \times r_{\text{comoving}}$. The relationship between a and cosmological redshift z (an important observable) is given by

$$1 + z = \frac{a_0}{a(z)}, \quad (1.8)$$

where the scale factor today is commonly set to unity: $a_0 = 1$.^{*} Redshift quantifies the fractional increase in wavelength of photons due to the expansion of space, and is therefore a direct probe of expansion.

The expansion rate at a given point in the Universe’s history is characterized by the Hubble parameter,

$$H(t) \equiv \frac{\dot{a}(t)}{a(t)}. \quad (1.9)$$

The present-day value H_0 , the Hubble constant, defines the present expansion rate. Accurate measurements of H_0 can be split into two observational methods: early- (CMB) and late-time (local) measurements. Late-time measurements rely on cosmic distance ladders, and measure $H_0 \approx 73\text{--}74$ km/s/Mpc [19, 20, 21]. However, CMB measurements (e.g. from Planck [22]) find $H_0 \approx 67\text{--}68$ km/s/Mpc, with the tension between the observations now at a significance of $\sim 5\sigma$.

It is possible that the ‘Hubble tension’ will be resolved with the discovery of systematic errors in one or both methods, but thus far there is no evidence to suggest poor modelling in either. It is also possible that both methods have made accurate measurements, and that new physics is required to explain the different values. Other methods of measuring H_0 (e.g. through GWs, in combination with either electromagnetic information, or galaxy catalogues [23, 24, 25]) are beginning to be used, but do not yet have sufficient accuracy to weigh in on the tension.

The Friedmann equations, which can be derived from the Einstein field equation under the

^{*}Note that parameters with the subscript 0 denote the value of that parameter at present time, t_0 .

FLRW metric, define the evolution of expansion with time:

$$H^2 = \left(\frac{\dot{a}}{a}\right)^2 = \frac{8\pi G}{3}\rho - \frac{kc^2}{a^2} + \frac{\Lambda c^2}{3}, \quad (1.10)$$

$$\frac{\ddot{a}}{a} = \frac{4\pi G}{3}\left(\rho + \frac{3p}{c^2}\right) + \frac{\Lambda c^2}{3}, \quad (1.11)$$

where ρ denotes the total energy density, p the pressure, a the scale factor, G Newton's gravitational constant, k the curvature parameter (as above), and Λ the cosmological constant. The total energy density includes contributions from radiation, baryonic matter, dark matter, and dark energy. The energy density of each component scales differently with a (and t , by extension), leading to different epochs of the universe being dominated by different contributions.

The critical density ρ_c is defined as the energy density which gives $k = 0$:

$$\rho_c = \frac{3H^2}{8\pi G}. \quad (1.12)$$

ρ_c is a useful reference scale to express the relative contributions to the total energy density from different components; the density parameter Ω_i for component i is defined

$$\Omega_i \equiv \frac{\rho_i}{\rho_c}, \quad (1.13)$$

such that the $\Omega_{\text{total}} = \sum_i \Omega_i = 1$ for a flat, > 1 for a closed, and < 1 for an open universe.

1.6 Dark Matter

There are observed disparities between our baryonic-matter models and observations, which have led towards theories of ‘dark’ matter which only interacts gravitationally. Purely baryonic models of the Universe predict a certain energy density for baryons Ω_b , but only the total matter density Ω_m is available experimentally. We observe that $\Omega_m - \Omega_b \neq 0$, and label the remaining density Ω_{DM} .

This tension between initial theory and observation has led to physically meaningful cosmology—we consider that dark matter does physically exist. The H_0 tension discussed in section 1.5 may well be explained through a similarly physical extension to our cosmology, or it may simply resolve as an error in our current experiments. In the case of dark matter, we are confident that our observations can be trusted, aided by the consistency of the measurement of Ω_{DM} through different experimental methods.

The best evidence we have for the existence of dark matter comes from galactic behaviours which cannot be explained with observed baryonic masses alone. We can, for example, calculate the Virial mass of a galaxy and compare it with the visible baryonic mass [26, 27], or compare the expected and observed rotation curves [26, 28].

Constraints on the characteristics of dark matter, as well as its overall abundance, can also be calculated from, e.g., freeze-out and nucleosynthesis calculations [29, 30], as well as upper limits from dark matter detection experiments [31].

We can parametrize different cosmological components with an equation of state: $w = p/\rho$. A range of dark matter candidates have been proposed, from dark matter particles of various types, to massive objects, some of which are considered in the work detailed in Chapter 6. The standard cosmological model is restricted to ‘cold’ dark matter—that is, dark matter which is non-relativistic during structure formation, with the equation of state $w = 0$. There are theories which put forward ‘warm’ and ‘hot’ dark matter as alternative options, but these are not observationally favoured over the standard cold dark matter [32, 33].

1.7 Dark Energy

As well as dark matter, modern cosmology requires a second invisible component to explain the accelerating expansion of the Universe. Contributing $\sim 70\%$ of the total energy in the Universe, ‘dark energy’ acts with negative pressure, and can be described by the cosmological constant Λ in the Einstein field equation.

The equation of state for Λ is $w = -1$, which is the equation of state for dark energy used in the Λ CDM model of cosmology. However, variations on the equation of state have been put forward for testing. The most common variation is a simple first-order expansion, $w = w_0 + w_a(1 - a)$, which is not physically motivated [34, 35, 36]. There is some observational support for this parametrization of the equation of state [37, 38].

1.8 Thermal History of the Universe

The thermal history of the Universe is fundamentally connected to its expansion—as the Universe expands, its temperature decreases:

$$T(z) = T_0(1 + z), \tag{1.14}$$

where $T_0 \approx 2.725$ K is the observed temperature of the CMB today.

At different temperatures, the Universe has different properties, and goes through a number of cosmological epochs to reach its present state. This section starts at the beginning of the Universe, and traces the thermal history to present day. Unless stated otherwise, I adopt the Planck 2018 cosmological results [22] for any parameters throughout this thesis.

1.8.1 The Very Beginning

Observations of the CMB (see Chapter 2) provide evidence for a universe which has expanded from some initial singularity—the Big Bang. The Big Bang theory states that the Universe began in a hot, compressed, dense state, which would leave observable thermal radiation (the CMB). By the time we make observations today, this radiation has been redshifted due to the expansion of the Universe.

In the very first moments of the Universe—the Planck epoch, which lasts for the first Planck time (10^{-43} s)—all four fundamental forces are expected to be unified, with quantum gravitational effects dominating, but this era is beyond the scope of the standard cosmological model.

After the Planck epoch follows the Grand Unification epoch (10^{-43} s to 10^{-36} s) during which the Strong force separates from the electroweak forces [39]. This process could be the cause for the baryon asymmetry which we observe today, where significantly more matter than antimatter exists [40, 41]. It is also possibly the trigger for inflation, which occurs next.

At these very small time scales, however, nothing can be well-theorized—the precise state of the Universe at this time is largely invisible to modern observations, as inflation washes out almost all irregularities.

1.8.2 Inflation

Inflation is a period of exponential growth theorized to occur in the early Universe, as an explanation for several problems which face the standard cosmological model: the flatness, horizon, and monopole problems. The flatness problem asks why the Universe today is observed to have $\Omega_{\text{total}} \approx 0$, since any initial curvature would have increased over time, and we have no reason to predict that the Universe would have begun with very precise flatness. The horizon problem asks how areas of the CMB can be causally linked, given their distance from each other. And the monopole problems asks why we have not observed relic particles like magnetic monopoles formed in the early Universe.

Inflation solves each of these problems, by taking a small, causally-connected, region of space and expanding it so quickly that curvature and heavy relics are effectively diluted to zero. To act as a solution to the three problems, the Universe must expand by a factor of (at least) 10^{78} [42, 43]. Inflation begins at 10^{-36} s and ends somewhere between 10^{-33} s and 10^{-32} s.

Different theories of inflation have different driving forces behind this expansion, but most involve a scalar field called the inflaton. The short-lived decay of the inflaton particle heats the Universe and produces the particles we see today, and inflation ends following this decay. This occurs with a constant energy density for the inflaton field, but a strong reduction in those for radiation, matter, and curvature. Inflation thereby smooths out any pre-existing fluctuations in the Universe, and flattens spacetime curvature.

Since the expansion during inflation is predicted to be exponential, initial quantum fluctuations in the inflaton field grow to macroscopic levels. This sources scalar (i.e. density) and tensor (GW) perturbations in the matter content of the Universe, which later affect the CMB and growth of LSS.

The scalar and tensor perturbations can be described by the power spectra:

$$P_s(k) = A_s \left(\frac{k}{k_*} \right)^{n_s-1}, \quad (1.15)$$

where P_s is the scalar power spectrum, A_s is the amplitude of primordial scalar perturbations at the chosen pivot scale $k = k_*$ (typically 0.04 Mpc^{-1}), $n_s \approx 1$ is the scalar spectral index, and

$$P_t(k) = A_t \left(\frac{k}{k_*} \right)^{n_t}, \quad (1.16)$$

where P_t is the tensor power spectrum, A_t is the amplitude of tensor perturbations, and n_t is the tensor spectral index. The ratio of the two power spectra, the tensor-to-scalar ratio r , is an important observable, as P_s itself can be well constrained through the CMB, but we rely on r to evaluate P_t :

$$r(k) = \frac{P_t(k)}{P_s(k)}. \quad (1.17)$$

Different theories of inflation make different predictions for these parameters—single field slow roll (SFSR) requires the consistency relation $n_T = -r/8$. Note that $n_T < 0$ gives a red spectrum (i.e. more power at low frequencies), $n_T > 0$ a blue spectrum (i.e. more power at high frequencies), and $n_T = 0$ a scale-invariant spectrum.

1.8.3 The Early Universe

Following inflation, the Universe is dominated by photons and other high-energy particles in the form of a hot plasma. The Universe goes through several epochs:

- The **quark epoch**, where the plasma takes the form of a quark-gluon plasma. After cooling, the quarks combine to form protons and neutrons.
- The **hadron epoch**, where protons and neutrons interact through the weak force.
- **Neutrino decoupling**, where neutrinos freeze out of thermal equilibrium with the plasma as the Universe expands and cools. At this point, weak interaction rates fall below the expansion rate, neutrinos therefore cease interacting with other particles, and free-stream across the Universe. This forms the cosmic neutrino background (CNB).
- The **lepton epoch**, where electron-positron pairs annihilate. The resulting injection of energy to the photon-baryon plasma gives the subsequent CMB a temperature boost over the CNB.
- **Big Bang nucleosynthesis (BBN)** follows the freeze-out of weak interactions and produces small numbers of nuclei, with relative abundances dependent on the baryon-to-photon ratio. This is predominantly helium, but atoms up to boron are created. Our measurements of these nuclei in the Universe today, particularly the helium-to-hydrogen ratio, are generally representative of early abundance ratios.[†]

In these epochs, pressure and temperature remain too high for many nuclei to form—such high energy particles move too quickly for the Strong force, with its very low interaction distance, to draw them together. As such, the Universe is almost entirely a charged plasma of photons, protons and electrons.

The abundance of free electrons leads to Thomson and Compton scattering—the elastic and inelastic collisions between particles and photons, respectively. Photons therefore have a very small mean free path and are essentially locked into local regions of the plasma. This causes the plasma to be ‘opaque’. In the absence of atoms, photons are absorbed and emitted over the whole frequency spectrum, which produces blackbody radiation.

The initial quantum fluctuations persist through to this era in the form of spatial variations in the matter and dark matter distributions.

[†]Except for lithium which we observe at a much lower abundance than BBN models predict [44]—the ‘lithium problem’.

1.8.4 Recombination

This era continues until around 378,000 years after the Big Bang, when the temperature has fallen enough for hydrogen (and a small number of helium and deuterium) nuclei to combined with electrons to form atoms—‘recombination’. With almost all electrons now bound in atoms, fewer photons are scattered and the mean free path increases dramatically, allowing photons to travel much further—‘photon decoupling’. The photons from this era form the surface of last scattering ([SLS](#)),[‡] which is as far back as electromagnetic information about the Universe can be retained.

1.8.5 Reionization and Structure Formation

Following recombination, hundreds of millions of years follow where no new light is produced—the ‘dark ages’. Eventually, the Universe cools enough that gravitational collapse produces the first stars and galaxies. These emit photons with sufficiently high energies to reionize the neutral hydrogen that then makes up the intergalactic medium ([IGM](#)).

Reionization is, like recombination, not instantaneous; by the end of reionization, the Universe is again in an overall ionized state. The duration of reionization is not tightly constrained, but upcoming experiments like [HERA](#) and [SKA-low](#) hope to be able to do this [[45](#), [46](#)].

The Universe continues to build [LSS](#), with dark matter gravitational potentials continuing to grow from the initial density perturbations seen at recombination. Baryon acoustic oscillations ([BAOs](#)) in the [LSS](#) of the Universe, traced out across the cosmic web, can be correlated with the [CMB](#) (see section [2.3.2](#)).

A second tension which Λ CDM faces is the $S8$ tension, where early- (through the [CMB](#)) and late-time ([LSS](#) surveys) disagree on the value of the amplitude of matter fluctuations $S8 = \sigma_8 \sqrt{\Omega_m/0.3}$ (where σ_8 is the rms. of matter fluctuations). That S_8 describes the amplitude of matter clustering makes it a key probe of structure growth. Indirect measurements from the [CMB](#) gives a higher value for $S8$ than [LSS](#) surveys, which measure S_8 through cosmic shear and galaxy clustering [[22](#), [47](#), [48](#)]. As with the Hubble tension, unresolved systematics may be behind the tension, or alternatively new physics.

[‡]The [SLS](#) is more commonly referred to as the ‘last scattering surface’ ([LSS](#)), but to avoid confusion with large-scale structure ([LSS](#)) I use [SLS](#) instead in this thesis.

1.9 The Stochastic Gravitational Wave Background

In the same way that we have an electromagnetic wave background—the [CMB](#)—we also have a gravitational wave background. Unlike the [CMB](#), which was generated during a single event (recombination), the [SGWB](#) is a superposition of low-amplitude gravitational waves from energetic astrophysical and cosmological events throughout cosmic history. The inflationary tensors discussed in section [1.8.2](#), for example, produce a contribution to the [SGWB](#). Other cosmological and astrophysical sources which might contribute are discussed in Chapter [4](#).

By definition, the [SGWB](#) carries no phase information, and measurements can only be made in terms of its strength at different frequencies—the strain h .

The first detection of the [SGWB](#) was made recently at pulsar timing array ([PTA](#)) frequencies [[49](#), [50](#), [51](#), [52](#), [53](#)], with the tightest constraints provided by [NANOGrav](#) [[49](#), [50](#)]. Thus far, measurements are not accurate enough to determine the source of this signal.

Once the [SGWB](#) has been well-characterized across multiple frequency bands, it has the potential to become a powerful cosmological tool, in complement to the [CMB](#). Crucially, it offers access to early epochs which are hidden from electromagnetic probes by the [SLS](#). As such, studying the [SGWB](#) may illuminate new physics—from inflationary dynamics to high-energy phase transitions—which are otherwise inaccessible.

A significant portion of this thesis (Chapters [4–7](#)) focuses on modelling, forecasting, and interpreting the [SGWB](#) in both astrophysical and cosmological contexts, using upcoming observations from [GW](#) detectors.

1.10 Summary and Discussion

Standard modern cosmology falls under the Λ CDM model, which has remained a solid framework to explain most of the phenomena observed in the Universe today. With the founding principles of homogeneity, isotropy, and [GR](#), and the addition of cold dark matter and dark energy (in the form of a pure cosmological constant Λ), Λ CDM explains the evolution of the Universe through the observed [CMB](#) and [LSS](#).

But Λ CDM is not a concrete series of perfectly interlinked theories. In reality, various tensions have arisen between different elements of cosmology—see the discussions of the H_0 and $S8$ tensions (sections [1.5](#) and [1.8.5](#)). Λ CDM also requires the existence of dark matter and dark energy, but does not offer a physical explanation of either.

Phenomena such as [GWs](#) and the [SGWB](#), which have only recently been detected, provide

opportunities for further testing of Λ CDM as a model. GWs may offer opportunities for resolving the H_0 tension, while the SGWB carries information about the Universe from an earlier time than has been observationally accessible before.

Chapter 2

The Cosmic Microwave Background

This chapter provides a more detailed overview of the [CMB](#), with a focus on the generation of its temperature and polarization anisotropies. It covers the physical origins of these fluctuations, the observational history of the [CMB](#), the mathematical formalism used to extract cosmological information from [CMB](#) data, and the instrumental and analytical challenges involved in its measurement.

Portions of this chapter are adapted from a set of lecture notes I wrote in 2019 [[1](#)], which in turn drew from a literature review I wrote the same year [[2](#)]. Other sources are cited directly in the text.

2.1 Introduction

Emitted during the epoch of recombination ($z \approx 1100$), the cosmic microwave background ([CMB](#)) is thermal radiation left over from the early Universe. It was discovered by chance by Penzias and Wilson in 1964 [[54](#)], and provides a unique window into the Universe as it was approximately 380,000 years after the Big Bang. By this time, the Universe has cooled sufficiently for protons and electrons to combine into neutral hydrogen, decoupling photons from matter and allowing them to travel freely through space.

Today, we observe this radiation as an almost uniform microwave background with a temperature of approximately 2.725 K. However, detailed observations reveal minute anisotropies at the level of one part in 100,000. These fluctuations are relics of the density variations in the primordial plasma that eventually gave rise to galaxies and large-scale structure ([LSS](#)).

The photons we observe today tell us not only about the physical processes which generated the [CMB](#) during recombination, but also inform us about the conditions of the Universe at early times, and its subsequent evolution.

2.2 The Surface of Last Scattering

The surface of last scattering ([SLS](#)) represents the boundary at which the Universe transitioned from opaque to transparent—see the discussion in section [1.8.4](#). This transition creates a ‘surface’ from which photons scattered, and then travelled (almost) directly to us today.

We observe this surface as a spherical shell surrounding us, with its radius given by the comoving distance that the photons have travelled since recombination. Photons from the [SLS](#) are redshifted by the time we observe them, due to cosmic expansion.

The spherical shell has a thickness, since recombination is not an instantaneous process. This thickness corresponds to the range of optical depth over which photons last scatter. To understand this, we can make the analogy of the opacity of a cloud (which will depend on the number of water droplets along the line of sight). With sufficiently high optical depth, the cloud is effectively opaque, and we can only observe the surface closest to us. However a lower optical depth gives more transparency, and we can see further into the cloud. The [SLS](#) corresponds to the surface where the optical depth (from Thomson scattering) switched from high to low, and represents the furthest back in time we can directly observe the Universe with electromagnetic information (although the [CMB](#) does still give us information about earlier epochs).

2.3 Temperature Anisotropies

The anisotropies observed in the [CMB](#) today can be traced back to physical processes in the very early Universe. During this time, the Universe is filled with a photon-baryon plasma. This plasma experienced gravitational interactions and pressure forces that led to the propagation of acoustic oscillations—sound waves. Denser regions experienced stronger gravitational compression, which increased the temperature and pressure of the photon fluid.

When recombination occurred, the photons decoupled, meaning that their mean free path increased drastically and they effectively ceased to interact with matter. From this point onwards, photons streamed freely through the Universe. The temperature distribution of the photons at recombination was frozen in, preserving the imprint of these oscillations.

2.3.1 Quantum Fluctuations and Inflation

Although the Cosmological Principle leads us to assume that the Universe in its earliest state is homogenous and isotropic, we still expect some inhomogeneity in the form of quantum fluctuations. During the period of inflation, these are stretched exponentially to reach macroscopic

scales. These, in turn, give rise to density perturbations in the plasma that fills the Universe during recombination.

The initial quantum fluctuations are understood to have Gaussian probability distributions, and in simple models this Gaussianity persists through to the density perturbations. However, models with non-Gaussianity have been proposed, and there are active searches for non-Gaussianity in CMB data [55, 22].

2.3.2 Baryon Acoustic Oscillations (BAOs)

Prior to photon decoupling, photons and baryons exist in a plasma which is tightly coupled through Thompson scattering (see discussion in section 1.8.4). Gravitational potential wells, caused by overdensities in the dark matter distribution, draw in baryonic matter. Photons are coupled to baryons via scattering, and they are effectively dragged into the wells as well. This compresses the plasma from the baryon overdensity, and the subsequent radiation pressure of photons acts to resist the over-compression. This generates baryon acoustic oscillations (BAOs) in the plasma, which propagate as spherical pressure waves at relativistic sound speeds.

BAOs leave a characteristic imprint in both the CMB and the distribution of LSS. When photons decouple at recombination (discussed below in section 2.3.3), the propagation of acoustic waves halts, and overdensities are frozen into place. As a result, there is a slight preference for matter to be distributed at the BAO scale—a fixed comoving scale (~ 150 Mpc) which can be observed in galaxy clustering. Since the imprints of BAOs in the CMB and LSS lie at the same angular scales, cross-correlations between the CMB and LSS are a useful cosmological tool.

2.3.3 Recombination

The temperature of the CMB as it is observed now is linked to baryon density in local regions of the plasma, specifically the densities as they were when CMB photons were first free to travel across the Universe. Photon decoupling occurs during recombination, at which point the mean free path of photons increases dramatically—photons cease to be trapped in their local regions and are free to travel across the Universe almost completely unimpeded. They carry with them information about their local regions at the time at which the photons last scattered.

At the point of Recombination, photon frequencies relate to the conditions of the local regions they were previously locked into: a higher local density ρ leads to higher baryon pressure P ; this generates a higher local temperature T , which itself leads to photons having higher

energies, and thus also higher frequencies ν . So the photons, now free to travel throughout the Universe, carry with them information about the initial conditions of the photon-baryon plasma through their frequencies. By measuring the temperature of the CMB now, we can reconstruct the state of the photon-baryon plasma at recombination.

The Sachs-Wolfe Effect (SWE)

The Sachs-Wolfe effect (SWE) [56] introduces temperature fluctuations in the CMB via gravitational red- and blueshifting, and is the main source of anisotropies at large angular scales ($\gtrsim 10^\circ$). It consists of two parts: the Non-Integrated SWE and the Integrated SWE (covered in section 2.3.4).

The Non-Integrated SWE occurs during recombination, when CMB photons are redshifted as they escape uneven gravitational potential wells created by the uneven distributions of dark matter. This uneven distribution leads to variable redshifting, causing anisotropies. These effects appear on large angular scales due to the relatively large size of dark matter overdensities compared to those of the photon-baryon fluid.

Damping

CMB photons last scattered at different times, depending on the temperature and density of their local regions, which gives the SLS its finite thickness. During this period, photons diffuse from higher to lower temperature areas. The diffusing photons carry electrons with them, which subsequently attract protons. This adiabatically damps the density fluctuations in the plasma, and makes the matter more isotropic at the smallest angular scales ($\ell \gtrsim 800$). This process is called Silk damping [57].

2.3.4 Secondary Anisotropies

The physical process above (BAOs, Silk damping and the Sachs-Wolfe effect) generate primary anisotropies in the CMB—i.e., anisotropies which are generated during recombination, when the CMB itself is generated. Additional, secondary, anisotropies affect the CMB photons as they travel to us.

The secondary anisotropies present in the observed CMB are generated through the interaction of CMB photons with the evolving large-scale structure in the Universe. The dominant effects are described below.

Reionization

The ionized plasma generated during the epoch of reionization (see section 1.8.5) continues to make up the majority of the intergalactic medium (IGM), and scatters CMB photons through Thomson scattering. This leads to an erasure of temperature anisotropies at angular scales above 10° , by a factor of $e^{-\tau}$, where τ (the optical depth to reionization) is a measure of the opacity of line-of-sight observations of the CMB. This can, therefore, be a probe of the duration of reionization.

Scattering during reionization also generates large-scale polarization, which is discussed further in section 2.4.3.

The Integrated Sachs-Wolfe (ISW) Effect

The integrated Sachs-Wolfe (ISW) effect results from gravitational redshifting of CMB photons as they traverse evolving gravitational potentials between recombination and the present day. If a gravitational potential evolves while a photon is passing through it, the photon can gain or lose energy. This modifies the observed temperature of the CMB on large scales.

The ISW effect can be divided into two components: early- and late-time ISW. Early-time ISW occurs shortly after recombination, when residual radiation contributes to the decay of gravitational potentials. As photons exit these decaying wells, they gain energy, contributing to anisotropies. Late-time ISW occurs when dark energy dominates the energy density of the Universe. The accelerated expansion causes gravitational potentials to decay during photon transit.

A related but subtler contribution comes from the Rees-Sciama (RS) effect, which occurs when gravitational potentials evolve nonlinearly (such as during the collapse of structures like galaxy clusters) [58]. Although smaller in amplitude than the ISW effect, the RS effect can contribute to CMB anisotropies on smaller angular scales.

The Sunyaev-Zeldovich (SZ) Effect

The Sunyaev-Zeldovich effect (SZ) effect modifies the CMB via inverse Compton scattering, where high-energy electrons—typically found in galaxy clusters—transfer energy to CMB photons, boosting their frequency.

These electrons gain energy either thermally (from high temperatures) or kinematically (from bulk motion). The SZ effect is split into the thermal Sunyaev-Zeldovich (tSZ) and kinematic Sunyaev-Zeldovich effect (kSZ), which generate different spectra—the tSZ gives a

net increase or decrease on the temperature at a given frequency, where it shifts low-energy photons to higher energies; the [kSZ](#) shift the entire spectrum, retaining the blackbody spectral shape.

The [SZ](#) effect appears on small angular scales and has a relatively low amplitude, making it challenging to isolate from other signals. However, since it is independent of redshift, it can be used to probe galaxy clusters across all epochs.

2.4 Polarization

The [CMB](#) we observe now is polarized to a small degree ($\sim 5\,\mu\text{K}$) as the result of scalar and tensor perturbations, during recombination and later on.

Thomson scattering prior to recombination induces linear polarization in the photons within the primordial plasma. While a completely isotropic field of radiation in a region of homogeneous matter would pick up a net polarization of zero, temperature quadrupole anisotropies induce linear polarization. As discussed in section [2.3](#), initial scalar perturbations generate density fluctuations in the primordial plasma, and so the photons pick up linear polarization.

Tensor perturbations, for example from a gravitational wave background ([GWB](#)) propagating during recombination, induce polarization through the same mechanism. The squeezing and stretching of space during tensor perturbations blue- and redshifts photons in the direction of the polarization of the gravitational wave ([GW](#)), hence generating additional temperature quadrupoles.

Polarization of radiation can be broken down into the four Stokes parameters, I , Q , U , and V , where I gives the total intensity (i.e. temperature), Q and U give the differences between the intensities of two directions of linearly polarization (such that U has axes at 45° to Q), and V the intensity of circular polarization. Note that Thomson scattering does not generate circular polarization and hence the standard model predicts $V = 0$.

Importantly, Q and U are not invariant under rotations of the co-ordinate system, and hence they are complicated to deal with on the full sky. [CMB](#) polarization can instead be described in terms of E - and B -modes, where E -modes have gradient but zero curl, and B -modes curl but zero gradient. This parametrization is invariant under rotations.

2.4.1 *E*-mode Polarization

E-modes, which are generated by scalar and tensor perturbations, strongly dominate the polarization of the CMB—see figure 2.2. We have good measurements of *E*-modes from existing CMB experiments.

2.4.2 *B*-mode Polarization

B-modes are generated by tensor perturbations (GWs) only and by the lensing of pre-existing *E*-modes (see section 2.4.4 below). *B*-modes in the CMB have not yet been well-measured, but present and upcoming experiments (e.g. Simons Observatory (SO), see Chapter 3) hope to do this.

Measurements of *B*-modes will act as a constraint on the primordial GWB—any GWB already propagating by recombination is equivalently a series of tensor perturbations which will induce *B*-modes in the CMB. We can work backwards from measurements of *B*-modes to possible models which could generate them (see discussion in Chapter 4).

2.4.3 Reionization

Scattering during reionization suppresses the polarization signal by the same factor it does the temperature signal ($e^{-\tau}$) but it also generates new polarization over a large angular scale. This polarization signal (which contains both *E*- and *B*-modes) has an amplitude proportional to τ and the amplitude of the reionization temperature quadrupole.

2.4.4 Gravitational Lensing

CMB photons encounter deflection from LSS gravitational potentials on their path to us as observers. This causes distortions of the CMB at small scales.

Lensing affects both temperature and polarization, both by smearing anisotropies and by the introduction of small non-Gaussianities.

Its more dramatic effect, though, is the conversion of pre-existing *E*-mode polarization in CMB photons to *B*-mode polarization. These act effectively as a foreground for the primordial *B*-modes, and must be removed through delensing techniques (see section 2.6.5). Figure 2.1 shows the strength of lensing vs primordial *B*-modes, and that the multipole range where primordial *B*-modes are most detectable is at around $20 \lesssim \ell \lesssim 100$.




Figure 2.1: Summary of current measurements and upper limits of the B -mode polarization power spectra of the CMB. Figure credit: BICEP/Keck Collaboration [59].

2.4.5 Synchrotron and Dust Foregrounds

Galactic synchrotron and dust emissions are the dominant polarization foregrounds at microwave frequencies, and must be well understood and modelled in order to extract a primordial B -mode signal.

Synchrotron emissions, generated by energetic electrons in the magnetic fields of the Galaxy, have a broad spectrum characterized by $I_\nu \propto \nu^{-\beta}$ (for $\beta \sim 3$, with the precise value determined by the frequency and particular line of sight). It has a high polarization fraction, $\gtrsim 20\%$ (compared with the E -mode polarization fraction of roughly 10%, and the B -mode fraction which is much smaller), and is strongest at low frequencies (below 70 GHz).

Galactic dust generates thermal radiation as it is hit by incident starlight, which can be described by an approximately blackbody spectrum: $I_\nu \propto \nu^{\beta_d} B_\nu(T_d)$, where $|\beta_d| \ll 1$ is the dust spectral emissivity index, B_ν is the standard Planck function, and T_d is the temperature of the dust.

This radiation is also polarized as a result of the Galaxy's magnetic fields, in this case because the dust aligns itself along the magnetic field lines. It is strongest at frequencies above 100 GHz.

2.5 Analysis of CMB Anisotropies

The anisotropies of the CMB contain rich information about the early Universe. These fluctuations are characterized statistically through their angular power spectra, which describe the variance of signals across different scales.




Figure 2.2: Summary of current measurements and upper limits of the temperature and polarization power spectra of the CMB. Figure credit: Planck Collaboration [60].

2.5.1 Temperature Power Spectrum

We define the temperature fluctuations in a given direction $\hat{n}$:

$$\frac{\Delta T(\hat{n})}{\bar{T}} = \frac{T(\hat{n}) - \bar{T}}{\bar{T}}, \quad (2.1)$$

where $\bar{T}$ is the mean temperature.

Since we observe the [CMB](#) as a spherical shell, it is useful to expand the temperature fluctuations in spherical harmonics:

$$\frac{\Delta T(\hat{n})}{\bar{T}} = \sum_{\ell=0}^{\infty} \sum_{m=-\ell}^{\ell} a_{\ell m} Y_{\ell m}(\hat{n}). \quad (2.2)$$

where the expansion coefficients are

$$a_{\ell m} = \int d\Omega Y_{\ell m}^*(\hat{\mathbf{n}}) T(\hat{\mathbf{n}}). \quad (2.3)$$

Note that the multipole $\ell = 0$ denotes the sky-averaged mean, and $\ell = 1$ the dipole, which is typically removed since it is dependent on local velocity—the dipole is primarily caused by Doppler shift from the motion of the Solar System relative to the [CMB](#) rest frame.

The power spectrum is defined as $C_{\ell} = \langle |a_{\ell m}|^2 \rangle$, and an unbiased estimator of that is given by

$$\hat{C}_{\ell}^{TT} = \frac{1}{2\ell + 1} \sum_{m=-\ell}^{\ell} |a_{\ell m}|^2. \quad (2.4)$$

This equation shows the cosmic variance which is present due to the finite number of observable modes (i.e. $2\ell + 1$). Assuming Gaussianity, $\hat{C}_{\ell}^{TT}$ is Chi-squared distributed with $2\ell + 1$ degrees of freedom. The cosmic variance is therefore

$$\text{var}(\hat{C}_{\ell}^{TT}) = \frac{2}{2\ell + 1} (C_{\ell}^{TT})^2. \quad (2.5)$$

Per figure [2.3](#), we can derive a range of cosmological parameters from the temperature power spectrum. The position of the first acoustic peak (at around one degree) measures the sound horizon, and provides constraints on the curvature of the Universe. The heights and positions of the subsequent peaks provide insights into the baryon-to-photon ratio (which affects the relative peak amplitudes) and the total matter density (affecting the overall shape of the spectrum). The damping tail at small angular scales contains information about the processes occurring




Figure 2.3: The effects of varying cosmological parameters on the CMB temperature angular power spectrum—figure credit: Wayne Hu & Scott Dodelson [61].

during recombination (e.g. details of Silk damping and the thickness of the SLS).

By precisely measuring these features, we can tightly constrain fundamental parameters including the Hubble constant, baryon density, dark matter density, and dark energy content, thereby shaping our understanding of the Universe’s composition, geometry, and evolution.

2.5.2 Polarization Power Spectra

This formalism is used similarly for E - and B -mode polarization, with spin-2 spherical harmonics $_{\pm 2}Y_{\ell m}$ [62]:

$$\hat{C}_{\ell}^{EE} = \frac{1}{2\ell + 1} \sum_{m=-\ell}^{\ell} |a_{E,\ell m}|^2, \quad (2.6)$$

$$\hat{C}_{\ell}^{BB} = \frac{1}{2\ell + 1} \sum_{m=-\ell}^{\ell} |a_{B,\ell m}|^2, \quad (2.7)$$

where the coefficients are given by

$$a_{E,\ell m} = -(a_{2,\ell m} + a_{-2,\ell m})/2, \quad (2.8)$$

$$a_{B,\ell m} = i(a_{2,\ell m} - a_{-2,\ell m})/2, \quad (2.9)$$

with

$$a_{2,\ell m} = \int d\Omega {}_2Y_{\ell m}^*(\hat{\mathbf{n}})(Q + iU)(\hat{\mathbf{n}}), \quad (2.10)$$

$$a_{-2,\ell m} = \int d\Omega {}_{-2}Y_{\ell m}^*(\hat{\mathbf{n}})(Q - iU)(\hat{\mathbf{n}}). \quad (2.11)$$

There is a cross-correlation between temperature and E -modes, which stems from their generation by the same density fluctuations during recombination. Their cross-spectrum (C_ℓ^{TE}) is, therefore, a useful tool for comparison with each of their auto-spectra. In contrast, there is no correlation between B -modes and temperature or E -modes, due to parity—temperature and E -modes are parity-even, while B -modes are parity-odd, and so their cross-spectra vanish in a parity-conserving universe.

The cosmic variances for the power spectra for polarization are similarly

$$\text{var}(\hat{C}_\ell^{EE}) = \frac{2}{2\ell + 1} (C_\ell^{EE})^2, \text{var}(\hat{C}_\ell^{BB}) = \frac{2}{2\ell + 1} (C_\ell^{BB})^2. \quad (2.12)$$

The cosmic variance for the cross-spectrum between temperature and E -modes is

$$\text{var}(\hat{C}_\ell^{TE}) = \frac{2}{2\ell + 1} [(C_\ell^{TE})^2 + C_\ell^T C_\ell^E]. \quad (2.13)$$

2.6 Observational Techniques and Data Reduction

Moving from the observable signal on the sky to [CMB](#) science requires careful methods. Some of these are built in to the experiments themselves, and some are carried out once data has been collected.

2.6.1 Scanning Strategy

There are different options for how experiments scan the observable sky. A scan strategy is chosen to optimize for certain characteristics, given the priorities of the experiment. These could be to maximize coverage or depth, to aim for uniformity across the sky, to minimize

systematics or atmospheric noise, or to boost cadence for non-CMB observations.

My work on scan strategies for both of SO's different telescopes is detailed in Chapters 8 and 9.

2.6.2 Multi-frequency Observations

The different microwave foregrounds which contaminate CMB data each dominate at different frequencies—e.g. Galactic synchrotron radiation is strongest at low frequencies, while Galactic dust is strongest at high frequencies.

Observing across multiple frequencies enables separation of the CMB from each of the foregrounds, provided that we can model their spectra.

This is particularly important for *B*-mode detections, where the polarized synchrotron and dust foregrounds can be mistaken for primordial *B*-modes—see figure 2.4.




Figure 2.4: CMB foregrounds at different frequencies—figure credit: Planck Collaboration and ESA.

2.6.3 Noise Modelling

Beyond foregrounds, noise from the atmosphere and instruments must be modelled, to then be deconvolved from the data.

Noise in CMB data arises from multiple sources. There is white noise, primarily due to detector and photon shot noise, which dominates at small angular scales. $1/f$ noise is caused by, e.g., slow drifts in detector gain, temperature fluctuations and instrumental instabilities. It can generate correlations and striping in the maps. Scan synchronous signals are correlated with the telescope’s scanning motion, e.g. from effects like mechanical vibrations or ground pickup. Additional correlated noise between detectors, e.g. from atmospheric fluctuations, can be present. Finally, beam uncertainties and pointing errors behave as noise-like systematics, and are commonly treated within the noise model.

2.6.4 Component Separation and Map-making

Statistical methods (e.g. Bayesian component separation) are used to isolate the CMB signal from the identified foregrounds and noise. Then, cleaned maps of temperature and polarization can be produced.

2.6.5 Delensing

Delensing involves reconstructing the CMB lensing potential, usually by applying quadratic estimators to temperature or E -mode data, or by using external tracers (e.g. the cosmic infrared background (CIB) or galaxy surveys). This reconstructed potential is then combined with observed E -mode maps to estimate and subtract the lensing-induced B -modes. Reconstruction is imperfect, and current methods can only reduce lensed B -modes by around 30–50% [63, 64].

2.6.6 Power Spectrum Estimation

Finally, the power spectra as defined in section 2.5 are extracted from the cleaned maps. Ideally this is done using maximum likelihood (ML) estimators that approximate a fully Bayesian analysis, but pseudo- C_ℓ estimators (which are computationally simpler) are also used.

2.7 History of CMB Observations and Results

2.7.1 Initial Discovery and Early Measurements

The existence of some form of primordial ‘relic’ radiation was theorised in 1948 by Alpher and Gamov [65, 66], but the idea did not initially gain scientific momentum—it contradicted the steady state model of the Universe, the contemporary mainstream cosmology.

In 1964, an accidental discovery of a faint isotropic microwave signal was made by Penzias and Wilson [54]. This, first, observation of the CMB was strong but not conclusive evidence for the Big Bang theory. The pair were later awarded the Nobel prize for their discovery.

In the 1970s and 80s, a series of ground-based and balloon experiments sought to improve the sensitivity of our measurements of the CMB, looking to characterize both the overall temperature signal and any anisotropies present.

Their sensitivities were insufficient to detect anisotropies in the temperature, but provided upper limits and established the possibility of precise measurements of the CMB in the future.

2.7.2 Satellite Experiments

The first CMB satellite experiment was NASA’s COBE, launched in 1989 for a four year observing period. It made a precise measurement of the temperature spectrum of the CMB, showing it to be a near-perfect blackbody at a temperature of 2.726 ± 0.010 K [67] (conclusive support for the Big Bang theory over a steady state model), and also made the first observation of a temperature anisotropy at a scale of 1 in 100,000 [68].

COBE was followed by WMAP, operating between 2001–2010. The temperature maps it produced were at significantly higher sensitivity and angular resolution than COBE, sufficient to detect a series of acoustic peaks in the angular power spectrum, and to derive cosmological parameters [69]. Its measurements of polarization (with measurements of E -modes, but only upper limits of B -modes) gave the first detection of CMB polarization.

The most comprehensive CMB experiment to date was Planck. Launched in 2009 for a four year observing run, its final data release was published in 2018 [22]. Planck improved on the measurements made by WMAP in both sensitivity and resolution, providing full-sky maps of temperature and E -mode polarization anisotropies. The temperature power spectrum measured by Planck is at almost cosmic variance levels, and subsequent improvements from other experiments to the temperature power spectrum will have to come in the form of measurements at higher ℓ —see figure 2.2.

2.7.3 Ground-Based and Suborbital Experiments

Alongside the three satellites, ground-based and suborbital balloon experiments have been run. While these cannot provide full-sky maps, and have to deal with atmospheric noise, they have probed the CMB at higher resolutions, often focusing on small angular scales and improved polarization measurements.

Ground-based experiments include: [DASI](#) [70], [BICEP-Keck](#) [71, 59] and [SPT](#) [72] (based in Antarctica); and [CBI](#) [73], [CLASS](#) [74], [POLARBEAR](#) [75] and [ACT](#) [76] (Chile). Both locations are chosen as high, dry areas to reduce atmospheric noise. Balloon experiments include [SPIDER](#) [77] and [BOOMERanG](#) [78].

[DASI](#) was the first experiment to detect polarization in the [CMB](#), while [CBI](#) provided the first detailed E -mode spectrum. Subsequent experiments have continued to make improved observations of the [CMB](#), and upper limits on the B -mode power spectrum, or equivalently on the tensor-to-scalar ratio r .

2.7.4 Upcoming and Future Experiments

The next generation of [CMB](#) experiments hope to make detections of B -modes, as well as continuing to improve on temperature and E -mode observations with more sensitive measurements at better resolutions and higher ℓ .

Simons Observatory ([SO](#)) in Chile (discussed in detail in Chapter 3) is in initial science observations ([ISO](#)), and expects to begin its first observing run soon [79]. [CMB-S4](#), with telescopes in both Chile and Antarctica, aims to bring further improvements to polarization [80].

[JAXA](#)'s [LiteBIRD](#) is a satellite experiment, expected to launch within the next decade [81, 82]. It will bring improved polarization measurements at large angular scales.

2.8 Summary and Discussion

The [CMB](#) acts as a tool for cosmologists to probe the early Universe, particularly as it was during recombination. Ever-increasingly sensitive experiments are reaching cosmic variance limits on measurements of its temperature and E -mode power spectra, while present and future experiments (including [SO](#)) expect to move from upper limits to accurate measurements of its B -mode polarization.

Analysis of the [CMB](#) remains an active field of research in its own right, but measurements of [CMB](#) B -modes will also play a role in our understanding of the primordial [GWB](#), a major component of the overall stochastic gravitational wave background ([SGWB](#)).

Chapter 3

The Simons Observatory

This chapter gives an overview of the [SO](#) experiment, as part of which collaboration I have carried out work on forecasting and scan strategy (detailed in Chapters [8](#) and [9](#)). This chapter is based largely on the public Simons Observatory ([SO](#)) and Simons Foundation websites [[83](#), [84](#)] and the 2018 [SO](#) white paper [[79](#)], as well as more specialized papers, which are cited in the text.

3.1 Introduction

Simons Observatory ([SO](#)) is a cosmic microwave background ([CMB](#)) experiment, currently in initial science observations ([ISO](#)), which is situated at an altitude of 5190 m on Cerro Toco in the Atacama Desert in Chile. It aims to make detailed measurements of the temperature and polarization of the sky at sub-millimetre frequencies. The initial [SO](#) configuration comprises a set of three 0.42 m Small Aperture Telescopes ([SATs](#)) and one 6 m Large Aperture Telescope ([LAT](#)). First light for each of the [SATs](#), and the [LAT](#) receiver, was achieved across a few months in late 2024, with first light for the full [LAT](#) setup following in February 2025. Full observations are expected to start later this year.

The sky regions which [SO](#) will cover with the [LAT](#) and [SATs](#) coincide with many other surveys, including significant overlap with [LSST](#), [DES](#), [DESI](#) and Euclid. Six frequency bands will be probed, from 27 to 280 GHz.

[SO](#) begins with an initial five year observing run, with further observing runs to be made after upgrades and additions to the hardware. Advanced Simons Observatory ([ASO](#)), along with Simons Observatory: UK ([SO:UK](#)) and Simons Observatory: Japan ([SO:JP](#)), will bring upgrades to [SO](#) infrastructure for its second observing run, due to begin in 2030.

[ASO](#) will provide additional detectors for the Large Aperture Telescope Receiver ([LATR](#)),

and also the installation of solar panels to fulfil 70% of the energy requirements, which will add power reliability and a subsequent reduction of telescope downtime by around 10%.

SO:UK is slated to provide two additional SATs (with SO:JP contributing a third), an optics tube for the LATR, a data centre based at Manchester University to deliver ‘science ready’ data products, and analysis algorithms for data processing. Future international contributions may also provide additional hardware and resources to the SO configurations.

3.2 The Telescopes

For its initial observing run, SO has its 6 m LAT and three 0.42 m SATs, similar in setup to ACT and BICEP3, respectively. 60,000 detectors are split equally between the LAT and the set of three SATs, with each detector able to detect one orthogonal linear polarization, and a single frequency band. Four detectors make up a pixel, such that both orthogonal polarizations and two frequency bands can be measured [85, 86]. More detectors will be added, along with additional SATs, for later observing runs, but the initial setup already contains more detectors than all previous CMB experiments combined and represents a sizeable increase over other contemporary arrays—SPT has $\sim 16,000$ detectors [87].

Observing the CMB over the greatest range of angular scales, from arcminute to degree scales, requires two different sizes of telescope [80]—the larger size of the LAT contributes high resolution for the smallest angular scales, while the SATs contribute to the largest angular scales, and are less sensitive to atmospheric effects. In terms of measuring polarization, this means that the SATs’ measurements of B -modes will be used to constrain the tensor-to-scalar ratio r (see section 1.8.2). The LAT will be the instrument capable of making detailed measurements of lensing B -modes, which can then be subtracted from the SAT polarization maps to make possible detections of the primordial gravitational wave background (GWB) though its imprint on B -modes, and its precise measurements of temperature power spectrum damping at small scales would show the effects of the primordial GWB.

3.2.1 The Large Aperture Telescope

The LAT is a modified crossed-Dragone telescope. The design of the LAT was done alongside the almost identical design for CCAT’s FYST (previously named ‘CCAT-prime’).

A crossed-Dragone telescope, first proposed by Dragone in 1978 [88], places the secondary mirror off-axis to the primary mirror to avoid blocking any incident light. This requires a




Figure 3.1: (a) A cutaway rendering of the LAT, showing the positions of the primary and secondary mirrors, as well as the LATR. (b) A rendering of the LATR. Figure credit for both: SO Collaboration.

secondary mirror of similar size to the primary mirror—the LAT’s secondary mirror is 5.9 m in diameter, compared to the primary’s 6 m diameter. This larger secondary mirror necessarily increases the difficulty and cost of manufacture, but the setup also reduces systematics and increases the field of view (FoV), while maintaining the polarization of incoming light (which contains cosmological information) through to the receiver.

The LAT makes further improvements to the available FoV (reaching 7.8° at a wavelength of 3 mm) by switching to aspheric mirror panels—a standard crossed-Dragone telescope has a parabolic primary mirror and hyperbolic secondary. However, both of the LAT’s mirrors are quasi-hyperbolic (with each of the 77/69 panels in the primary/secondary mirrors having a different shape) to reduce comatic aberration (an effect which introduces distortion to objects observed off-axis). FYST and the LAT are the first crossed-Dragone telescopes built which improve on the classical mirror shapes [89].

The LAT is equipped with a 2.4 m receiver, the LATR, which is the largest cryogenic millimetre-wave telescope receiver to date. The instruments in the LATR must be kept at 0.1 K, and to fulfil this the whole LATR is in vacuum. The LATR is supported by the corotator, which enables the receiver to maintain its relative orientation to the mirrors when the observing elevation is changed—see figure 3.1b. There are currently seven optics tubes (OTs) in the LATR, with space for six more. The OTs observe across six frequencies: low- (30 and 40 GHz), mid- (90 and 150 GHz), and ultra high-frequencies (230 and 280 GHz).

The 30,000 transition edge sensor (TES) detectors, split between these OTs, are at approximately the ultimate sensitivity limit. As such, the only way to make more sensitive




Figure 3.2: The [LAT](#)’s first light image of Mars, which demonstrates complete end-to-end functionality. This map was made using 743 of the 24,000 detectors, with a [SNR](#) of > 4000 per detector. Figure credit: SO Collaboration.

measurements of the [CMB](#) will be to increase the number of detectors. [ASO](#) will bring the number of detectors in the [LATR](#) up to 62,000 with the addition of another six [OTs](#), including an ultra-high frequency [OT](#) from [SO:UK](#). This will enable the [LAT](#) to map the sky twice as fast.

The [LAT](#) scans the sky at a specified and variable elevation, scanning back and forth as the Earth rotates. The baseline scan strategy included in the [SO](#) white paper [79] was a classical strategy which observed four fields, detailed in reference [90]. As discussed in detail in Chapter 8, my work on optimizing the scan strategy gives an improvement to the scan against a range of figures of merit. See figure 8.5 as an example of a map of the observable sky.

The [LAT](#) achieved first light in late February 2025—figure 3.2 shows the first light image of Mars. There follows a period of commissioning, before the [LAT](#) begins [ISO](#), and finally full science operations later in the year [91].

3.2.2 The Small Aperture Telescopes

The SATs are 0.42 m in diameter, with FoVs of 35° . The SATs are designed for measurements over large angular scales, $30 \leq \ell \leq 200$, where we expect the peak of any primordial B -mode spectra to be most visible, from the combined GWB and lensing foreground spectra [92]—see figure 2.1. The SATs are equipped with rotating half-wave plates (HWPs) so that CMB polarization can be better separated from slow-moving unpolarized atmospheric fluctuations, and systematic effects.

Over the initial five year observing run, the combined sensitivity of the SATs will reach $2 \mu\text{K-arcmin}$, at subdegree resolutions. Our deepest CMB polarization data currently comes from BICEP3, which was at $2.8 \mu\text{K-arcmin}$, carried out at degree-scale resolutions [71]. The baseline scan strategy used in the white paper [79], and detailed in [90], covers a patch of sky in the South field, split into tiles, optimized for even coverage and depth. The work discussed in Chapter 9 improves upon this strategy against a series of figures of merit, balancing improvements to the sensitivity to the tensor-to-scalar ratio r with instrumental constraints.

SO:UK will provide two additional mid-frequency SATs. These will differ from the initial SO SATs slightly: they will use kinetic inductance detectors (KIDs) instead of TESs and a two-lens configuration (instead of the current three-lens setup). There is also scope for making the HWPs from metal-mesh instead of sapphire. SO:JP will provide one additional low-frequency SAT.

3.3 Science Goals of SO

SO has six primary science goals, aiming to further investigate: primordial perturbations (see section 3.3.1); the effective number of relativistic species N_{eff} (section 3.3.2); neutrino mass (section 3.3.3); deviations from the cosmological constant Λ (section 3.3.4); galaxy evolution (section 3.3.5); and the duration of reionization (section 3.3.6). There are also a number of secondary science goals (section 3.3.7), as well as a goal of producing maps and catalogues derived from the CMB, to be publicly shared with the astrophysics and cosmology communities (section 3.3.8).

3.3.1 Primordial Perturbations

SO will probe primordial perturbations with three metrics: the tensor-to-scalar ratio r , the power spectrum of scalar perturbations, and non-Gaussian perturbations.

Tensor-to-Scalar Ratio

Recall from section 1.8.2 that the tensor-to-scalar ratio r gives information about the relative strength of primordial tensor perturbations.

The initial forecasting for SO’s sensitivity to r was carried out for the 2018 white paper [79], and gives a baseline sensitivity of $\sigma(r) = 0.003$ for a fiducial model with $r = 0$ (improved to $\sigma(r) = 0.002$ with improved forecasting and scan strategy, discussed in Chapter 9). This compares with our current upper limits of $r \leq 0.06$ and $r < 0.044$ [22, 93], as discussed in section 4.3.1.

SO will constrain r through its measurements of large-scale B -modes by the SATs. Discussed in more detail in section 5.3.3, a primordial **GWB** leaves imprints in the B -mode polarization of the **CMB**. These are expected to be stronger than the lensing foreground at $\ell \lesssim 100$ [94], which lies in the sensitivity range of the SATs ($30 \lesssim \ell \lesssim 300$). The LAT will play an important role in delensing.

Scalar Perturbations

Planck has constrained the amplitude of scalar perturbations in the early universe to within 3% [22].

SO expects to measure the amplitude to within 0.5% at $k = 0.2 \text{ Mpc}^{-1}$, extending our observations to smaller scales than Planck could reach, allowing us to compare predictions for the power spectrum from a variety of inflationary models over a wider range.

This will be done through the temperature and E -mode polarization auto- and cross-power spectra from the LAT.

SO will also constrain the scalar index n_s (see section 1.8.2) as a secondary science goal.

Non-Gaussian Perturbations

Multi-field models of inflation induce non-Gaussianities in the initial perturbations, which can be characterised by their amplitude f_{NL} . Measurements of $|f_{\text{NL}}| \gtrsim 1$ would rule out single-field models.

The local non-Gaussianity $f_{\text{NL}}^{\text{local}}$ will be constrained by SO in three ways, through correlations between LSST galaxy cluster catalogues, and the observed lensing (1) and kinematic Sunyaev-Zeldovich effect (**kSZ**) (2) seen in the **CMB**, as well as through the observed temperature and polarization bispectra (3).

Dependent on forecasting details, and the level of delensing reached, SO expects to reach

sensitivities between $\sigma(f_{\text{NL}}^{\text{local}}) = 3$ and $\sigma(f_{\text{NL}}^{\text{local}}) = 1$, to compare with our best current constraint of $f_{\text{NL}}^{\text{local}} = -0.9 \pm 5.1$ from Planck [95, 22].

As secondary science goals, [SO](#) will also constrain $f_{\text{NL}}^{\text{equilateral}}$ and $f_{\text{NL}}^{\text{orthogonal}}$.

3.3.2 Effective Number of Relativistic Species

The standard value (i.e., assuming only the three standard neutrino flavours) for the effective number of relativistic species is given by $N_{\text{eff}} = 3.046$ [96, 97, 98], where

$$\rho_{\text{total rad.}} = \rho_{\gamma} \left[1 + \frac{7}{8} \left(\frac{4}{11} \right)^{\frac{4}{3}} N_{\text{eff}} \right]. \quad (3.1)$$

However, this increases if there are more light species than the three standard neutrinos. Our current best constraints on N_{eff} come from Planck, with $N_{\text{eff}} = 2.99 \pm 0.17$ [22].

[SO](#) expects to improve these constraints by considering the temperature and E -mode polarization power spectra from the [LAT](#): N_{eff} affects the positions of the peaks, and the damping tails, of the TT , EE , and TE power spectra auto- and cross-correlations [99]. Forecasting from the 2018 baseline scan strategy gives the constraint as $\sigma(N_{\text{eff}}) = 0.07$. But since [SO](#) will be sample variance limited across a large range of ℓ , a higher observed sky fraction could improve this even further—see Chapter 8 for details on the updated scanning strategy.

The addition of light non-scalar particles in thermal equilibrium would induce variations of $\Delta N_{\text{eff}} \geq 0.047$. The constraint from [SO](#) is sensitive enough to rule out any models with at least three additional particles, at the 95% confidence level.

3.3.3 Neutrino Mass

The standard model of particle physics predicts massless neutrinos, but various extensions to the standard model confer masses to neutrinos.

A measurement of $\Sigma m_{\nu} \gtrsim 0.1$ eV would indicate a non-zero mass for neutrinos. Our best cosmological constraint on neutrino mass comes from Planck ($\Sigma m_{\nu} \leq 0.12$ eV [22]), while our best result independent of cosmological models comes from the neutrino experiment [KATRIN](#) gives a weaker constraint of $\Sigma m_{\nu} \leq 0.8$ eV at the 90% CL [100].

The expected [SO](#) constraint is $\sigma(\Sigma m_{\nu}) = 0.04$ eV, using a combination of three methods: combining both the [SO](#) lensing power spectrum (1) and thermal Sunyaev-Zeldovich ([tSZ](#)) maps (2) with [DESI](#) baryon acoustic oscillation ([BAO](#)) measurements; and Sunyaev-Zeldovich effect ([SZ](#)) cluster counts (3).

3.3.4 Deviations from a Cosmological Constant Λ

SO plans to evaluate deviations from a pure cosmological constant Λ (with constant equation of state $w = w$) with two metrics: the amplitude of matter perturbations σ_8 , and the Hubble constant H_0 .

Amplitude of matter perturbations at high redshifts

Our current CMB constraint on the amplitude of matter fluctuations, σ_8 , comes from Planck, at $\sigma_8 = 0.8101 \pm 0.0061$, from a combination of the temperature, E -mode polarization, and lensing power spectra [101, 22]. Non-CMB results come from DESI clustering analyses, with their best result as $\sigma_8 = 0.805 \pm 0.012$ [102].

SO will measure σ_8 through the correlation of the reconstructed lensing spectrum with LSST galaxy density maps.

SO plans to measure σ_8 out to redshifts of $z = 6$, with the best constraints expected to be at the 1–2% level over a range of redshift bins up to $z = 4$. This forecasting is very dependent on the sky coverage, and SO sensitivity—see figure 29 in reference [79] for details.

Separately, optical and weak lensing experiments (including DESI and Euclid) are expected to increase their sensitivity to σ_8 at $z < 1$, and as such the higher redshift SO observations are intended to work in complement.

Hubble constant

SO expects to reduce the uncertainty on H_0 , as derived from the CMB, to half a percent. This will improve on our best CMB constraints on H_0 which are at $\approx 0.74\%$ from Planck ($H_0 = 67.4 \pm 0.5 \text{ km s}^{-1} \text{ Mpc}^{-1}$ [22]). See section 1.5 for discussion of non-CMB measurements, in tension with this value.

The best SO constraint $\sigma(H_0) = 0.3$ on H_0 (in units of $\text{km s}^{-1} \text{ Mpc}^{-1}$) will come from a combination of the TT , EE , and TE power spectra, which are sensitive to H_0 . However each of these can also constrain H_0 individually (at levels of $\sigma(H_0) = 0.6/0.5/0.5$, respectively) which will act as quasi-independent checks on the measurement.

Dark Energy and Modified Gravity

SO will be able provide novel high redshift constraints on deviations from Λ , or equivalently deviations from ΛCDM , including from modified gravity or dark energy models—as secondary

science goals, [SO](#) will measure the dark energy equation of state coefficients w_0 , w_a , as well as the growth rate at high redshifts.

3.3.5 Galaxy Evolution

[SO](#) will measure two parameters of galaxy evolution: feedback efficiency η_{feedback} , and degree of non-thermal pressure support p_{nt} . The former is a measure of how much energy star formation, supernovae and active galactic nuclei ([AGNs](#)) inject into the intergalactic medium ([IGM](#)), while the latter is a measure of the support to the [IGM](#) generated by processes within the gas.

There are currently no strong constraints on either parameter. [SO](#) expects to be able to constrain them to 3% and 8%, respectively, looking at [kSZ](#) and [tSZ](#) in the [LAT](#) temperature maps, in combination with galaxy positions given in the [DESI](#) spectroscopic survey.

3.3.6 Duration of Reionization

As discussed in section [1.8.5](#), the epoch of reionization is not well constrained—the best constraint we have on its duration is an upper limit $\Delta z < 2.8$ at the 95% CL, from the combination of Planck and South Pole Telescope ([SPT](#)) data [[103](#)].

[SO](#) expects to be among the first to make real measurements of the duration of reionization, and is forecast to reach a constraint of $\sigma(\Delta z) = 0.4$, with a fiducial model of $\Delta z = 1$. This will be done by looking for excess variance in the temperature anisotropies generated by the [kSZ](#), as seen in the [LAT](#) temperature and polarization maps. Indeed, [SO](#) may achieve much better sensitivity to Δz , since this forecast implements a maximum multipole cut-off of $l_{\text{max}} = 3000$ which will likely be exceeded after foreground removal.

As secondary science goals, other parameters which [SO](#) will be able to constrain include the ionization efficiency ζ , and the mean free path of ionizing photons λ_{mfp} .

3.3.7 Secondary Science Goals

[SO](#) has a number of additional secondary science goals. Some of these link to the six primary goals, and are discussed in the relevant sections above. The remainder are discussed here.

These additional goals include: constraining Big Bang nucleosynthesis ([BBN](#)) through the primordial helium fraction Y_P for varying values of N_{eff} ; measuring the interaction cross-section between dark matter ([DM](#)) and baryons, σ_p ; constraining the fraction of ultra-light-axion dark matter, Ω_a/Ω_d ; and providing external validation of the shear bias, $m_{z=1}$, for [LSST](#).

By virtue of its high cadence scanning strategies (see Chapters 8 and 9), **SO** will also enable sub-millimetre observations of the Galaxy, Galactic and extragalactic sources, as well as transient sources. **SO** will produce daily transient maps, and release individual source light curves to the astrophysics community within 30 hours of detection.

Transient sources which are known to be visible at sub-millimetre wavelengths include fast blue optical transients (**FBOTs**), gamma-ray bursts (**GRBs**), supernovae, and variable sources like **AGNs**. **SO**'s blind time-domain survey approach also offers the potential to discover new classes of sub-millimetre transient sources.

SO will also be able to perform follow-up studies of **GW** detections. These observations are complementary to optical and radio transient surveys, and **SO** will add an additional dimension in its contribution of polarization observations [104].

3.3.8 Legacy Catalogues

SO has a broader aim of producing publicly available maps and catalogues for wider astrophysics and cosmology work.

SO will provide **CMB** maps (of temperature and polarization, at each of the six frequencies probed) from the **LAT** and **SATs** at their respective sensitivities and resolutions, as well as **CMB** lensing coverage and potentials.

In combination with Planck maps at lower ℓ values, and at higher frequencies, **SO** will also generate component-separated maps. These will again include maps of **CMB** temperature and polarization, as well as maps of the cosmic infrared background (**CIB**), **SZ**, and Galactic dust and synchrotron emission.

Catalogues of an expected 16,000 galaxy clusters, 15,000 **AGNs** and 8,500 dusty star-forming galaxies will also be produced.

3.4 Summary and Discussion

During its first observing run, **SO** will provide significantly better constraints on a range of cosmological parameters, which will only improve with the subsequent upgrades that **ASO**, along with **SO:UK** and **SO:JP**, will bring.

In one form or another, the sensitivities of **SO** to all of these parameters depend on the **LAT** and **SAT** scanning strategies. This was the impetus behind the campaigns of work detailed in Chapters 8 and 9. Indeed the possibility for substantial galactic and transient science is new

to ground-based [CMB](#) experiments, with the introduction of high-cadence scans.

[SO](#) will make significant contributions to our understanding of the primordial [GWB](#) at ultra-low frequencies, with both the [LAT](#) and [SATs](#). See section [5.3.3](#) for more detailed discussion of how measurements of the [CMB](#) can be used to measure (primordial contributions to) the stochastic gravitational wave background ([SGWB](#)).

Chapter 4

The Stochastic Gravitational Wave Background

This chapter discusses the stochastic gravitational wave background (SGWB) in more detail, focusing on its astrophysical foregrounds—produced by inspiralling binary populations—and a range of theorized cosmological sources including inflation, PBHs, cosmic strings, and FOPTs. It explores the spectral shapes associated with each source across relevant frequency bands. Significant portions draw from my forecasting work, published in reference [105].

4.1 Introduction

Recall from section 1.9 that the stochastic gravitational wave background (SGWB) is the background of low-level gravitational waves (GWs) generated by high energy cosmological and astrophysical phenomena.

In July 2023, the first confirmed detections of the SGWB were published [49, 50, 51, 52, 53], with the tightest constraints coming from the NANOGrav collaboration’s 15-year data set [49, 50]. We expect additional detections over a range of frequencies from other experiments in the upcoming decades. Additional constraints from the cosmic microwave background (CMB) can be made at lower frequencies—see discussion in section 5.3.3.

The SGWB observed in the pulsar timing array (PTA) frequency range ($\sim 10^{-9}$ Hz) could be the result of any number of cosmological phenomena, an astrophysical foreground of inspiralling supermassive black hole binaries (SMBHBs) [106], or indeed some combination. Cosmological sources include inflation (both immediate tensor perturbations [107], and GWs associated with primordial black holes (PBHs) [108, 109]), first order phase transitions (FOPTs) [110, 111, 112],

cosmic strings [113], and other more exotic phenomena including a variety of dark matter candidates [114, 115, 116]. At other frequencies, we could also see contributions from any of these cosmological phenomena, and foregrounds from binary populations of different masses of compact objects. Some of the cosmological phenomena which have been proposed as sources for the SGWB are contained within the Λ CDM model of cosmology (e.g. inflation), while others extend it (e.g. PBHs), and some entail entirely new physics (e.g. cosmic superstrings). As such, observations of the SGWB over a wide range of frequencies will have implications for our models of cosmology.

There are multiple avenues for observing SGWB sources, in addition to making direct observations of the GW spectra: we can constrain some individual phenomena through other physical signals; whatever SGWB signal was propagating at recombination will leave imprints on the B -mode polarization of the CMB; and present day GW foregrounds will affect the CMB temperature power spectrum (discussed further in section 5.3.3). But GW experiments remain our most comprehensive method of observing the whole SGWB spectrum across a wide frequency range.

The strength of the SGWB can be described by its energy density

$$\Omega_{\text{GW}}(f) \equiv \frac{1}{\rho_c} \frac{d\rho_{\text{GW}}}{d \ln f}, \quad (4.1)$$

where $\rho_{\text{GW}} \propto h^2$ is the energy density in GWs and ρ_c is the critical density [117].

The energy density of can be converted to other measures:

$$\Omega_{\text{GW}}(f) = \frac{2\pi^2}{3H_0^2} f^3 S_h(f) = \frac{2\pi^2}{3H_0^2} f^2 [h_c(f)]^2 = \frac{8\pi^2}{3H_0^2} f^4 |\tilde{h}(f)|^2, \quad (4.2)$$

where S_h is the (one-sided) power spectral density (PSD), $\tilde{h}(f) = \sqrt{|\tilde{h}_+|^2 + |\tilde{h}_\times|^2}$ is the Fourier transform of the time-domain signal strain $h(t)$, and $h_c(f) \equiv 4f^2 |\tilde{h}(f)|^2$ is the characteristic strain [118, 119]. S_h and $\tilde{h}$ have units of Hz^{-1} , while Ω_{GW} and h_c are dimensionless.

4.2 Astrophysical Foregrounds

In this thesis we consider the GW spectra of inspiralling binary populations to be a foreground contaminant for those of cosmological phenomena. Different binary masses contribute at different frequencies. At PTA frequencies the relevant binary systems are SMBHBs, at space-based interferometer frequencies white dwarf (WD) and stellar origin black hole (SOBH) binaries con-

tribute, and binaries containing [SOBHs](#) and neutron stars ([NSs](#)) contribute at ground-based interferometer frequencies—see figure [5.1](#).

The [GW](#) strain h_c we observe from a single binary at a given frequency f depends only on the distance to the system D , and chirp mass $\mathcal{M} = \frac{(m_1 m_2)^{3/5}}{(m_1 + m_2)^{1/5}}$ of the system:

$$h_c(f) = \frac{2(G\mathcal{M})^{5/3} \pi^{2/3} f^{-2/3}}{c^4 D}, \quad (4.3)$$

We can move from h_c to the energy density using equation [4.2](#) and then write the power spectrum for a single binary as a power law:

$$\Omega_{\text{GW}} \propto \frac{\mathcal{M}^{10/3}}{D^2} \times f^{2/3}. \quad (4.4)$$

There is no a priori lower limit on the frequencies emitted by each population, but a rough upper limit can be found by assuming a regime where both objects in a binary are spinless and of equal mass—that is, ignoring extreme mass ratio inspiral ([EMRI](#)).*

Taking Kepler’s third law, $\omega^2 \propto M^2/a^3$, where ω is orbital velocity, M is the mass of the object being orbited, and a is the orbital separation: we can input the *minimum* orbital separation to establish the *maximum* possible orbital frequency. For compact objects ([NSs](#) and [WDs](#)), $a^{\min}$ is simply twice their radius, while for black holes, the Schwarzschild radius $r_s = 2GM/c^2$ (the radius of the event horizon for a point mass M) can be substituted. Finally noting that the frequency of the associated [GWs](#) is twice the orbital frequency, we are left with

$$f_{\text{GW}}^{\max} \propto \omega^{\max} \propto \frac{M^{1/2}}{R^{3/2}} \quad (\text{compact objects}), \quad (4.5)$$

$$f_{\text{GW}}^{\max} \propto \omega^{\max} \propto \frac{1}{M} \quad (\text{black holes}). \quad (4.6)$$

The combined power spectrum for a population of inspiralling binaries is a superposition of each individual binary (retaining the same overall power law spectral shape as that of an individual binary) and as such can be calculated from a population mass function.

*[EMRI](#) commonly lead to highly eccentric orbits [[120](#)][—eccentricity in binaries boosts \[GWs\]\(#\) at higher frequencies \[\[121\]\(#\)\], while extreme eccentricity \(\$\gtrsim 0.9\$ \) results in attenuation of \[GW\]\(#\) strength across the frequency range \[\[122\]\(#\)\].](#)

4.3 Cosmological Sources

A variety of cosmological phenomena which might contribute to the [SGWB](#) has been suggested, spanning a huge range of frequencies, amplitudes, and spectral shapes—see [figure 4.1](#), which plots the spectra of different sources chosen to demonstrate this point. Sources which have been proposed include inflation (see [section 4.3.1](#)), [PBHs](#) ([section 4.3.2](#)), cosmic strings ([section 4.3.3](#)), [FOPTs](#) ([section 4.3.4](#)), and exotic phenomena including cosmic superstrings and axions/axion-like particles ([ALPs](#)) ([section 4.3.5](#)). The spectra of these cosmological sources can be modelled as single or broken power laws, as can be seen in [figure 4.1](#). The precise connections between the physical parameters of a phenomenon and the spectral parameters of the power laws are complex, but there are correlations between the amplitude, peak and slope(s) of the power law and the physical parameters of each phenomenon.

In all cases, evaluating the sensitivity of experiment to spectral parameters enables us to trace back sensitivities to physical parameters, either analytically or statistically.

4.3.1 Inflation

The same amplification of quantum fluctuations during a period of accelerating inflationary expansion which generates anisotropic structure across the Universe via scalar fluctuations also generates [GWs](#) through the magnification of tensor perturbations. These tensor perturbations then contribute a background of [GWs](#) to the [SGWB](#) [[107](#)]. The resulting [GW](#) spectrum is expected to be broad across frequencies, though the exact shape depends on the precise model of inflation involved.

We describe the primordial spectrum of tensor perturbations which arises from inflation as a power-law:

$$P_t(k) = r A_s(k_*) \left(\frac{k}{k_*} \right)^{n_T}, \quad (4.7)$$

where $r = P_t/P_s$ is the tensor-to-scalar ratio (P_s is the scalar power spectrum), A_s is the amplitude of primordial *scalar* perturbations at a given scale $k = k_*$, and n_T is the tensor spectral index.

Different theories predict their own values of r and n_T , both of which we have yet to properly constrain. Data from the most recent Planck release gives us an upper bound on $r < 0.06$ [[22](#)]. A tighter bound has been derived from the combination of the Planck results and the [BICEP2/Keck 2015](#) data to give an upper limit of $r < 0.044$ [[93](#)].[†] As discussed in

[†]Note that this calculation was done under the assumption of single field slow roll ([SFSR](#)) inflation, which requires the consistency relation $n_T = -r/8$. As such, it is not fully generalizable, and in some cases the

section 3.3.1, Simons Observatory (SO) expects to measure r with a sensitivity of $\sigma(r) \sim 0.002$.

To observe these primordial tensor fluctuations as red-shifted GWs today, we measure the present GW energy density Ω_{GW} . Assuming the standard radiation-, matter-, then dark-energy-dominated eras, and at PTA frequency, this can be approximated as [123]

$$\Omega_{\text{GW}} \propto \frac{\Omega_m^2 r A_s}{H_0^2 k_{\text{eq}}^2} \left(\frac{k}{k_*} \right)^{n_T} \propto \left(\frac{f}{f_*} \right)^{n_T}, \quad (4.8)$$

where $k_{\text{eq}} \approx 0.073 \Omega_m h^2 \text{ Mpc}^{-1}$ is the matter-equality wavenumber, $k_* = 0.05 \text{ Mpc}^{-1}$ is the reference pivot scale, corresponding to a reference frequency $f_* \approx 10^{-16} \text{ Hz}$. Hence the GW power scales with r , with its shape dictated by n_T .

4.3.2 Primordial Black Holes

PBHs are black holes which might develop during the very early stages of the universe. PBHs would predominantly develop before matter-domination, and so do not stem from the collapse of massive objects (unlike astrophysical black holes, which are born of stellar collapse);[‡] this means that they can be exceedingly small—as small as 10^{-8} kg at their inception. But any PBHs extant today must have been at least $\sim 10^{11} \text{ kg}$ initially, since black holes lose mass through Hawking radiation at a rate inversely proportional to their mass, and any lighter black holes would have evaporated by today [125].

A PBH forms when the density contrast of a spherical region as it reenters the horizon, $\delta \equiv \delta\rho/\rho$, surpasses some critical value δ_c . For a radiation dominated era, this could be somewhere between $\delta_c = 1/3$ and 0.8 , dependent on the shape of the given density perturbation—for a Gaussian perturbation, $\delta_c \approx 0.45$, while a top-hat perturbation has $\delta_c \approx 0.7$ (see, e.g., references [126, 127] for analysis on the effects of different shaped peaks). For a matter dominated era, $\delta_c \propto a$, and we hence predict a much higher rate of PBH formation in the very early universe.

The mass of the PBH is approximately that of the horizon mass at its formation. In a radiation dominated era, this is proportional to the age of the universe t :

$$M_{\text{PBH}} \approx M_{\text{Horizon}} \approx 10^{15} \text{ g} \left(\frac{t}{10^{-23} \text{ s}} \right). \quad (4.9)$$

There are multiple proposed mechanisms for generating the density contrasts which lead to the formation of PBHs. These include topological defects, and very specific types of FOPT. The

previous constraint $r < 0.06$ is more appropriate.

[‡]Although supermassive black holes (SMBHs) may be generated by direct collapse, rather than stellar collapse [124].

most well-motivated mechanism is inflation, as a natural consequence of scalar perturbations. The precise nature and evolution of the density contrast involved will depend on the inflationary model.

At linear order, the tensor and scalar modes of a perturbation are independent. However, at higher orders, this is not true—the same scalar perturbations which lead to PBHs also induce GWs.[§] It should be noted that the PBHs themselves are not generating the GWs; rather, both phenomena are consequences of the same initial scalar perturbations.

In the same way that PBHs form when the relevant spherical region reenters the horizon, so too do GWs only begin to propagate when their wavelength reenters the horizon. Thus, we can explicitly link the frequency of the GWs not only to the time of their formation via the horizon size, but also to the mass of the coeval PBHs. The present day frequency is then approximately [128]:

$$f_{\text{GW}} \sim 10^{-9} \text{ Hz} \left(\frac{M_{\text{PBH}}}{30M_{\odot}} \right)^{-1/2}. \quad (4.10)$$

Increasing the mass of the PBHs shifts the peak of the GW spectrum to lower frequencies. The amplitude of the spectrum is increased by increasing f_{PBH} , the proportion of dark matter attributed to PBHs.

We can also constrain PBHs by considering, e.g., Hawking radiation, positron interactions, micro-lensing, and distortions to the CMB through Silk damping [129]. The constraints on the abundance of PBHs at various mass ranges from a variety of experiments are explored in [130].

4.3.3 Cosmic Strings

Cosmic strings are theoretical linear topological defects produced in the early universe, whose presence and properties depend on the symmetries of the fields in the underlying model of particle physics. They form kinks and cusps as they vibrate, leading to the production of separate loops, which produce GWs as they decay. A population of cosmic strings would then generate a gravitational wave background (GWB) [13]. These cosmic strings would be produced following a symmetry breaking event [113] at the end of inflation [131]. Strings are parameterized by their tension $G\mu$, where μ is their mass per unit length. In units where $c = 1$, $G\mu$ is dimensionless.

The peak amplitude of the resulting GW power spectrum is proportional to $\Gamma \times (G\mu)^2$, where Γ is the number density of loops [132]. Cosmic string spectra are almost completely flat

[§]This also means that the original tensor modes which generate GWs will induce additional density perturbations. From upper limits we have on PBH abundance, this also constrains the amplitude of the tensor perturbations.

across the whole GW experiment frequency range—see figure 4.1.

We have upper bounds on the string tension from null results at various frequencies. Upper limits from LIGO/Virgo restrict the tension to $G\mu < 10^{-8}$ [133], and measurements of the CMB require $G\mu < 10^{-7}$ [134]. A combination of these constraints with additional factors from baryon acoustic oscillations (BAOs) and gravitational lensing requires that $G\mu < 4 \times 10^{-9}$ [135]. Reference [136] considers various models of cosmic strings which satisfy these constraints.

4.3.4 First-Order Phase Transitions

Cosmological phase transitions relate to transitions which alter properties at a universal scale. Typically these involve symmetry breaking events, which occur when the temperature of the universe drops below some critical temperature [137].

Phase transitions are characterized by their order: second-order phase transitions shift the universe from one state to the next in a continuous way. FOPTs, which have discontinuous change, proceed via bubble nucleation—bubbles of the new state nucleate stochastically across the volume, expanding and merging until the phase transition is complete [138].

The standard model of particle physics predicts second-order cosmological phase transitions for the QCD and electroweak phase transitions, but only FOPTs generate GWs. An observation of GWs from FOPTs would tell us both about the history of the Universe, but also requires new physics [139]. FOPTs for a range of cosmological phase transitions have been proposed as contributing sources of the SGWB, including the electro-weak and QCD phase transitions.

FOPTs can be characterized by four parameters: their strength α , the rate of bubble nucleation Γ , the temperature when bubble nucleation begins T_n , and the speed of the expansion of the bubbles (the wall-velocity) v_{wall} . There is not a simple analytical model to describe the GW spectrum from a FOPT, nor are all FOPTs similar enough to have a single model. Depending on the four parameters, the transition (and hence also GW spectrum) is qualitatively different.

There are three elements of FOPTs which generate GWs: bubble collisions, sound waves, and turbulence [137]. GW spectra from FOPTs have to be simulated, since it is impossible to derive their spectra analytically (although semi-analytic treatments have been carried out alongside simulations). The resulting spectra depend on α , Γ , T_n , and v_{wall} , but strongly depend on the dominant effect—bubble collisions dominate only where there is very strong supercooling; sound waves for weak phase transitions; and turbulence for strong.

Bubble Collisions

Bubble nucleation occurs in some background of cosmic fluid, and there is interaction between the expanding bubbles and this background fluid. This generates both sound waves and turbulence. We separate this effect from the bubble collisions in order to deal with bubble collisions (semi-)analytically.[¶]

Bubble collisions are a violent process, and the energy involved is sufficient to generate observable GWs [110]. While it is possible to evaluate the GW spectra from two colliding bubbles entirely analytically with scalar field analysis by relying on the symmetry of the system, the same cannot be done for multiple bubbles. To this end, the Envelope Approximation was developed [141]; it has been shown to be accurate by a 3D scalar field simulation [142], and is widely used when dealing with FOPTs. The Envelope Approximation mirrors the findings of the original scalar field analysis, such that the GW spectrum depends only on the vacuum energy and radii of the colliding bubbles [141].

We consider only the ‘envelope’ of bubble surfaces around the outside of a bubble system. That is, we ignore sections of bubble surfaces which lie within other bubbles. We also assume infinitesimally thin bubble walls, motivated by the near-light speeds of expansions which cause Lorentz contraction. There are alternatives to the Envelope Approximation—see reference [143], for example.

Sound Waves

The effects of sound waves caused by the interaction of expanding bubbles with the background fluid are too complex to be determined analytically, and are instead simulated.[‡]

The qualitative properties of the sound waves can be determined by comparison of v_{wall} with c_{sound} in the fluid—where $v_{\text{wall}} > c_{\text{sound}}$ you have a detonation (the sound wave shell lags behind the expanding bubble shell), otherwise you have a deflagration (where the sound shell is pushed along in front of the bubble shell) [144].

[¶]Actually, the interaction between bubble and background has an effect on the bubble wall speed v_{wall} . While the literature commonly uses a constant $R_{\text{wall}} = t \times \text{constant}$, the effect of friction suggests that $R_{\text{wall}} \approx \sqrt{t^2 + \text{constant}}$ would be more accurate. This tends to the constant expansion rate for large t . R^3 is the important factor for the GW spectra, and so this does have an effect on the generated spectra [140].

[‡]Although reference [110] attempts a semi-analytic treatment alongside a 3D simulation involving both the scalar field and the compression waves.

Turbulence

Continued interactions between the expanding bubbles and the sound waves propagating in the fluid creates stochastic turbulence.

Turbulence generates GWs for much longer than either other process—sometimes up to many Hubble times** after the phase transition has completed [145]. Carrying out full scale simulations for such long periods are computationally expensive, and so approximations such as the Bulk Flow Model are often used [146].

General GW Spectra

Where bubble collisions dominate, the spectrum is largely dictated by the average separation of bubbles r_* (which can be derived from Γ and v_{wall}); here an increase in r_* increases the amplitude of the spectrum. Where sound waves dominate, increasing r_* or α increases the amplitude and shifts the peak to lower frequencies; increasing T_n increases the peak frequency; and increasing v_{wall} slightly decreases the amplitude while modifying the shape of the spectrum slightly. Describing spectra from turbulence qualitatively is more complicated, but primarily the amplitude and frequency of the peak increase as α increases.

4.3.5 Exotic Phenomena

Other phenomena which might contribute to a SGWB have been proposed, including cosmic superstrings and axions/ALPs.

Cosmic Superstrings

Under string theory, cosmic superstrings would be generated by the expansion of fundamental superstrings via inflation. They would generate GWs in the same manner as cosmic strings, discussed in section 4.3.3 [147, 148], and would therefore have spectra of a similar shape. Unlike a detection of cosmic strings, a detection of superstrings would be evidence for string theory [149]. Observational upper limits on the SGWB place bounds on the superstring tension, for various reconnection probabilities $p^{\dagger\dagger}$: $G\mu < 10^{-12}$ for $p \approx 10^{-3}$, $G\mu \approx 10^{-10}$ for $p \approx 10^{-1}$, and $G\mu \approx 10^{-8}$ for $p \approx 1$ [152].

**Where in this context, the Hubble time $t_H = 1/H(t)$ gives a characteristic time scale for the duration of FOPTs.

$^{\dagger\dagger}p$ refers to the probability that two strings will reconnect if they touch. Cosmic strings have $p = 1$ as standard, while cosmic superstrings typically have much lower reconnection probabilities. This can lead to strong SGWB signals [150, 151], and the different values of p are a good way to distinguish cosmic superstring spectra from those of cosmic strings.

Axions and Axion-like Particles

Axions are theorized low mass ($\sim 10^{-5}$ – 10^{-3} eV) spin- and chargeless particles, which follow from the Peccei-Quinn solution to the strong CP problem^{‡‡} [153, 154], which could also act as cold dark matter [155]. Axion-like particles are similar in mass, and largely arise with the introduction of string theory to the standard model [156]. The oscillations of an ALP field in the early universe would generate a GWB at some level [157], with specifics of the spectral shape depending on the precise model.

4.4 Summary and Discussion

This chapter has discussed a non-exhaustive list of cosmological phenomena which could contribute to the SGWB. Their different spectral behaviours can be seen in figure 4.1, with qualitative links between their spectra and the physics behind them discussed in the text.

Of the phenomena discussed here, some fit into Λ CDM cosmology, while others are considered new physics. Inflation in some form is part of Λ CDM, and the PBHs it might generate do not cause any contradiction to the model. Likewise, cosmic strings, axions and ALPs do not form part of Λ CDM, but do not necessarily contradict it. However, FOPTs do directly contradict the second-order phase transitions specified by the standard particle physics model. And an observation of cosmic superstrings would provide evidence for string theory.

Ultimately we expect to find a SGWB which is a combination of the spectra from multiple phenomena, plus the astrophysical foreground, which will be complicated to separate out. Some work has been done on template-free separation of cosmological signals from astrophysical foregrounds and transient sources (see, e.g., [158, 159, 160]), but it is still in early stages.

It is worth noting that some of these phenomena could also be observed and/or constrained through other effects. Indeed this is the driving force behind the work detailed in Chapters 6 and 7, which forecast the sensitivities of experiments to different the spectra of different phenomena, and therefore lays out which phenomena it would be best to constrain through other means—given that there are other ways of looking at (some of) these phenomena, it’s useful to establish how much we will be able to constrain parameters with SGWB observations alone (in advance of getting them).

^{‡‡}Which asks why QCD does not appear to violate the combined symmetry of charge conjugation (C) and parity (P), despite it being mathematically possible.

Chapter 5

Observing the SGWB

This chapter reviews the observational history of [GWs](#) and the [SGWB](#), covering direct detections and upper limits from [GW](#) detectors, as well as indirect constraints from the [CMB](#). It discusses the current generation of [GW](#) detectors and outlines plans for future detectors, noting the gap in frequency coverage which remains. Significant portions of the text are drawn from my paper, reference [\[105\]](#).

5.1 Introduction

Since 2023, we have had confirmed detections of the stochastic gravitational wave background ([SGWB](#)) at low frequencies [\[49, 50, 51, 52, 53\]](#). Ultimately we hope to move from these initial narrow frequency observations of the [SGWB](#) to a wider spectrum. Current and planned gravitational wave ([GW](#)) experiments span more than 12 decades of frequency. Figure [5.1](#) shows the current and predicted sensitivities of the [GW](#) experiments considered in this paper, and shows the large gap in observable frequency between pulsar timing arrays ([PTAs](#)) and space-based detectors; a range of possible intermediate-frequency experiments have been suggested (see, e.g., references [\[161, 162, 163\]](#) for discussions) but none are yet expected to be implemented. Many of these experiments will not come online for over a decade, and it is therefore worth knowing in advance which are likely to be useful for observing which phenomena *before* we have access to their results. This allows us to establish, for example, whether we should wait for [GW](#) results or prioritize looking for additional physical constraints (see relevant subsections of section [4.3](#)) in the interim. That we should model sensitivities in advance of experiments coming online, is the motivation for the work discussed in Chapters [6](#) and [7](#).

5.2 Gravitational Wave Experiments

All GW detectors work in fundamentally the same way—exploiting the constancy of the speed of light to determine whether a given distance changes over time. Recall from section 1.3 that a GW stretches or squeezes space itself as it travels. When a GW passes through a detector, the ‘arm’ of a detector is stretched (squeezed) by a tiny amount, leading to light passing along the arm reaching the end slightly later (earlier) than expected. Different sizes of GW detector are sensitive to different frequencies of GWs—the larger the arm length of the detector, the lower the frequencies it can reach, as can be seen in figure 5.1. More details and references for each experiment considered in this thesis are given below, and summarized in table 5.1.

At mid- and high-frequencies, we use ground- and space-based interferometers which rely on lasers, while at low frequencies we use PTAs, which exploit the properties of pulsars. At ultra-low frequencies we can use cosmic microwave background (CMB) experiments to search for the impressions of the primordial component of the SGWB on the temperature and polarization signals of the CMB. As can be seen in figure 5.1, there is a gap in observing frequencies over the range $\sim 10^{-7.5}$ – $10^{-5.5}$ Hz.

Laser interferometers are based on the 1887 Michelson-Morley set-up [167]—two perpendicular ‘arms’ of equal length, with a beam-splitter set up such that a beam of coherent light is split along each arm, before being reflected back by mirrors. If the recombined beam is still coherent, the path travelled by each half-beam has remained equal in length. If the light instead arrives incoherent, there has been an effective change in the length of one of the arms caused by a GW passing through and stretching/squeezing space through its perturbation of the metric. Analysis of the exact phases of each beam can give us precise changes in length, and if we can remove noise effects, we can deduce the presence of GWs passing through. Modern interferometers can have either perpendicular or triangular setups.

Pulsars are neutron stars with strong magnetic fields, which are also spinning with very short periods. If the rotational and magnetic axes are distinct, we can observe pulses of electromagnetic radiation with (almost) perfectly consistent periods—once per rotation, when the jet is in the line of sight from Earth. Periods range from milliseconds to tens of seconds.

We can use the most consistent pulsars to measure waves in spacetime along the line of sight to the pulsar. A PTA takes regular observations of a group of pulsars, and analyzes their timing data. The best pulsars for observing the SGWB have periods < 30 ms [168]—these millisecond pulsars can show us GWs between 10^{-9} and 10^{-6} Hz, where the lower limit comes from total observing time of a few decades ($f_{\min} \approx 1/T_{\text{obs}}$), and the upper limit comes from the




Figure 5.1: Current results and eventual sensitivities for the experiments considered in this thesis. Dashed lines show the instantaneous noise curves Ω_{nominal} , detailed in Appendix B. Solid lines show the eventual noise curves Ω_{final} which take into account the improvement of sensitivity with observing time and the presence of any foregrounds. The cross hatched regions show the eventual sensitivities in the absence of any foregrounds—see section 6.4.1 for discussion of Ω_{nwt} (‘noise-with-time’ (nwt)) and its difference to standard power law integrated sensitivity curves (PLISCs). Note that Ω_{final} for SKA is shown after 5/10/20 years, and see section 6.4.1 for the issues with Ω_{nominal} for PTAs. The diagonally hatched region shows the upper limit given by LVK-O3’s lack of SGWB detection. The most recent NANOGrav results [49] are shown by the blue and purple shaded regions, and are discussed in more detail in figure 7.5a. Dotted lines show the predicted instantaneous astrophysical foregrounds from different binary populations, using equation 6.4 fitted to models for WD-WD [164], and BH-BH, NS-NS and BH-NS binaries [165, 166].

cadence of observations (which is on the order of fortnightly observations), as well as timing noise from the pulsars themselves.

Experiment	Freq. Range $\log_{10}(f)$	Arm Length (km)	Operational	T_{obs} (yr)	Ref(s).
Ground-Based Experiments:					
LVK (collab.)	[1.3, 3.6]	3–4	2019/23/27 [†]	1/2/3 [†]	[169]
CE	[0.7, 3.5]	40	2035–40*	5	[170]
ET	[1, 3.3]	20	2035*	5	[171 , 172]
Space-Based Experiments:					
LISA	[−5, −1]	2.5×10^6	2037*	4	[173]
DECIGO	[−3, 1]	1,000	2030s / 40s*	3	[174]
BBO	[−2, 1]	50,000	TBA*	10	[175 , 176]
Pulsar Timing Arrays:					
NANOGrav	[−8.75, −7.5]	N/A	2004	15+	[177]
IPTA (collab.)	[−9, −7.5]	N/A	2016 [‡]	12.5+	[178]
SKA	[−9, −7]	N/A	2028*	20	[179]

Table 5.1: Relevant properties of the [GW](#) experiments considered in this thesis.

* Expected to begin operations. [†] Operational phases 3/4/5, respectively. [‡] First data release.

5.2.1 Ground-Based Interferometers:

Current Interferometers:

Currently active ground-based detectors are all involved in the LIGO-Virgo-KAGRA ([LVK](#)) collaboration, which is in its fourth observing run (O4), scheduled to conclude on October 7, 2025. The fifth observing run (O5) is expected to start in late 2027.

[aLIGO](#) has been collecting data since 2015, as the successor to [LIGO](#). [LIGO](#) ran from 2002–10 and did not detect any [GWs](#), while [aLIGO](#) has made over 300 detections of individual [GW](#) events, but no detection of the [SGWB](#).

[aLIGO](#)’s two 4 km interferometers, in Livingston and Hanford, USA, are separated by ~ 3000 km. This leads to a small time delay between their observations of the same [GW](#) events, which can be used to determine the location of the [GW](#) source on the sky.

With a very similar set-up, [aVirgo](#), a 3 km interferometer in Italy, has been collecting data in conjunction with [aLIGO](#) since 2017. The initial [Virgo](#) detector ran between 2007–11, and did not make any [GW](#) detections.

[KAGRA](#), a 3 km interferometer located underground in the Kamioka mine, Japan, joined the [LVK](#) network in 2020. It is the first large-scale [GW](#) experiment built underground, and also the first to use cryogenically cooled mirrors, which are intended to improve its sensitivity,

especially at lower frequencies.

Planned Interferometers:

Future ground-based detectors aim to enhance sensitivity to GWs at high frequencies. These include upgrades to LVK, as well as new experiments including the Indian Initiative in Gravitational-Wave Observations (INDIGO), the Einstein Telescope (ET), and Cosmic Explorer (CE).

INDIGO is a proposed interferometer, based on the aLIGO design, to act as a third detector for the aLIGO network. Locating this third detector in India will greatly improve source localization, due to its separation from the two U.S. detectors, bringing an order of magnitude improvement on source localization.

CE is a proposed U.S. experiment with arm lengths of 40 km, while ET, a European project, will have arm lengths of 20 km. CE and ET are designed to work in conjunction with one another to increase the signal-to-noise and improve localization of individual sources on the sky [170, 180, 181].

5.2.2 Space-Based Interferometers:

Much larger, space-based interferometers will be constructed, with arm lengths up to scales of millions of kilometres, sensitive to mid-frequency GWs. Laser Interferometer Space Antenna (LISA) is one such project, while Deci-hertz Interferometer Gravitational wave Observatory (DECIGO) and Big Bang Observer (BBO) have also been proposed as additional space-based detectors.

All of these consist of a trio of spacecraft in an equilateral triangle formation, sitting at Lagrange points in the Earth's orbit.

LISA's demonstration mission, LISA Pathfinder, was run in 2015. It operated successfully in space for a full year, achieving better than required sensitivity levels for the pathfinder mission. LISA will have its craft at separations of 2.5×10^6 km, while DECIGO's will be at 10^3 km. It is likely that a smaller demonstration experiment, B-DECIGO, will be launched first, with craft separated by only 100 km.

BBO is a proposed experiment which would operate in a similar frequency band to DECIGO but with improved sensitivity. It follows a similar design to LISA, but with multiple overlapping detectors.

5.2.3 Pulsar Timing Arrays

Active [PTAs](#) include those in the International Pulsar Timing Array ([IPTA](#)) collaboration ([NANOGrav](#), the European Pulsar Timing Array ([EPTA](#)), the Parkes Pulsar Timing Array ([PPTA](#)), the Indian Pulsar Timing Array ([InPTA](#))) and the Five-hundred-metre Aperture Spherical radio Telescope ([FAST](#)). The Square Kilometre Array ([SKA](#)) will also act as a [PTA](#) when it comes online.

Current PTAs:

In June 2023, [PTAs](#) jointly announced strong evidence for a detection of the [SGWB](#), marking a significant milestone in [GW](#) astronomy [[49](#), [50](#), [51](#), [52](#), [53](#)].

The North American Nanohertz Observatory for Gravitational Waves ([NANOGrav](#)) originally worked with data from two telescopes: the Arecibo and Green Bank telescopes. Following the collapse of the Arecibo telescope in 2020, [NANOGrav](#) has had to expand to using other observatories, including [FAST](#) and ultimately [SKA](#).

[EPTA](#) uses a system of five ~ 100 m telescopes—Westerbork Synthesis Radio Telescope in the Netherlands; Effelsberg Telescope in Germany; Lovell Telescope in England; Nançay Radio Observatory in France; and Sardinia Radio Telescope in Italy. Combining the data, this gives [EPTA](#) the equivalent of a 200 m telescope, since the effective combined observing area of a [PTA](#) is approximately equivalent to the total observing area [[182](#)].

[PPTA](#) collects data from the Parkes Telescope in Australia. While it observes with by far the smallest telescope used for [PTA](#) astronomy, [PPTA](#) has several days of observation time each month, compared with [NANOGrav](#)'s ~ 1.5 days per month.

[FAST](#) in China has many science goals, and acting as a [PTA](#) is one of them. Since first light in 2016, it has discovered over 44 new pulsars and is expected to be able to carry out higher precision timing than any of the individual [PTAs](#) involved in [IPTA](#).

Upcoming PTAs

The [SKA](#) is under construction, with full operations expected to start before 2030—its SKA-Low telescope achieved first light in 2025.

The design of [SKA](#), with a total telescope observing area of ~ 1 square kilometer, uses multiple receivers spread over thousands of kilometers, a process called aperture synthesis. There will be core clusters in South Africa and Australia. [SKA](#) is expected to be 50 times more sensitive than current detectors, at frequencies all the way from 50 MHz–30 GHz. It should also

be able to scan the sky at incredibly fast rates. One of its main science goals is to test GR through pulsar timing.

5.3 History of SGWB Observations and Results

5.3.1 Observations from PTAs

In 2023, [NANOGrav](#)’s 15 year data releases (hereafter ‘NANOGrav15’) gave the tightest constraints on the first confirmed observation of the [SGWB](#) [49, 50]. This observation was corroborated by other [PTAs](#) in a simultaneous release of data: [EPTA](#) (including pulsars from [InPTA](#)) [51], [PPTA](#) [52], and the Chinese Pulsar Timing Array ([CPTA](#)) [53]. A combined analysis in the upcoming [IPTA](#) third data release will provide more sensitive constraints on the observations [50].

[PTA](#) data are typically released in terms of the constraints on a power-law spectrum with a reference frequency of $f_{\text{yr}} = \text{yr}^{-1}$:

$$h_c(f) = A_{\text{CP}} \left(\frac{f}{f_{\text{yr}}} \right)^\alpha, \quad (5.1)$$

(with A_{CP} denoting the amplitude of a power spectrum due to some ‘common process’) and the cross-power spectral density of the timing deviations between two pulsars:

$$S_{ab}(f) = \Gamma_{ab} \frac{A_{\text{CP}}^2}{12\pi^2} f_{\text{yr}}^{-3} \left(\frac{f}{f_{\text{yr}}} \right)^{-\gamma_{\text{CP}}}, \quad (5.2)$$

where Γ_{ab} is the Hellings-Downs correlation coefficient between the two pulsars [183]. In the case where pulsar timing residuals are due entirely to a common [GW](#) process, the two equations are linked such that $\alpha = (3 - \gamma_{\text{CP}})/2$, and the fit of results to the Hellings-Downs curve is therefore a test for the presence of the [SGWB](#). We can then obtain the [GW](#) energy density Ω_{GW} using equation 4.2:

$$\Omega_{\text{GW}}(f) = \frac{2\pi^2}{3H_0^2} f^3 S(f). \quad (5.3)$$

The [PTA](#) analysis results in table 5.2 are given in both joint A - γ posteriors, and A posteriors with a specified γ for supermassive black hole binaries ([SMBHBs](#)) (the dominant astrophysical foreground population at PTA frequencies)—per section 4.2 and from equation 4.4, we obtain a value $\gamma_{\text{SMBHBs}} = 13/3$. Note that the results given in table 5.2 do not show the adherence of the set of pairs of pulsars to the Hellings-Downs curve, but rather constraints on the total

power spectrum *given that* adherence has already been established. The NANOGrav15 results are also plotted in figure 5.1, in the novel ‘bowtie’ style which is described in Chapter 7.

	NANOGrav	EPTA		PPTA		CPTA	MeerKAT
$\log_{10} A$	$-14.61^{+0.12}_{-0.11}$	$-14.61^{+0.11}_{-0.15}$		$-14.69^{+0.05}_{-0.05}$		$-14.7^{+0.9}_{-1.9}$	$-14.28^{+0.21}_{-0.21}$
γ	$3.2^{+0.6}_{-0.6}$	$4.19^{+0.73}_{-0.63}$	$2.71^{+1.18}_{-0.73}$	3.87	4.02	$[0, 6.6]$	$3.60^{+1.31}_{-0.89}$
α	$-0.1^{+0.3}_{-0.3}$	$-0.60^{+0.32}_{-0.36}$	$0.15^{+0.36}_{-0.60}$	-0.44	-0.51	$[-1.8, 1.5]$	$-0.3^{+1.3}_{-0.45}$
$\log_{10} A$	$-14.19^{+0.22}_{-0.24}$	$-14.54^{+0.28}_{-0.41}$	$-13.94^{+0.23}_{-0.48}$	-14.5	-14.56	$-14.4^{+1.0}_{-2.8}$	$-14.25^{+0.21}_{-0.36}$

Table 5.2: Analysis results from each PTA given at the 95% confidence level (NANOGrav [50, 49], EPTA (including pulsars from InPTA) [51], PPTA [52], CPTA [53]), MPTA [184]. Values in the upper section consider a fixed value of $\gamma = 13/3$ ($\alpha = -2/3$), i.e., a population of SMBHBs. Values in the lower section are joint A - γ and A - α posteriors. Values in bold indicate which of γ and α the original joint posteriors were given in—conversions have been made for convenience of comparison. Where no errors are listed, errors were not provided in the original references. Ranges given in square brackets indicate uniform sampling across the range (insufficient observing time for constraints). The left column for EPTA shows its results for its full 24.7 year dataset, and the right shows just the 10.3 year subset which was carried out with modern instrumentation (both sets are in combination with InPTA data). The left column for PPTA relies only on pulsar auto-correlations, while the right uses both auto- and cross-correlations.

5.3.2 Upper Bounds from Ground-Based Detectors

The most up-to-date upper bounds from operational ground-based detectors comes from the LVK-O3 data release [185]. Table 5.3 gives the upper limits on the SGWB signal at a reference frequency of 25 Hz from that work, for three values of the power-law index:^{*} $m = 0$ (a flat spectrum, corresponding to, e.g., cosmic strings or single field slow roll (SFSR) inflationary models), $2/3$ (compact binary inspiralling population), and 3 (historically used value, e.g., some supernova, see [185] and references therein). Note that m is named α in the original LVK analyses; it has been changed here to avoid confusion with α used in Section 5.3.1, and for consistency with our results later.

^{*}NB this considers a power law at a much higher reference frequency than for PTAs, and there is no suggestion that the two should match.

	$\Omega_{\text{GW}} \times 10^{-8}$	
m	Uniform prior	Log-uniform prior
0	1.7	3.5
2/3	1.2	3.0
3	0.13	0.51

Table 5.3: Current upper limits from LVK-O3, at the 95% credible level [185]. Results are given for uniform and log-uniform priors on $\Omega_{\text{GW}}(f_{\text{ref}} = 25 \text{ Hz})$, where the former is more conservative and the latter is more sensitive to weak signals.

5.3.3 Constraints from the CMB

In CMB temperature observations, present-day GWs would induce an additional quadrupole in the CMB; GWs propagating at the formation of the CMB would have caused damping on the smallest angular scales.

GWs propagating at early time act to increase the effective number of relativistic species N_{eff} :

$$N_{\text{eff, GW}} = \int_{f_{\text{min}}}^{f_{\text{max}}} df \frac{\Omega_{\text{GW}}(f) h^2}{f}, \quad (5.4)$$

where here $h = H_0/100 \text{ km/s/Mpc}$ is the reduced Hubble constant. N_{eff} is constrained by the CMB, since increased N_{eff} causes observable damping of high- ℓ peaks in the CMB temperature power spectrum. N_{eff} also affects the abundance of light elements, and so we can also use Big Bang nucleosynthesis (BBN) to constrain it. Considered together, CMB+BBN requires that the additional contribution is $\Delta N_{\text{eff}} \lesssim 0.4$ [22].

Upper limits on ΔN_{eff} and r can be converted into constraints on Ω_{GW} . Assuming that $\Delta N_{\text{eff}} = N_{\text{GW}}$:

$$\Omega_{\text{GW}}^{N_{\text{eff}}} h^2 \approx 5.6 \times 10^{-6} N_{\text{GW}}, \quad (5.5)$$

at frequencies between 10^{-10} and 10^6 Hz . And

$$\Omega_{\text{GW}}^r(f) = \frac{3r A_s \Omega_r}{128} \left(\frac{f}{10^{-17} \text{ Hz}} \right)^{n_T} \left[\frac{1}{2} \left(\frac{f_{\text{eq}}}{f} \right)^2 + \frac{16}{9} \right], \quad (5.6)$$

where $f_{\text{eq}} = \sqrt{2} H_0 \Omega_m / 2\pi \sqrt{\Omega_r}$. Conversions often fix either the consistency relation $n_T = -r/8$, or $n_T = 0$. The predicted sensitivities of Simons Observatory (SO) to the SGWB can then be evaluated. First from N_{eff} : $\Omega_{\text{GW}}^{N_{\text{eff}}} \lesssim 8 \times 10^{-7}$, which improves by an order of magnitude on the current constraint, $\Omega_{\text{GW}}^{N_{\text{eff}}} \lesssim 5 \times 10^{-6}$.

Any gravitational wave background (GWB) already propagating by the time of CMB generation (i.e. at recombination) would introduce B -modes into an otherwise B -mode-free CMB [186]. Separating out B -modes from polarization foregrounds by objects between us and the CMB (dust emissions within the galaxy, and lensing by massive objects further away) should leave us with observable B -modes from GWs; we expect to detect these at angular scales of $\sim 0.1\text{--}1^\circ$ [186].

Since we have yet to detect the effect of a GWB in CMB measurements, the temperature, B -mode polarization and lensing from Planck and BICEP/Keck place an upper limit of $\Omega_{\text{GW}} \lesssim 10^{-12} (k/0.1 \text{ Mpc}^{-1})^3$ for $k \gtrsim 0.1 \text{ Mpc}^{-1}$ (corresponding to frequencies between roughly 10^{-16} – 10^{-14} Hz) [187]. The baseline goal for SO's constraint on r ($\sigma(r) = 0.002$) gives $\Omega_{\text{GW}}^r \lesssim 7 \times 10^{-13}$,

while the improved constraint from the new Small Aperture Telescope (SAT) scanning strategy would improve this by a factor of 1.2 (see Chapter 9).

5.4 Summary and Discussion

We currently have measurements of the SGWB over one small frequency range, with upper limits at other frequencies while we wait for more sensitive experiments to provide us with observations.

Currently running and upcoming detectors should dramatically increase our observations of GWs and the SGWB. When the space-based interferometers are launched, LISA's sensitivity should be $\Omega_{\text{GW}} \sim 5 \times 10^{-13}$ at 10^{-3} Hz [188], while DECIGO should be sensitive to $\Omega_{\text{GW}} \sim 2 \times 10^{-16}$ at 0.1–10 Hz [13].

In the longer term, ET and CE should be around 10 times more sensitive to the SGWB than current ground-based interferometers, which would also enable them to pick out individual binary signals down to stellar mass ranges, allowing the remaining cosmological signal to be more easily detected. Five years of detections should then give us $\Omega_{\text{GW}} \sim 10^{-13}$ at the same ~ 10 –100 Hz range as LVK [189].

When PTAs are at their next stage, it should be possible to begin the extrication of specific SMBHB foreground signals from the underlying SGWB: IPTA is expected to reach the right level of sensitivity to pull out 3 SMBHB signals by the end of 2025, and 13 by 2030 [190]. More signals will be isolated as SKA subsequently begins observations.

We have upper limits from the CMB via Planck, but more sensitive experiments (e.g. SO) aim to provide detections or much tighter upper limits on the SGWB at ultra-low frequencies.

Piecing together constraints on the proposed phenomenological sources of the NANOGrav signal, it is immediately clear that neither specific models of source, nor even the type of source, can be pinpointed yet. We will need more and better data, at PTA frequencies, and at higher and lower frequencies. This will begin to appear over the next few years, and within the next decade, it should become possible to place much tighter constraints on the sources behind the SGWB.

Part II

Forecasting Experimental Sensitivities to the SGWB

Chapter 6

Modelling the SGWB

This chapter outlines the methodology for modelling the SGWB and forecasting the sensitivity of GW detectors to its possible sources. It describes various parametrizations of SGWB spectra, the incorporation of astrophysical foregrounds—with an adapted and generalized fitting function—and the treatment of instrumental noise. It introduces Fisher matrix techniques for evaluating the sensitivities of experiments (and combinations thereof) to fiducial SGWB spectra. Significant portions of the text are drawn from my forecasting papers, references [105, 191].

6.1 Introduction

Our overall goal is to determine the constraining powers of current and future gravitational wave (GW) experiments, individually and in combination, on different sources of the stochastic gravitational wave background (SGWB). We model the spectra of SGWB sources as simple functions, in the presence of astrophysical foregrounds.

6.2 Instrumental Noise

The power spectral density (PSD) of a single detector $S_{\text{ins}}(f)$ is

$$S_{\text{ins}}(f) \equiv \frac{N(f)}{\mathcal{R}(f)}, \quad (6.1)$$

where $N(f)$ is the instrumental noise and the expression defines $\mathcal{R}(f)$, the detector response function, such that S_{ins} is the PSD required to produce a signal equal to the instrumental noise. A value of $\mathcal{R}(f) = 1$ denotes an ideal experiment.

Recall from equation 4.2 that the energy density of the SGWB, Ω_{GW} , can be given in

alternative measures:

$$\Omega_{\text{GW}}(f) = \frac{2\pi^2}{3H_0^2} f^3 S_h(f) = \frac{2\pi^2}{3H_0^2} f^2 [h_c(f)]^2 = \frac{8\pi^2}{3H_0^2} f^4 |\tilde{h}(f)|^2, \quad (6.2)$$

The noise spectrum of an instrument can be defined similarly. $S_{\text{ins}}^{\text{eff}} = S_{\text{ins}}/\sqrt{n_{\text{det}}}$ is the effective noise PSD of a network of n_{det} detectors, and Ω_{ins} is then $S_{\text{ins}}^{\text{eff}}$ expressed in terms of an energy density.*

$$H_0^2 \Omega_{\text{ins}}(f) = \frac{2\pi^2}{3} f^3 S_{\text{ins}}^{\text{eff}}(f) = \frac{2\pi^2}{3} f^2 [h_{\text{ins}}(f)]^2 = \frac{8\pi^2}{3} f^4 |\tilde{n}(f)|^2, \quad (6.3)$$

where $\tilde{n}(f)$ is the Fourier transform of the time-domain noise strain $n(t)$, and $h_{\text{ins}}(f) \equiv 4f^2 |\tilde{n}(f)|^2$ is the noise strain [118].

6.3 Astrophysical Foregrounds—Binary Systems

Astrophysical foregrounds at different frequencies are generated by inspiralling binary populations of different masses (see discussion in section 4.2).

We model the different populations of binaries with one generalized function:

$$\Omega_{\text{fg}} = A f^{2/3} \left[\exp\left(-(f/f_{\text{ref}})^4\right) + D \sin^2(\pi f/f_{\text{ref}}) \right] \times [1 + \tanh(E(f_* - f))], \quad (6.4)$$

where a set of coefficients A , D , E , f_* , and f_{ref} for each population is given in Appendix A. The coefficients are found through a least-squares fitting 6.4 to the models specified in Appendix A. Equation 6.4 is adapted from a similar white dwarf (WD) binary model [192], which I developed to provide a more general fit across different populations.

A more comprehensive spectral shape can be constructed with the inclusion of merger and ringdown spectra, which is detailed in [193]. This affects the shape of the spectrum to the right of the merger frequency (i.e., to the right of the cutoff for inspiral signals), an approximation of which frequency is given in section 4.2. For many of the experiments considered here this behaviour is below the relevant sensitivities. See reference [194] for examples of how the spectral shape changes around the peak, and references [195, 196] for more detail on extended spectra of supermassive black hole binaries (SMBHBs) which might be visible in the LISA band.

*Note that some papers work with $\Omega_{\text{ins}}(S_{\text{ins}})$ rather than $\Omega_{\text{ins}}(S_{\text{ins}}^{\text{eff}})$, and care should be taken with factors of n_{det} .

The total observed power is therefore given by

$$\begin{aligned}\Omega_{\text{total}} &= \Omega_{\text{GW}} + \Omega_{\text{noise}} \\ &= \Omega_{\text{GW}} + \Omega_{\text{ins}} + \Omega_{\text{fg}}.\end{aligned}\tag{6.5}$$

In the case where multiple sources contribute to the [SGWB](#) at a given frequency, Ω_{GW} can be taken to refer solely to a source of interest, with the contributions of the remaining source(s) contained within the overall Ω_{noise} , which is taken to be known and fixed. Alternatively, Ω_{GW} can be taken as the combined power of the sources, in which case we must fit for the parameters of each such source, but this comes at the cost of lower sensitivity to, or even degeneracies between, those parameters.

6.4 Fisher Matrices and Sensitivities to the SGWB

Fisher matrices can be used to learn about the sensitivities of an experiment to specified parameters. These parameters could be ‘physical’ (where the parameters link the theory being investigated directly to the signal being observed) or ‘phenomenological’ (where the parameters simply characterize the spectral shape of the signal).

The Fisher matrix can be written such that its elements are the expectation values of the second derivatives of the log-likelihood $l_G = \ln P(D|\theta)$ with respect to some set of parameters θ ,

$$F_{ij} = \left\langle \frac{\partial^2 l_G}{\partial \theta_i \partial \theta_j} \right\rangle.\tag{6.6}$$

The inverse matrix $C = F^{-1}$ is then the covariance matrix of the parameters θ .

In line with extant [GW](#) experiment practices, we consider logarithmically-binned frequency observations: the width of each bin, defined by its (logarithmically) central frequency f_b , is given by δf_b , for a total of N_b bins. $n_b = \delta f_b T_{\text{obs}}$ then gives the number of frequencies accessible in each bin. The data here, D , are the [PSDs](#):

$$D_i \equiv |\tilde{d}(f_i)|^2,\tag{6.7}$$

where $\tilde{d}(f)$ are the (complex) Fourier transforms of the time domain data $d(t)$.

If we assume that $n_b \gg 1$ for all bins, the central limit theorem allows us to move to a

Gaussian approximation for the likelihood [197, 198]:[†]

$$p(\bar{D}_b|\theta) = \prod_{b=1}^{N_b} \left(\frac{n_b}{2\pi S_b^2(\theta)} \right)^{\frac{1}{2}} \exp \left(-\frac{1}{2} \frac{n_b(\bar{D}_b - S_b(\theta))^2}{S_b^2(\theta)} \right), \quad (6.8)$$

where $\bar{D}_b$ is the weighted mean for D in bin b , $S_b(\theta) = (1/n_b \sum 1/S(\theta, f))^{-1}$ is the spectral density of bin b , having averaged the variances $S(\theta, f)$ over the accessible frequencies in the bin.

We can make a less conservative forecast by incorporating the number of independent detector channels of a given experiment, N_{channel} , such that $n'_b = N_{\text{channel}} \times n_b$. As a naive estimate, for this thesis we take $N_{\text{channel}} = 2$ (which is true for LISA [197]). This could be more accurately estimated in future work.

The Gaussian approximation to the log-likelihood, incorporating this factor of 2, can then be written

$$l_G = -\frac{1}{2} \sum_{b=1}^{N_b} \frac{2n_b(\bar{D}_b - S_b)^2}{S_b^2} - \sum_{b=1}^{N_b} \ln S_b + \text{constant}, \quad (6.9)$$

such that the Fisher matrix is

$$F_{ij} = \sum_{b=1}^{N_b} \frac{2n_b}{S_b^2} \frac{\partial S_b}{\partial \theta_i} \frac{\partial S_b}{\partial \theta_j}. \quad (6.10)$$

We can move from S_b to $\Omega(f_b)$ using equation 4.2, to obtain

$$F_{ij} = \sum_{b=1}^{N_b} \frac{1}{\sigma_{b,\Omega}^2} \frac{\partial \Omega_{\text{GW}}}{\partial \theta_i} \frac{\partial \Omega_{\text{GW}}}{\partial \theta_j}, \quad (6.11)$$

where $\sigma_{b,\Omega}^2$ is the variance of Ω_{GW} in bin b ,[‡] given by

$$\frac{1}{\sigma_{b,\Omega}^2} = T_{\text{obs}} \frac{2\delta f_b}{\Omega_{\text{total}}^2}. \quad (6.12)$$

Hence the variance depends not only on the total observation time of the experiment T_{obs} , the width of the frequency bins δf_b , and total number of frequency bins N_b , but also (discussed further in section 6.4.1) on Ω_{total} —we cannot observe the GW signal power in isolation but rather the total spectrum, polluted by the total noise contribution $\Omega_{\text{noise}}(f)$ including foregrounds and any other sources of GW power in addition to instrumental noise (see section 6.3).

[†]Reference [197] begins with the full chi-squared distribution, and goes on to describe the relevant Fisher matrix for this distribution. Per reference [199], the additional term is incorrect, and the full field term evaluation of the Fisher matrix shows that the Gaussian approximation remains correct.

[‡]Note that this gives the variance when only one experiment is involved. Additional experiments which observe in the same frequency bin lead to a reduction in error. See section 6.7.

The signal-to-noise ratio (SNR) ρ is a commonly used figure of merit for the detection of discrete sources, but is less useful when considering the SGWB. The SNR evolves with observing time T_{obs} , since the total signal contribution from the SGWB increases with time. The SNR can be written in terms of either Ω or S :

$$\begin{aligned}\rho^2 &= T_{\text{obs}} \int_{f_{\text{min}}}^{f_{\text{max}}} df \frac{\Omega_{\text{GW}}^2(f)}{\Omega_{\text{noise}}^2(f)} \\ &= T_{\text{obs}} \int_{f_{\text{min}}}^{f_{\text{max}}} df \frac{S_{\text{GW}}^2(f)}{S_{\text{noise}}^{\text{eff}^2}(f)}.\end{aligned}\tag{6.13}$$

But when measuring a power spectrum, the error is determined by the total observed variance, which has contributions from both signal and noise. This can be calculated from equation 6.12, and notably depends on Ω_{total} and not Ω_{noise} . As such, the SNR is not directly used in our results.

6.4.1 The Evolution of Noise with Time

The evolution of noise with time leads us to plot multiple curves describing the instrumental noise contributions of each experiment: the ‘nominal’ noise curve as given by the PSDs shown in Appendix B, and a ‘noise-with-time’ (nwt) curve which demonstrates the effective increase in sensitivity given by running an experiment for its predicted T_{obs} . We define these nwt curves as

$$\Omega_{\text{nwt}}(f_b, T_{\text{obs}}) = \frac{\Omega_{\text{noise}}(f_b)}{\sqrt{T_{\text{obs}} \delta f_b}}; S_{\text{nwt}}(f_b, T_{\text{obs}}) = \frac{S_{\text{noise}}^{\text{eff}}(f_b)}{\sqrt{T_{\text{obs}} \delta f_b}}.\tag{6.14}$$

These nwt sensitivity curves are intended as an alternative to the commonly used power law integrated sensitivity curves (PLISCs) (see [119]), which show a curve above which the SNR meets some threshold value. As discussed above, the SNR is not a proxy for the variance when dealing with the SGWB, and as such PLISCs are not an optimal tool to visualise the Fisher matrix-derived sensitivities of experiments. The nominal and nwt sensitivity curves for each experiment considered are shown in figure 5.1.

The observing frequency range of a given experiment, $\Delta f = f_{\text{max}} - f_{\text{min}}$, depends on the nature of the experiment. However the minimum *possible* frequency also depends on T_{obs} , and is given by $1/T_{\text{obs}}$, regardless of the frequencies eventually reachable by the experiment. Hence an observing time of 1 year gives a minimum possible frequency of $\sim 3 \times 10^{-8}$. For space- and ground-based experiments, this is well below the frequency range being probed. However for pulsar timing arrays (PTAs), which probe frequencies as low as 10^{-9} , this means that their observing frequency range evolves with time. As such, the ‘nominal’ noise curve for PTAs

(International Pulsar Timing Array (IPTA), Square Kilometre Array (SKA)) should not be considered to give an accurate description of the initial performance (in both frequency range, and amplitude). Multiple Ω_{nwt} curves are plotted for SKA in figure 5.1 at different observing durations, to show its evolving frequency range. Figure 6.1 shows Ω_{nominal} and Ω_{nwt} for SKA and LISA for observing periods of 5/10/20 and 0.5/2/4 years, respectively. You can see that the frequency range for SKA evolves with time, but does not for LISA, and that Ω_{nominal} for SKA suggests a much better sensitivity than it will achieve.




Figure 6.1: Ω_{nominal} (dashed lines) and Ω_{nwt} (solid lines) for SKA and LISA. The top axis gives indications of minimum frequencies for a range of T_{obs} . Note that SKA can only start to make measurements after ~ 4 months of observing.

6.5 Power Law Approximations to Cosmological Spectra

When assessing the sensitivities of different experiments to SGWB signals, it is more useful to look at their sensitivities to *spectral shapes* rather than to individual phenomena. This is particularly true when we don't have tight constraints on the physical parameters of any of the phenomena, and thus no specific predicted GW spectrum. It is also convenient to have a generic function describing approximate spectra for phenomena whose spectra can only be simulated.

Moving from these theorized and simulated GW spectra in terms of physical parameters to spectra in terms of spectral parameters is not immediately simple; the shapes of the spectra of possible phenomena are not always an obvious fit to simple curves.

6.5.1 Simple Power Laws

In my first forecasting paper [105], we test the sensitivity of each GW experiment to the parameters for a single power law, defined by its amplitude A and power-law index (or logarithmic slope) m .

$$\Omega = A \left(\frac{f}{f_0} \right)^m. \quad (6.15)$$

A power law can be defined with a specified ‘pivot point’ f_0 , such that its amplitude A at that point is constant but its power m can vary—see figure 6.2. We can rewrite a power law as a straight line by moving to logarithmic parameters:

$$Y = Y_0 + m(X - X_0), \quad (6.16)$$

where $Y = \log_{10}(\Omega)$, $Y_0 = \log_{10}(A)$, $X = \log_{10}(f)$, $X_0 = \log_{10}(f_0)$.

6.5.2 Quadratic Deviations from Power Laws

In my second paper [191], we move to more complicated spectral shapes. We first consider whether a given experiment (or combination) will be able to make basic constraints on curvature present in the spectrum, before evaluating the sensitivity to more accurate spectral shapes.

We can consider a deviation from the simple power law in equations 6.15 and 6.16, by introducing a quadratic term in log-space:

$$Y = a(X - X_0)^2 + Y_0, \quad (6.17)$$

where Y , Y_0 , X , and X_0 are defined as above. Note that this describes a parabola with a directrix parallel to the x -axis—that is, it tends to a straight line with $m = 0$, in the limit that $a \rightarrow 0$. Although we do consider power laws with slopes $m \neq 0$, this simplified model is sufficient for the purpose of testing experimental sensitivities to curvature in an observed spectrum.

6.5.3 Double Broken Power Laws

We then test for sensitivities to a double broken power law (DBPL), a broken power law with breaks at two distinct frequencies (and therefore three power law segments), in contrast to a single broken power law (SBPL) which has only one break (and two power law segments)[§]—see figure 6.2b. Much of the previous work on sensitivities of different GW experiments has characterized spectra with SBPLs. For some phenomena, this is a reasonable fit, but not for all. In [197], a DBPL was used to test the sensitivities of LISA to first order phase transitions (FOPTs). That DBPL was therefore chosen specifically to fit the spectra of FOPTs which might be observable at the LISA frequency band.

A generic DBPL[¶] is a good fit to all of the SGWB phenomena being considered in this thesis. For a fully generic function, we need eight spectral parameters:

$$\Omega(a, b, f_1, f_2, j, k, y, z; f) = A_2 \left(\frac{f}{f_2} \right)^j \left(\frac{1 + \left(\frac{f_2}{f_1} \right)^y}{1 + \left(\frac{f}{f_1} \right)^y} \right)^{\frac{j-b}{y}} \left(\frac{b+k}{b+k-1 + \left(\frac{f}{f_2} \right)^z} \right)^{\frac{b+k}{z}}, \quad (6.18)$$

where f_1 and f_2 define the positions of the two breaks in the function, j and k define the power laws at the left and right of the function, and b the power law between the breaks. y and z define how smooth the transition between the sections of the function are, at the left and right breaks respectively. A_2 is the amplitude at f_2 . This ends also up being the peak if the equality

$$\frac{\left(\frac{f_2}{f_1} \right)^y}{1 + \left(\frac{f_2}{f_1} \right)^y} = \frac{j-1}{j-b} \quad (6.19)$$

is satisfied. The function could be modified to require this, but it loses generality—partly in that the peak centres closer to f_1 for some phenomena, and partly because different sets of parameters fit simulated spectra better with peaks close to, but not at, the breaks in the function. Note that this is not simply a generalization of the [197] DBPL.

[§]See the ‘FOPTs1a’ and ‘FOPTs2a’ lines in figure 4.1 for an example of a DBPL and SBPL, respectively.

[¶]which approximates a SBPL in the limit that $f_1 = f_2$, as used in equation 6.18.




Figure 6.2: Examples of power law functions

(a) A simple power law, with its amplitude specified at the pivot point (star) and varying slopes.

(b) Examples of broken power laws—blue dot-dash curve is a [SBPL](#), and red dash and pink solid curves are [DBPLs](#).

6.6 Power Law Sensitivities

6.6.1 Simple Power Laws

We can rewrite equation 6.11 to obtain the Fisher matrix in Y rather than Ω :

$$F_{ij} = \sum \frac{1}{\sigma_{b,Y}^2} \frac{\partial Y}{\partial \theta_i} \frac{\partial Y}{\partial \theta_j}, \text{ where } \frac{1}{\sigma_{b,Y}^2} = \ln(10)^2 T_{\text{obs}} \frac{2\delta f_b \Omega_{\text{GW}}^2}{\Omega_{\text{total}}^2}. \quad (6.20)$$

For the power law defined in equation 6.16, this equates to

$$F = \begin{bmatrix} \sum_b \frac{1}{\sigma_b^2} (X - X_0)^2 & \sum_b \frac{1}{\sigma_b^2} (X - X_0) \\ \sum_b \frac{1}{\sigma_b^2} (X - X_0) & \sum_b \frac{1}{\sigma_b^2} \end{bmatrix}, \text{ and } C = \begin{bmatrix} \sigma_m^2 & \sigma_{m,Y_0}^2 \\ \sigma_{m,Y_0}^2 & \sigma_{Y_0}^2 \end{bmatrix}. \quad (6.21)$$

The form of this matrix allows for a simple diagonalization of both the Fisher and covariance matrices, de-correlating the parameters, by specifying the pivot frequency:

$$X_0 = \sum_b \frac{1}{\sigma_b^2} X_b / \sum_b \frac{1}{\sigma_b^2}, \quad (6.22)$$

to give

$$F = \begin{bmatrix} \sum_b \frac{1}{\sigma_b^2} X^2 - \frac{\left(\sum_b \frac{1}{\sigma_b^2} X\right)^2}{\sum_b \frac{1}{\sigma_b^2}} & 0 \\ 0 & \sum_b \frac{1}{\sigma_b^2} \end{bmatrix}, \text{ and } C = \begin{bmatrix} \sigma_m^2 & 0 \\ 0 & \sigma_{Y_0}^2 \end{bmatrix}. \quad (6.23)$$

The quantity X_0 is simply the weighted average of frequencies given their variances, and represents the point on the spectrum which can be best constrained.

6.6.2 Deviations from Power Laws

With the same process, we obtain the Fisher matrix in Y for the quadratic log-spectrum defined in equation 6.17:

$$F = \begin{bmatrix} \sum_b \frac{1}{\sigma_b^2} (X - X_0)^4 & \sum_b \frac{1}{\sigma_b^2} (X - X_0)^2 \\ \sum_b \frac{1}{\sigma_b^2} (X - X_0)^2 & \sum_b \frac{1}{\sigma_b^2} \end{bmatrix}, \quad \text{and } C = \begin{bmatrix} \sigma_a^2 & \sigma_{a,Y_0}^2 \\ \sigma_{a,Y_0}^2 & \sigma_{Y_0}^2 \end{bmatrix}. \quad (6.24)$$

The value of X_0 is again free to be specified. The quadratic forms of the off-diagonal elements mean that it is not possible to choose an X_0 which diagonalizes the matrix like above. In fact, the value of X_0 which minimises the covariances σ_{a,Y_0}^2 is the same value which diagonalises the power law Fisher matrix above (given in equation 6.22). This follows the same logic as above, noting that the vertex of the parabola sits at the point on the equivalent straight line spectrum which can be best constrained.

6.6.3 Double Broken Power Laws

For this Fisher matrix, we leave log-variables and return to equation 6.11. Each term of the Fisher matrix can then be evaluated analytically.

All eight spectral parameters are needed for full generalization of the function, but there is some degeneracy which can be exploited to reduce the number of parameters evaluated for sensitivity, and reduce the size of the Fisher matrix required: j and y are reasonably degenerate to one another, as are k and z , and so we fix y and z . This leaves only six parameters to evaluate sensitivities to: $[A_2, f_1, f_2, b, j, k]$. Of these, some can be fixed for a given source (b , and often $[j, k]$) while the remainder are chosen from a specified range.

6.7 Combining Multiple Experiments

A benefit of the Fisher matrix formalism is the simplicity of considering the combined sensitivities of multiple experiments. This can be done simply by summing the Fisher matrices of each experiment:

$$F_{\text{combined}} = F_{\text{exp 1}} + F_{\text{exp 2}} ; \quad C_{\text{combined}} = F_{\text{combined}}^{-1}. \quad (6.25)$$

Note that this explicitly considers independent experiments, which are observing separate realizations of the sky. In reality, each experiment observes the same sky, and there are cross-correlations to consider. However, in line with established cosmic microwave background (CMB) science, this methodology is a very good approximation.

C_{combined} can then be used to find the improvement on σ_Y^2 and σ_m^2 of a combination of experiments, compared to each experiment individually. Note that equation 6.25 requires a consistent fiducial GW spectrum across the experiments being considered.

We can also calculate the new variance of a single frequency bin when a combination of experiments are considered, $\sigma_{b,Y,\text{combined}}^2$:

$$\sigma_{b,Y,\text{combined}}^2 = 1 / [1/\sigma_{b,Y,1}^2 + 1/\sigma_{b,Y,2}^2 + 1/\sigma_{b,Y,3}^2 \dots] . \quad (6.26)$$

The reduction in variance of a single frequency bin of combined experiments compared to a given experiment 1, is then given by

$$R = \frac{\sigma_{b,Y,\text{combined}}^2}{\sigma_{b,Y,1}^2} = 1 + \sigma_{b,Y,1}^2/\sigma_{b,Y,2}^2 + \sigma_{b,Y,1}^2/\sigma_{b,Y,3}^2 \dots . \quad (6.27)$$

6.8 Adding an Observational Prior

With the Fisher matrix formalism, the inclusion of previous observations[‡] as a prior is identical to the inclusion of additional experiments—it can be done just through the addition of an extra Fisher matrix per section 6.7, this time with the inverse of the observed error bars in the places of $1/\sigma_m^2$ and $1/\sigma_Y^2$. In this case, the fiducial model must match the mean observed values for Y and m .

We can instead consider the case where the fiducial GW spectrum under consideration deviates from the prior. By taking the distributions of both the fiducial model and the observed prior to be Gaussian and described by θ_p, σ_p and θ_d, σ_d , respectively, we can find their combined distribution by taking their product. We aim to get an approximately Gaussian combined distribution. Dropping prefactors, the product of the two initial Gaussian distributions can be written as a Gaussian with mean and variance θ_*, σ_*^2 using

$$\frac{(\theta - \theta_p)^2}{\sigma_p^2} + \frac{(\theta - \theta_d)^2}{\sigma_d^2} = \frac{(\theta - \theta_*)^2}{\sigma_*^2} + \chi^2, \quad (6.28)$$

[‡]In this thesis, the NANOGrav15 results.

where

$$\frac{1}{\sigma_*^2} = \frac{\sigma_p^2 + \sigma_d^2}{\sigma_p^2 \sigma_d^2}, \quad \theta_* = \frac{\theta_p \sigma_d^2 + \theta_d \sigma_p^2}{\sigma_p^2 + \sigma_d^2}, \quad \chi^2 = \frac{(\theta_p - \theta_d)^2}{\sigma_p^2 + \sigma_d^2}. \quad (6.29)$$

The product depends only on θ_* , σ_*^2 when $\theta_p \approx \theta_d$. Hence the Fisher formalism can only be used to calculate the constraining power of experiments with an observational prior in this case, i.e., when considering a fiducial model which is close to the prior model. In the forecasting work in Chapter 7, we limit ourselves to cases where $|\theta_p - \theta_d|$ lies within the reported 95% confidence level.

6.9 Summary and Discussion

Spectra for the range of phenomena considered in this thesis can be modelled, with varying levels of approximation, by simple power laws, power laws with an additional quadratic term, and single/double broken power laws.

Over small frequency ranges (e.g. within the frequency range of a single experiment), simple power laws are a good fit to most spectra. When considering spectra across the wider experimental frequency range, broken power laws are a better fit to the spectra of most sources. However, before moving directly to considering these more complicated spectral shapes, it is worth evaluating the sensitivity to a proxy for general curvature in the spectrum, a .

On top of these cosmological spectra, we model the astrophysical foreground from binary populations with a new generalized function (equation 6.4). Improvements to the accuracy of this function could be made, to move from the sharp drop off at frequencies higher than the peak, to incorporate the GWs emitted during the ringdown phase. For our purposes, this lies below the sensitivities of the experiments considered, and so would not affect our results.

Noise models for each experiment for which forecasting is done can be found in Appendix B.

Chapter 7

Forecasting Direct Sensitivities to the SGWB

This chapter presents the results from my first [SGWB](#) forecasting paper [\[105\]](#), applying the Fisher matrix formalism described in [Chapter 6](#) to evaluate experimental sensitivities to simple power law models of the [SGWB](#). In particular, it establishes the sensitivities of current and planned [GW](#) experiments to the amplitudes and slopes of different fiducial spectra, identifying pivot frequencies and highlighting how experimental sensitivities evolve over time. There is a focus on the improvements which come from combinations of experiments, compared to their individual constraining powers.

The analysis also explores how incorporating the [NANOGrav](#) 15-year results can enhance experimental constraining power. This work sets the stage for the analysis of sensitivities to more complex spectral models which will be covered in my upcoming paper, [\[191\]](#), preliminary results for which are given at the end of this chapter.

7.1 Introduction (Single Power Laws)

This work uses diagonalized Fisher matrices to evaluate the sensitivities of different gravitational wave ([GW](#)) experiments to the spectral parameters of a range of cosmological phenomena, modelled as single power laws in the presence of realistic astrophysical foregrounds.

Specifically, we consider the improvements to sensitivities that using multiple experiments in conjunction can provide, as well as including the NANOGrav15 results as a prior.

7.2 Results (Single Power Laws)

We consider first how well an individual experiment can constrain a single power law—what are its sensitivities to the slope and amplitude? Running a fiducial spectrum through the Fisher matrix gives us the minimum error on the (log) amplitude Y_0 and slope m (σ_{Y_0} and σ_m , respectively) for a given experiment. Examples are shown in figure 7.1, where σ_{Y_0} is shown as a shaded band above and below the fiducial spectrum, and σ_m is shown as the shaded region between the upper and lower limits on slope, added to the limits on the amplitude. Both errors are centred on the pivot point, $f_0 = 10^{X_0}$, plotted as a star. This new method, displaying the decorrelated errors with the combined shaded region, does not fully capture the combination of errors but is a reasonable approximation.

We use 20 logarithmically-spaced bins per experiment—changing the number of bins gives negligible impact on the results, provided that there are at least three bins. There is likewise negligible impact if linearly-separated bins are used instead. Where bins are shared between multiple experiments (where we are evaluating bin-by-bin errors for combinations of experiments, for example) the total number of bins is chosen to approximately match the density of bins across frequency space in the individual analyses.

The fiducial models used in our results are chosen primarily to illustrate a generalization of our findings, but are all taken from plausible physical spectra which have been put forward as stochastic gravitational wave background (SGWB) contributions from a range of phenomena, and can all be found in Appendix C.

7.2.1 Individual Experiments

Where the constraints on a spectrum are good, i.e., low σ_{Y_0} and σ_m , the shaded regions present as narrow bands along the fiducial line. Conversely, poor constraints generate wide ‘bowtie’ shaped errors. Sensitivities of a selection of experiments to an example cosmic string spectrum (strings1d, [132]) are shown in figure 7.1.

The improving sensitivities of the three ground-based experiments (LVK-O4, LVK-O5, CE) shown in figures 7.1a–7.1c can be seen both in the move of the shaded sensitivity regions to lower amplitudes, as well as in the shrinking of the ‘bowties’ which represent the errors on the slope and amplitude—the ‘bowtie’ collapses to a narrow band in cases where the fiducial spectrum can be very well constrained. Figure 7.1d then shows the sensitivity of LISA to the same fiducial spectrum at lower frequencies, as a comparison.




Figure 7.1: The sensitivities of [LVK-O4](#), [LVK-O5](#), [CE](#) and [LISA](#) to (strings1d). Shaded sensitivity regions match those in figure 5.1, and the grey dash-dot line is likewise the total instantaneous foreground.


- (a) LVK-O4: $f_0 = 40.8$ Hz, $\Omega_0 = 1.08 \times 10^{-10} \pm 1.17 \times 10^{-11}$, $m = -0.07 \pm 0.80$.
- (b) LVK-O5: $f_0 = 39.7$ Hz, $\Omega_0 = 1.09 \times 10^{-10} \pm 1.00 \times 10^{-11}$, $m = -0.07 \pm 0.23$.
- (c) CE: $f_0 = 12.2$ Hz, $\Omega_0 = 1.18 \times 10^{-10} \pm 1.32 \times 10^{-13}$, $m = -0.07 \pm 0.004$.
- (d) LISA: $f_0 = 0.004$ Hz, $\Omega_0 = 2.06 \times 10^{-10} \pm 2.57 \times 10^{-13}$, $m = -0.07 \pm 0.002$.

As discussed in section 6.4.1, a fiducial spectrum which crosses above the Ω_{nwt} line should be observable. But it is not true that the entire observable frequency range of an experiment can be well-constrained, as can be seen in the ‘well-constrained’ frequency ranges in figure 7.2—it can be seen that the width of the ‘well-constrained’ frequency range is all-important in determining whether the *shape* of the spectrum can be confidently assessed.


Recall from section 6.6 that the pivot point gives the frequency where the amplitude of the spectrum is best constrained. An interesting question is what the contribution of each frequency bin is to the amplitude of the fiducial spectrum— $\sigma_{Y_0,b}$ depends on the particular frequency bin f_b , and varies widely across the frequency ranges of each experiment. It is

therefore useful to probe the actual frequency range over which the experiment can confidently assess the amplitude of a given frequency bin, since the Fisher matrix formalism assumes the fiducial model across the entire prescribed frequency range. This is particularly important in assessing the shape of an observed spectrum, as a region of well-constrained Y_0 also gives the frequency range where the shape can be confidently determined. These ranges, specified as frequency bins where $\sigma_{Y_0,b}/Y_b \leq 1\%$, are shown in red.

Figure 7.2 shows that there are some cases where ‘bowtie’ results can be misleading with regards to the errors on the spectral shape: in cases where the fiducial model dips below the sensitivity of an experiment, they can over- (figure 7.2a) or under-estimate (figure 7.2b) the 1σ regions of uncertainty. Figure 7.2a is a useful demonstration that the ‘bowtie’ within the sensitivity region of an experiment is a good match to the shape of the individual error bars, but begins to fail outside of the sensitivity region. In this case the error bars show that there is a region of good constraint (shown in red) and a region where we can only establish an upper bound, while the ‘bowtie’ remains misleadingly narrow across the whole frequency space. Where the spectrum lies entirely below the ultimate sensitivity of an experiment, an upper bound rather than observation follows: the shape of the experiment’s sensitivity curve itself rules out not only a range of amplitudes, but also of slopes. In these situations, the ‘bowties’ which are generated are often not the most useful result. Figure 7.2b demonstrates that the error bars $\sigma_{Y_0,b}$ sketch out a more useful area than σ_{Y_0} and σ_m when evaluating the sensitivity of DECIGO to a single field slow roll (SFSR) model of inflation.



(a)



(b)

Figure 7.2: A visualization of the results given by σ_{Y_0} and σ_m (left plots) compared to $\sigma_{Y_{0,b}}$ (right). Where the fiducial model is plotted in red, $\sigma_{Y_{0,b}}/Y_b \leq 1\%$, else it is plotted in black.
(a) ET (strings1d): $f_0 = 35.5$ Hz, $\Omega_0 = 1.73 \times 10^{-12} \pm 1.40 \times 10^{-13}$, $m = -0.09 \pm 0.168$. Range constrained to 1%: 19.3–55.1 Hz.
(b) DECIGO (inflation1): $f_0 = 0.052$ Hz, $\Omega_0 = 1.91 \times 10^{-16} \pm 2.06 \times 10^{-15}$, $m = 0 \pm 7.16$. Range constrained to 1%: N/A.

7.2.2 Multiple Experiments

We next evaluate the improvements to the constraints on slope and amplitude by using multiple experiments in conjunction, and also the regime over which we can be confident about the shape of the spectrum being observed. We divide combinations of experiments into two categories: where the experiments overlap in frequency space, and where they are distinct—i.e., they lie on either side of a gap in observing frequency.

Overlapping Experiments

Within broad bands at low (pulsar timing) and high (ground-based) frequencies, upcoming experiments overlap in frequency space with the current experiments, to provide improved sensitivities. This can be seen in figure 5.1, where the sensitivities progress to lower amplitudes moving from [IPTA](#) $\rightarrow$ [SKA](#), and from [LVK-O3](#) $\rightarrow$ [LVK-O5](#) $\rightarrow$ [CE](#) and [ET](#).

A useful test is then whether the new experiment renders its predecessor irrelevant with regards to the [SGWB](#), or whether there is a noticeable improvement in constraining power when combining both experiments. Figure 7.3 considers the combined sensitivity of [IPTA](#)+[SKA](#) to a fiducial model representing primordial black holes (PBHs) (PBHs2, [200]) in two cases: where SKA has been observing for 5 years, and for the full 20-year observing run. When the two experiments' individual sensitivities are at similar levels, the inclusion of [IPTA](#) reduces the [SKA](#) errors σ_{Y_0} and σ_m by factors of 1.4 and 2.4, respectively. But by the time the sensitivity of [SKA](#) reaches even one decade lower in Ω_{GW} , the improvements added by [IPTA](#) are negligible.

Similarly, and now considering a fiducial model of cosmic strings with tension of $G\mu \approx 10^{-11}$ (strings1b, [132]), the less sensitive LVK-O4 can only offer [ET](#) negligible reductions in σ_{Y_0} and σ_m . Whereas [CE](#) with its similar sensitivity to [ET](#) confers reductions of 1.5 and 2.1, for cosmic strings with a tension of $G\mu \approx 10^{-15}$ (strings1d, [132]).

Per the Fisher matrix formalism, two mathematically identical experiments run in conjunction would give an improvement of $\sqrt{2} \approx 1.4$ over each experiment individually.* [CE](#) and [ET](#) are not identical, and in fact give a reduction in σ_{Y_0} of *at least* $\sqrt{2}$ in combination with each other, but give even more improved results for σ_m . This is largely as a result of the widening of the most sensitive frequency range.

These results are generalizable: a new experiment with better sensitivity does not gain significant improvements to its constraining power with the retention of its predecessor (although it should be noted that there are other science cases for continuing observations with older experiments, including their use in locating (non-[SGWB](#)) sources on the sky, given sufficiently high signal-to-noise ratio ([SNR](#)), and the mitigation of systematic errors). But combinations of experiments with similar sensitivities can be very useful in reducing the errors. Future experiments with similar sensitivities and frequency ranges are planned to run in conjunction: [CE](#)+[ET](#) (which are specifically intended to run as a combined experiment, like [LVK](#), to obtain better spatial resolution on individual sources), and [BBO](#)+[DECIGO](#) (which are not planned as a combined experiment, but do have significant frequency overlap).

*With the caveat that this relates to fully independent experiments looking at independent realizations of the sky—refer back to discussion in section 6.7.




Figure 7.3: Improvement in constraining power that IPTA confers to SKA where SKA has been observing for 5 years (top) and 20 years (bottom) to (PBHs2).

(a) (IPTA, SKA-5yrs, combined): $f_0 = (5.6, 14, 8.1) \times 10^{-9}$ Hz, $\Omega_0 = 2.9 \times 10^{-10} \pm (11.0, 13.4, 8.5) \times 10^{-11}$, $m = -0.01 \pm (0.80, 1.14, 0.47)$. $f_0 = 7.9 \times 10^{-9}$ Hz, $\Omega_0 = 2.9 \times 10^{-10} \pm (13.2, 13.3, 9.4) \times 10^{-11}$, $m = -0.01 \pm (1.09, 1.14, 0.48)$.

(b) (IPTA, SKA-20yrs, combined): $f_0 = (5.6, 10, 8.8) \times 10^{-9}$ Hz, $\Omega_0 = 2.9 \times 10^{-10} \pm (11.0, 5.8, 5.2) \times 10^{-11}$, $m = -0.01 \pm (0.80, 0.29, 0.26)$.

Distinct Experiments

With this confirmation that overlapping experiments confer additional sensitivity only when one is not much more sensitive than the other, we now turn to combinations of experiments which do not overlap.

Experiments covering a wide range of frequency space will give information on the overall shape of the SGWB spectrum, even within the frequency range of a single experiment. We look to see what the improvements will be when experiments at different frequencies come online and are evaluated in combination. In particular, we consider groups of experiments which straddle

gaps in observing frequency, to establish whether the improved sensitivities of experiments in combination can alleviate issues around unobserved frequencies.

The results are highly dependent on the fiducial model—a [SGWB](#) spectrum which is already very well constrained by a single experiment, but lies beneath the sensitivity region of another, is no better constrained by the combination of experiments. The cases where evaluating experiments in combination works well are those where the fiducial model is reasonably constrained by each experiment independently.

The example shown in figure [7.4](#) compares the individual and combined sensitivities of [LISA](#), [CE](#) and [ET](#) to a fiducial model representing cosmic strings with tension of $G\mu \approx 10^{-15}$ (strings1d, [\[132\]](#)). Each experiment individually constrains Ω by a similar amount, $\sigma_\Omega \approx 1.3 \times 10^{-13}$, but [LISA](#) constrains m considerably better than the ground-based experiments, with $\sigma_{m \text{ LISA}} = 0.08$ compared to $\sigma_{m \text{ CE, ET}} \approx 0.2$. These constraints reduce to $\sigma_\Omega \approx 8.4 \times 10^{-14}$ and $\sigma_m = 0.007$, improvements of 1.5 and 11, respectively.

In general, very good constraints can be conferred from one end of the spectrum to the other, but this relies on having confidence in the fiducial model as a single power law. In cases where we are confident in this shape of power law across a wide frequency range (e.g., some models of inflation and cosmic strings, or the $f^{2/3}$ sections of astrophysical foregrounds), this is reasonable. However, while experiments can work together across a frequency gap to constrain a fiducial model, they cannot generate any errors bars in the ‘unseen’ frequency bands between experiments (see the right-hand panels in figure [7.4](#)).




Figure 7.4: The sensitivities of **ET**, **CE** and **LISA** to (strings1d), shown both as experiments operating individually and as a group.

(a) (LISA,ET,CE): $f_0 = (0.005, 12.1, 35.5)$ Hz, $\Omega_0 = (3.9, 1.9, 1.7) \times 10^{-12} \pm (1.2, 1.3, 1.4) \times 10^{-13}$, $m = -0.09 \pm (0.08, 0.22, 0.17)$. Range constrained to 1%: (0.001–0.010, 8.13–21.4, 19.3–55.1) Hz.

(b) LISA+CE+ET combined: $f_0 = 0.04$ Hz, $\Omega_0 = 3.2 \times 10^{-12} \pm 8.4 \times 10^{-14}$, $m = -0.09 \pm 0.007$. Range constrained to 1%: 0.001–0.1, 8.13–55.1 Hz.

7.2.3 NANOGrav15 as a Prior

The addition of the NANOGrav15 results as a prior behaves identically to the addition of a distinct experiment, although any variation in the fiducial model compared to the posteriors from NANOGrav15 must be dealt with as laid out in section 6.8. And as justified in that section, we only consider fiducial models which lie within a reasonable distance from the posteriors. In the case of a single power law, this is particularly limiting, although this requirement is relaxed when more complicated spectral shapes are introduced (and will be considered in my upcoming paper [191]).

The NANOGrav15 results [49, 50] are given both in joint $A_{\text{CP}}\text{--}\gamma_{\text{CP}}$ posteriors, and a single A_{CP} posterior for fixed $\gamma_{\text{CP}} = 13/3$ (to match the expected supermassive black hole binary (SMBHB) spectral shape). At the reference frequency $f_{\text{ref}} = 1 \text{ yr}^{-1}$ the posteriors are independent, and can be converted to decorrelated errors on Y_0 and m . The two sets of posteriors (with and without fixed γ_{CP}) converted into our Y_0 and m parameters with standard linear error propagation are plotted in figure 7.5a. The hatched region in figure 7.5b then shows an extrapolation of the posteriors to higher frequencies, assuming an unrealistic spectrum characterized by a single power law at all frequencies—note that the blue shaded region here crosses into the LVK-O3 upper limit, demonstrating that this is an unrealistic extrapolation.




Figure 7.5: The NANOGrav15 results [49] plotted over GW experiment sensitivity regions, with the instantaneous foreground, for comparison. The blue shaded region shows the joint $A_{\text{CP}}\text{--}\gamma_{\text{CP}}$ posterior, and the purple region the A_{CP} posterior for fixed $\gamma_{\text{CP}} = 13/3$, both to the 95% confidence level. The diagonally hatched grey region shows the upper limit from LVK-O3.

Figure 7.6a shows the addition of the NANOGrav15 results as a prior for LVK-O4, for a fiducial model which matches the mean NANOGrav15 posteriors exactly. This represents a scenario where the SGWB spectrum presents as a single power law from pulsar timing array (PTA) to ground-based frequencies, and where it is also detected by the current operational phase of LVK. In this case, the inclusion of NANOGrav15 as a prior does not confer any improvements to the constraints that LVK-O4 would set, and indeed LVK-O4 would improve upon the NANOGrav15 constraints themselves.

A different case, if LVK-O4 were to make observations of a SGWB spectrum which deviated from the NANOGrav15 posteriors, is shown in figure 7.6b. In this scenario, per the discussion in section 6.8, the combination of the two initial power laws (the NANOGrav15 prior, and that observed by LVK-O4) can be thought of as constraints on a third power law spectrum, consistent with both sets of observations. Again in this case, the inclusion of NANOGrav15 as




Figure 7.6: The sensitivity of LVK-O4, with and without NANOGrav15 as a prior.
(a) LVK-O4(+NANOGrav15): $f_0 = 40.4$ Hz, $\Omega_0 = 6.6 \times 10^{-9} \pm 3.12$ (3.12×10^{-11}), $m = -0.1 \pm 0.01$ (0.01).
(b) LVK-O4(+NANOGrav15): $f_0 = 40.8$ Hz, $\Omega_0 = 7.3$ ($7.1 \times 10^{-9} \pm 3.12$ (3.11×10^{-11}), $m = -0.07(-0.07) \pm 0.12$ (0.11).

a prior only brings negligible improvements to the constraints.

7.3 Summary and Discussion (Single Power Laws)

Following on from the first detections of the SGWB at PTA frequencies, and in advance of any detections at other frequencies, we remain reliant on upper limits and forecast sensitivities in our understanding of the wider SGWB spectrum.

Previous works have primarily focused on the sensitivities of given experiments to one, or a range of, cosmological SGWB spectra (see, e.g., [197, 201, 202, 203, 204, 205, 206, 207, 208, 209]). A handful of works have considered the improvement to LISA’s sensitivity when a second experiment is combined [210, 211, 212], and reference [213] discusses the benefit of multi-frequency detections to an example PBH spectrum. Whereas in this paper we generalize to evaluating the respective sensitivities of current and upcoming GW experiments to SGWB spectra which map to generic single power laws, looking specifically at the improvement to sensitivity which is conferred by utilising multiple experiments in conjunction with each other in the presence of realistic foreground contamination.

Our results are given in terms of the constraints that the experimental setup confers on the amplitude and slope of the fiducial model, and the subsequent improvements on these when experiments are evaluated in combination with each other. We plot these as ‘bowtie’ shapes, which are a reasonable approximation to the true combination of amplitude and slope errors.

These bowtie shapes are a useful visualization of the results, for the region where they lie above our Ω_{nwt} curves. As discussed in section 7.2.1, they can over- and underestimate the error bars below the curves.

Our generalizable findings, discussed above in section 7.2.2, reveal that a more sensitive replacement experiment does not gain noticeable improvement in its constraints on the slope or amplitude from the retention of its predecessor, while the inclusion of another experiment at similar sensitivity levels often shows greater than the simplest $\sqrt{2}$ improvement.

We also show that the combination of experiments across a gap in observing frequency can confer good improvements to their individual constraints, under the assumption of a continuous fiducial model across the frequencies in question. We demonstrate the inclusion of the most recent [NANOGrav](#) results as a prior for fiducial models, and show that, with their present accuracy, they do not bring a noticeable improvement to signals detectable by [LVK-O4](#).

Preliminary Results for More Complex Spectra

The remainder of this chapter presents preliminary results from my upcoming [SGWB](#) forecasting paper [\[191\]](#). As with my first paper [\[105\]](#), this forecasts the sensitivities of experiments to [SGWB](#) spectra using the Fisher matrix formalism described in Chapter [6](#). This work moves from simple power laws to more complicated spectral shapes.

7.4 Introduction (More Complex Spectra)

We move now to considering more complex stochastic gravitational wave background ([SGWB](#)) spectral shapes, including those with quadratic deviations from linear power laws, and single/double broken power laws (which provide a representative fit for a wider range of predicted [SGWB](#) spectra). By extending the range of predicted [SGWB](#) spectra which can act as fiducial models, we will ultimately be able to use these techniques to construct a timeline for when it might be possible to make observations of different contributions of the [SGWB](#).

7.5 Results (More Complex Spectra)

We begin by considering the ability of each experiment to evaluate curvature (or lack thereof) in a fiducial spectrum, by considering sensitivity to a quadratic deviation from a linear power spectrum. This allows us to determine whether it is subsequently worth evaluating sensitivities to the more complex parameters which define a broken power law. We look at sensitivity to the curvature parameter a from equation [6.17](#), for fiducial spectra with both $a = 0$ (i.e. no curvature) and $a < 0$ (negative curvature, mimicking the broken power law shapes we will later probe).

As discussed in section [6.6.2](#), we cannot diagonalize the Fisher matrix for this spectral shape, but we can minimize the off-diagonal elements by fixing X_0 to the appropriate value. This does

not fully decorrelate the errors on a and Y_0 (σ_a and σ_{Y_0}), but minimizes their covariances.

Evaluating σ_a can inform us both about the capacity of an experiment to confirm a flat power law spectrum, as well as sensitivity to the characteristics of a pseudo-broken power law spectrum.




Figure 7.7: The sensitivities of **ET** to (strings1d): $f_0 = 37.0$ Hz, $\Omega_0 = 9.24 \times 10^{-13} \pm 7.41 \times 10^{-14}$, $a = 0.0 \pm 1.05$. Shaded sensitivity regions match those in figure 5.1, and the grey dash-dot line is likewise the total instantaneous foreground.

(a) Equivalent to ‘bowtie’ shape, showing the possible spectral shapes, assuming the shape described in equation 6.17.

(b) Errors bars on Y_0 .

7.5.1 Example Results

By analogy with the ‘bowtie’ and error bar visualizations introduced in section 7.2, equivalent visualizations are shown in figure 7.7. With an example flat fiducial spectrum ($a = 0$), figure 7.7a shows the range of curvature in the spectrum which remains possible. Figure 7.7b then shows the errors on the amplitude, $\sigma_{Y_0,b}$.

7.5.2 Confirming Flat Spectra

Figure 7.8 shows the individual and combined sensitivities of **CE** and **ET** to the same flat cosmic string spectrum as probed above (strings1d). Paying particular attention to σ_a , **CE** (despite its significantly inferior individual constraining power on a) confers a valuable improvement on the ability to confirm the flat fiducial spectrum.




Figure 7.8: The sensitivities of **CE** and **ET** to (strings1d). Shaded sensitivity regions match those in figure 5.1, and the grey dash-dot line is likewise the total instantaneous foreground.

(a) CE: $f_0 = 12.3$ Hz, $\Omega_0 = 9.24 \times 10^{-13} \pm 6.72 \times 10^{-14}$, $a = 0.0 \pm 2.10$.

ET: $f_0 = 37.0$ Hz, $\Omega_0 = 9.24 \times 10^{-13} \pm 7.41 \times 10^{-14}$, $a = 0.0 \pm 1.05$.

(b) Combined: $f_0 = 20.5$ Hz, $\Omega_0 = 9.24 \times 10^{-13} \pm 1.28 \times 10^{-14}$, $a = 0.0 \pm 0.46$.

7.5.3 Constraining Curved Spectra




Figure 7.9: The sensitivities of **IPTA** and **SKA** to (PBHs1). Shaded sensitivity regions match those in figure 5.1, and the grey dash-dot line is likewise the total instantaneous foreground.

(a) IPTA: $f_0 = 6.12 \times 10^{-9}$ Hz, $\Omega_0 = 2.38 \times 10^{-19} \pm 4.88 \times 10^{-10}$, $a = -5.0 \pm 3.17$.

SKA: $f_0 = 1.09 \times 10^{-8}$ Hz, $\Omega_0 = 1.36 \times 10^{-9} \pm 1.38 \times 10^{-10}$, $a = -5.0 \pm 0.70$.

(b) Combined: $f_0 = 9.58 \times 10^{-9}$ Hz, $\Omega_0 = 1.20 \times 10^{-9} \pm 2.63 \times 10^{-10}$, $a = -5.0 \pm 0.70$.

Figure 7.9 shows the individual and combined sensitivities of **IPTA** and **SKA** to a curved fiducial spectrum associated with Solar-mass primordial black holes (PBHs1). This is an example of a fiducial spectrum which lies well above the sensitivities of both experiments, but **SKA** with its wider frequency range and improved sensitivity can place much better constraints on the tails

of the fiducial spectrum. This suggests that, if the fiducial spectrum were instead modelled with a broken power law, [SKA](#) would likely be able to make good constraints on the power law tails—these are characteristic of specific cosmological phenomena, and so [SKA](#) could therefore narrow down potential sources.

7.6 Summary (More Complex Spectra)

Work is ongoing on evaluating the sensitivities of different experiments to this model of fiducial spectrum, before single broken power law ([SBPL](#)) and double broken power law ([DBPL](#)) spectra are considered.

Part III

Forecasting SO Sensitivities to the CMB and SGWB

Chapter 8

Large Aperture Telescope Scan Strategy

This chapter details a campaign of work on the [LAT](#) scan strategy, to which I made significant contributions as part of the calibration, systematics and sensitivity ([CSS](#)) working group in [SO](#), in collaboration with Sara Simon, Reijo Keskitalo, Zack Atkins, Adriaan Duivenvoorden, Frank Qu, and Zack Li. Significant portions of the text are based on reference [\[214\]](#), a paper which is in preparation for internal [SO](#) review. I made substantial contributions to the campaign and the paper, in both analysis and writing. My specific contributions to the analysis came in both the evaluations (and running) of the initial round of simulations, as well as the optimization of the three elevation scan.

8.1 Introduction

The science forecasting done for the Simons Observatory ([SO](#)) white paper [\[79\]](#) was carried out for a simple baseline Large Aperture Telescope ([LAT](#)) scanning strategy: a classical strategy with four sky patches, similar in style to previous ground-based astronomy experiments [\[90\]](#). This strategy does not take advantage of the [LAT](#)'s high scanning rate and capacity for considerable sky coverage each night, and would not allow for Galactic or transient science—science goals which, in part, drove the need to reevaluate the scanning strategy.

The brief for this campaign was to rank previously explored scan types ('Classical' and 'tiled') alongside new scan types designed for high cadence observations and wide, even coverage, in order to select a nominal scan strategy for the [LAT](#), against a series of figures of merit ([FoMs](#)).

The approach for this work began with the collection of technical constraints from the

technical working groups (TWGs). This informed the scan parameters chosen, and helped to gauge the risks of new scan types. Each scan strategy was then simulated in the time domain, for evaluation against a series of FoMs. The simulations were also passed to analysis working groups (AWGs) for assessment of their scientific performances against different science goals. A plan for testing promising scan strategies on sky during commissioning and initial science observations (ISO) was chosen after weighing the technical risks against expected scientific performance. This plan will select the baseline scan strategy for the next few years of observations.

To best provide data for SO’s various science goals (see section 3.3) it is essential that the sky is hit consistently by the scans—uneven hit rates will impact on the ability to analyze information at various angular scales. This issue is compounded if uniform tiling schemes are used, since they introduce sharp coverage drop offs around the boundaries between tiles. These edge effects impact uniformity and can bias power spectrum estimations, particularly at low multipoles. The tiling strategies which we explore have slightly modulated tile placement to reduce this effect.

Another important feature of the scan strategy is its cross-linking—from how many directions a given pixel is hit. Observations from multiple directions can create a strong null test. They are also beneficial to the polarization data. Each scan will hit a pixel in at least two directions, as the telescope rises and sets, but some scan types incorporate additional directions.

Transient science is dependent on the cadence of observations for each pixel. It is therefore necessary to optimize the overall cadence, as well as aiming for a large observed sky area with even coverage.

8.2 Proposed Strategies

The CSS working group in SO considered novel, previously unexplored scan strategies designed for high cadence observations and wide, even coverage. To compare overall performance, we compare several flavours of these new scans to previously explored strategies.

A previous study [90] considered two candidates: a ‘Classical’ scan approach which targets four key fields and observes them as they rise and set; and a ‘tiled’ approach which tiles the sky into $10^\circ \times 20^\circ$ (RA×Dec) rectangular tiles and optimizes observations to make the coverage as uniform as possible. Although the Classical scan is the most similar to the ACTPol scan strategy, which is the most similar cosmic microwave background (CMB) telescope to the LAT, it has a limited sky area and avoids the galaxy, ruling out Galactic science. The tiled approach improves the coverage (although again avoiding the galaxy), but still lacks uniformity and a

high cadence of observations.

We consider three flavours of scan strategies, beginning with the previously explored classical strategy above, improved tiling strategies (which target more even coverage, including Galactic observations) and moving on to new strategies which provide high cadence observations with wide sky coverage.

The [LAT](#) is involved with science affecting small angular scales (including looking for the primordial gravitational wave background ([GWB](#))). To collect data across the sky, the [LAT](#) will scan the sky for roughly ten months of each year,* at a specified (but not necessarily fixed) elevation, scanning horizontally with a specified azimuthal throw. There will be an avoidance range for the Sun, Moon, and brightest planets. An optimal scan strategy is essential in gathering as much data as possible.

Common properties of the strategies considered include:

- Observing elevations between 40° and 60° ;
- These elevations enable sky coverage between declinations ([Decs](#)) of -57° and 21° , equating to $\approx 60\%$ sky fraction ($\approx 40\%$ after Galactic masking), except for the Classical scan which has lower sky coverage—see the footprints of each scan in figures [8.1–8.7](#);
- Azimuthal throws of 110° — $\text{az}=[35^\circ, 145^\circ]$ and $\text{az}=[215^\circ, 325^\circ]$ for rising and setting scans, respectively, for all high cadence strategies;
- Uniform observing priority across the field;
- A nominal scan rate of 1 deg/s , for all strategies with constant speed;
- Maximum acceleration of 3 deg/s^2 ;
- A Sun avoidance radius of 45° ;
- No Moon or planetary avoidance (but all samples within 5° of the Moon, and 0.5° of Mars/Jupiter/Saturn are flagged and cut);
- And an observing period of one year, with the break in observing between 15/01 and 15/03.

Footprints of each of the scan strategies considered can be seen in the hit maps shown in figures [8.1–8.7](#), maps which show the depth of observations at each pixel.

*January and February of each year may not be used for scanning, as they typically have poor weather conditions. These two months may instead be used for maintenance.

Type 1 Scans

These are scans with constant elevation and constant azimuthal speed.[†] These consist of Classical and tiling scans, which covers all strategies previously used by other CMB telescopes, and ‘high cadence’ scans, which are novel.

The Classical scan considers a single field, subdivided into four sky patches. The boundaries of these patches are rigid. The patch to be observed is chosen first by which are ‘available’ (i.e. do not contain an area of Sun avoidance), and then by which patch has been least observed.

Tiling scans split the sky into a specified number of tiles, scanning one tile at a time, with the same method as for the Classical scan. For preliminary scans, we used two sizes of tile, seeking even coverage. With the ‘Large Tile’ and ‘Small Tile’ scans, the borders move slightly with every successful tile scan, with the aim of minimizing edge effects.

High cadence scans observe as much of the sky as possible each day/night. This prioritizes the cadence per pixel. For Type 1, this involves a constant scan rate (either at one specified elevation, or three). The fixed elevation leads to excess observation time at the extremes of the observed sky, due to the Earth’s rotation. As the sky moves westward, the telescope observes the edges (in Dec) more frequently than the centre, because the central regions move more quickly across the field of view. This uneven sampling leads to variance in map depth at the edges, which are likely to be excluded from analysis—in line with Atacama Cosmology Telescope (ACT) and others [215], SO expects to drop around 15% of the maps, along the edges [79].

1.0 Classical

This strategy divides the sky into four large fields, two north of site latitude, two south, across longitudes of either 0° or 180°. This is the only strategy which does not produce continuous coverage of the sky between the limiting declinations.

This strategy is most similar to AdvACT, but its limited sky coverage and avoidance of the Galaxy impact the science cases.

[†]Except during turnarounds, which are typically cut during analysis.




Figure 8.1: The hit map for the Type 1.0 (Classical) scan, with a Mollweide projection in Equatorial co-ordinates (as are all subsequent maps, unless stated otherwise).

1.1 Tiled

1.1.0 Large Tiles

This strategy divides the observable sky into $20^\circ \times 40^\circ$ tiles (RA $\times$ Dec), with modest oscillation of tile coordinates to reduce non-uniformity issues at the boundaries between tiles.

The scheduler chooses which tiles to observe by prioritizing uniform coverage. Tiled scans require observations at a number of elevations and azimuths throughout the day, which in practice could lead to additional observation efficiency losses not considered in this study, due to the need to re-bias the detectors after an elevation change.




Figure 8.2: The hit map for the Type 1.1.0 (large tiles) scan.

1.1.1 Small Tiles

This strategy is identical to 1.1.0 except with $10^\circ \times 20^\circ$ tiles.




Figure 8.3: The hit map for the Type 1.1.1 (small tiles) scan.

1.2 High Cadence

1.2.0 High Cadence, One Elevation

The single-elevation high cadence scans observe at an elevation of 40° to maximize the available sky area, and use the maximum azimuth throw. Two azimuth ranges are used: $az=[35^\circ, 145^\circ]$ and $az=[215^\circ, 325^\circ]$. This enables rising and setting observations, and observing at one azimuth when the other may be unavailable due to Sun avoidance constraints. The high cadence style scans observe almost the full available sky daily, excluding the area of Sun avoidance.




Figure 8.4: The hit map for the Type 1.2.1 (high cadence, 1 el.) scan.

1.2.1 High Cadence, Three Elevations

The three-elevation high cadence scan follows the same design as the one-elevation high cadence scan but observes at three set elevations: 40° , 50° , and 60° . This enables improved cross-linking at the cost of some uniformity and observation cadence.

Initially this strategy had observing times distributed evenly between each elevation, but I later optimized this for improved performance (see section [8.4.3](#)).




Figure 8.5: The hit map for the initial (non-optimized) Type 1.2.1 (high cadence, 3 el.) scan.

Type 2 and 3 Scans

Type 2 and Type 3 scans are also high cadence, but with varying scan rates—faster scan rates at the edges, and slower across the centre of the observed sky, to reduce the issue of Type 1 high cadence scans spending excess time observing the edges. Type 3 also introduces a sinusoidal ‘oscillating elevation’ for more uniform cross-linking.

2 High Cadence, Varying Speed

The Type 2 scan is similar to the Type 1.2.0 Single-Elevation High Cadence scans. It is at an observing elevation of 40° , but the azimuthal scan rate changes with azimuth to make the coverage more uniform near the map edges. The azimuthal scan speed on the telescope mount is a function of the azimuth θ_{az} , and varies between 0.55 and 0.96 deg/s. The speed function is given by $0.55/\sin(\theta_{az})$. The azimuthal acceleration during the main scan (i.e. not turnaround acceleration) does not exceed 2.24 deg/s^2 .




Figure 8.6: The hit map for the Type 2 (high cadence, varying speed) scan.

3 High Cadence, Varying Elevation

As Type 2 above, but with the addition of $\pm 1^\circ$ sinusoidal oscillations in elevation to give uniform cross-linking across the entire map.

In general, the [LAT](#) changes elevation by rotating the elevation structure (see figure 3.1a), with the corotator rotating the Large Aperture Telescope Receiver ([LATR](#)) in sync to maintain consistent orientation between the [LATR](#) and the mirrors. The elevation nod could be done either with the corotator held fixed and just the focal plane projection rotating during the oscillation, or with the corotator oscillating as well.

The frequency of the oscillation is limited by the maximum scan rate (1 deg/s); 0.15 Hz (equating to a period of 6.7s) is used here. The starting phase of the elevation oscillation is randomized.




Figure 8.7: The hit map for the Type 3 (high cadence, varying el.) scan.

8.3 Technical Constraints and Concerns

Technical constraints for the [LAT](#) are laid out in reference [216], and are summarized in table 8.1.

LAT Platform Considerations (Plat.)

Since there are no hard constraints expected from the platform control software, the [LAT](#) platform performance defines the absolute scan bounds on azimuth and elevation specified in table 8.1.

The platform itself only restricts elevations between -90° and 180° . However, a ground-keep region at some elevation will be required. [ACT](#) used a small amount data from an elevation of 30° , but most of its data came from $\geq 40^\circ$. As such, the elevation range is considered to be 40°

LAT Scan Constraints		Set By
Sun Avoidance Radius	30°	Plat.
Moon Avoidance Radius	none	Plat.
Azimuth:		
Range	−180° to 360°	Plat.
Maximum Speed	$\approx 1 \text{ deg/s} / \cos(\theta_{\text{el}})$	Det.
Maximum Turnaround Acceleration	6 deg/s ²	Plat.
Maximum Acceleration During Scan	6 deg/s ²	Plat.
Elevation:		
Range	40° to 140°	Plat.
Maximum Speed	3 deg/s	Plat.
Maximum Acceleration	6 deg/s ²	Plat.
Maximum Elevation Modulation Peak-to-Peak	3°	Det.

Table 8.1: A summary of the technical constraints for the [LAT](#) scan strategy. Table format is matched to table 1 in [\[216\]](#).

to 140°.

Sun and Moon avoidance radii have not been finalized—a Sun avoidance of 30° will be used initially, but will be adjusted if contamination is detected in [ISO](#) data. The pointing specifications are guaranteed up to a 90° Sun avoidance radius, although [ACT](#) used some data up to a 30° avoidance. The Moon will be cut during post-processing, with data flagged if it falls within 5° of the Moon.

Cryogenic Considerations

Any accelerations excite vibrations modes, which will cause heating in all the cryogenic stages over time, although the final (mK) stage will be the most important to protect—it is the most susceptible to heating since it has the smallest cooling power. The magnitude of the heating effect will be determined by in-situ testing, but qualitatively, a higher acceleration will have a larger heating effect, as will prolonged acceleration.

Elevation oscillation could further impact cryogenic performance through atmospheric loading variation, and changes to the efficiency of the pulse-tube coolers ([PTCs](#)).

There are concerns around the turnaround periods in the Type 1 scans—these have caused some heating issues for [ACT](#), which has a smaller (and less heating-sensitive) cryostat than [SO](#). These concerns become more complicated for Type 2 scans, which have almost constant non-zero acceleration. Type 3 scans increase the acceleration by introducing elevation oscillations. If the corotator oscillates at the same time as the elevation nod, the cryostat is affected by movement in two directions. It has been suggested that the [LATR](#) could instead be kept in the same rotation position, so that the only movement is in elevation. This stabilization of the

instrument might mitigate the cryogenic concerns, since the additional impact from corotator movement is eliminated.

Detector Considerations (Det.)

The scan speed is constrained by the detector time constants, which describe how quickly the detectors can respond to changes in incoming light—scanning too quickly can cause blurring and signal distortion. To avoid this, the frequency response of the detector, characterized by the 3dB roll-off frequency f_{3dB} ,[‡] should match the minimum sampling rate for the beamwidth.[§]

This constraint is most important at the highest frequencies, where the beam size is smallest and the telescope is sensitive to the smallest angular scales. At a minimum $f_{3dB} = 200$ Hz, the scan speed $v_{az} < 1 \text{ deg/s} / \cos \theta_{el}$, where θ_{el} is the observing elevation. Measurements show slower time constants than estimated, but the constraint above is sufficient to guide scan strategy development, provided we can tolerate some blurring in the 220/270 GHz channels.

Elevation oscillation in Type 3 scans will affect the loading (the effective integration time per sky region), at a level which is small but not insignificant: an oscillation of $\sim 5^\circ$ at an elevation of 50° makes a change of $\sim 10\%$. This variation in loading can lead to non-uniform sky coverage, as the telescope spends slightly more time observing certain regions, particularly those near the edges of the scan path. While this effect is modest, it could potentially be used for calibrating detector responsivity. The periodic change in loading creates a modulation pattern in the data, allowing for more precise calibration of the detectors by analyzing how their response varies with elevation. This could help correct for any elevation-dependent biases in the measurements.

Thermal variations from cryogenics will affect detector performance.

Optical Considerations

The Type 3 scan can be performed in two distinct ways: by varying the receiver co-rotator with the elevation nods or without varying the receiver co-rotator. Varying the co-rotator with the elevation nods enables the same optical performance as the nod is performed, but comes with higher technical risk and risk for additional vibrational modes and heating from the constant acceleration of the cryostat. Performing the elevation nods without the co-rotator would mitigate some of the technical risks, but it would come at the cost of slight variations

[‡]The point at which the detector’s frequency response drops to half of its maximum power.

[§]From the Nyquist condition, we need to be able to sample twice per beam area.

in the optical performance, including beam and pointing variations. The effects are small and can be modelled and removed, but uncertainties in their fits can cause additional systematic contributions. In both cases, the pointing can vary with the elevation nod due to the telescope sag. Variations in the polarization angle calibration are also a risk. If the pointing changes are accounted for, the polarization angle calibration is dealt with too (subject to some residuals)—we expect these effects to be small, and they can be modelled based on the optics. Note that more difficult polarization angle calibration could have implications for measurements of the bispectrum.

Technical Concerns

We assess the technical constraints and considerations of each scan type to inform the simulations and assess their technical risk. In general, the Type 1 scans have the lowest risk because they have been previously demonstrated. The Type 2 scans have additional risk from the varying scan speed in the azimuthal direction, which has not been demonstrated on previous experiments. The Type 3 scans have the highest technical risk as they not only have the varying scan speed in azimuth like the Type 2 scans but also have added risk from the constant elevation nods.

Many constraints from the instrument like vibrational noise and heating can only be fully determined from in-situ testing of the full system on sky. The [LAT](#) commissioning plan includes tests with varying turnaround accelerations for each scan type to tune the final scan parameters and to determine the final technical feasibility of each scan type.

8.4 Methodology

To make an initial assessment of the different scan types, we simulated observing runs to compare against three [FoMs](#): uniformity, cadence, and cross-linking, which were chosen based on the science requirements as discussed in [section 8.1](#). An initial set of simulations was analyzed to determine and quantify the benefits of each type, before the scans were put through more detailed simulations to better assess their scientific performances.

8.4.1 Figures of Merit

Uniformity of depth can be analyzed by looking at the distribution of hits over the map. Numerically, the standard deviation is a useful metric, but visual maps are often a more useful

indicator as they can show discrepancies across the sky.

To compare the cadence of scan strategies, we counted the number of distinct Julian dates each sky pixel was observed on. With the two-month observing break, the theoretical maximum for this variable is approximately 300 separate dates.

We consider the level of cross-linking by evaluating the ACT-style cross-linking statistic, described in [217]. This allows us to generate cross-linking maps, by using the projected scan angle to represent a polarization angle:

$$P_{\times} = \frac{\sqrt{Q_{\times}^2 + U_{\times}^2}}{I_{\times}}, \quad (8.1)$$

where $Q_{\times}$ and $U_{\times}$ are the Stokes parameters, and $I_{\times}$ is the intensity (here, the total number of hits). A value of $P_{\times} = 0$ represents uniform scan directions and $P_{\times} = 1$ a single scan direction (i.e. no cross-linking). Since power spectra analyses are done on cut maps, given some minimum cross-linking value (see figure D.1), this FoM is a priority.

Additional figures of merit considered were the sky fractions f_{sky} , which can be extracted from the hit maps:

$$f_{\text{sky}}^{\text{raw}} \equiv \frac{n_{\text{hit}}}{n_{\text{tot}}}, \quad f_{\text{sky}}^{\text{NN}} \equiv \frac{\langle N_{\text{hits}} \rangle^2}{\langle N_{\text{hits}}^2 \rangle}, \quad f_{\text{sky}}^{\text{SN}} \equiv \frac{\langle N_{\text{hits}} \rangle \langle N_{\text{hits}}^2 \rangle}{\langle N_{\text{hits}}^3 \rangle}, \quad f_{\text{sky}}^{\text{SS}} \equiv \frac{\langle N_{\text{hits}}^2 \rangle^2}{\langle N_{\text{hits}}^4 \rangle}. \quad (8.2)$$

8.4.2 Initial Simulation Run

The initial run of simulations spanned a full observing season, only simulating the $1/f$ instrumental noise, details of which can be found in reference [218]. All four of the LAT medium-frequency (150 Hz) optics tubes (OTs) were simulated, with pixel thinning applied at a factor of 1/16. The simulated time-ordered data were sampled at 10 Hz, providing sufficient temporal resolution to trace the telescope scanning strategy and sky signal variations. The resulting maps were generated using the HEALPix pixelization scheme with a resolution parameter of $N_{\text{side}} = 1024$, corresponding to a pixel size of approximately 3.4 arcminutes. The time-ordered data were processed with the Madam mapmaker [219], a destriping-based algorithm that estimates and removes low-frequency noise baselines. This step mitigates striping artifacts introduced by $1/f$ noise. Madam served as a computationally efficient placeholder for the more accurate maximum likelihood (ML) mapmaking pipeline that is under development for future data processing. A summary of the simulation results is given in table 8.3, with maps of the hits, cadence, and cross-linking given in Appendix D.1.

From these initial simulations it was clear that the Classical scan (Type 1.0) was heavily

disfavoured against all figures of merit, as well as eliminating Galactic science by virtue of its sky coverage. It was therefore ruled out from further consideration for the SO baseline scan strategy.

As hoped, the high cadence strategies (Types 1.2, 2, 3) brought noticeable improvements over the tiling scans in average cadence (see figure 8.9)—it was noted that Type 1.1 scans have sufficiently poor cadence properties to significantly hinder transient science.

The Type 3 scan gave much better cross-linking than any of the other scans, although the three-elevation scan (Type 1.2.1) had a marked improvement over the remaining flavours of scan (see figure 8.8). However, Type 1.2.1 had significantly poorer uniformity, as can be seen from the hit maps above, and it was decided that this flavour needed further refinement to improve the uniformity of the coverage without sacrificing the improved cross-linking.

It was also noted that more detailed simulations would be required for fuller analysis of the scientific performances of the scan types. As such, a further set of simulations was commissioned, to include atmospheric noise and ML map-making capabilities.

Scan Strategy		f_{sky} Values		
Scan Type		$f_{\text{sky}}^{\text{raw}}$	$f_{\text{sky}}^{\text{NN}}$	$f_{\text{sky}}^{\text{SS}}$
1.0	Classical	0.621	0.361	0.101
1.1.0	Large Tiles	0.616	0.561	0.465
1.1.1	Small Tiles	0.616	0.561	0.481
1.2.0	High Cadence, 1 el.	0.607	0.571	0.502
1.2.1	High Cadence, 3 el.	0.607	0.509	0.415
2	High Cadence, varying speed	0.607	0.577	0.538
3	High Cadence, oscillating el.	0.614	0.577	0.538

Table 8.2: The f_{sky} values for the initial simulation run.




Figure 8.8: Maps of the cross-linking statistic P_x for each scan strategy considered in the initial run of simulations. A value $P_x = 0$ gives perfect cross-linking, and $P_x = 1$ gives zero cross-linking (a pixel observed only from one angle).

Scan Type		Previously Demonstrated?	Comments
Type 1:			
1.0	Classical	Yes	Ruled out
1.1.0	Large Tiles	Yes	Poor uniformity
1.1.1	Small Tiles	Yes	Poor uniformity
1.2.0	High Cadence (1 el.)	No	Improved uniformity over Types 1.0, 1.1
1.2.1	High Cadence (3 el.)	No	Improved uniformity and cross-linking over Types 1.0, 1.1, 1.2.0
Type 2:			
2	High Cadence (varying speed)	No	Improved uniformity over Type 1 scans
Type 3:			
3	High Cadence (oscillating el.)	No	Improved uniformity over Type 1 scans, improved cross-linking over Type 2 scans

Table 8.3: A summary of the different scan types, following analysis of the initial round of simulations.




Figure 8.9: Maps of the cadence (number of individual days hit) for each scan strategy considered in the initial run of simulations.

8.4.3 Optimization of Type 1.2.1

While the three-elevation high cadence scan (Type 1.2.1) showed promise, offering improved cross-linking over the single-elevation scans while avoiding the technical concerns with Types 2 and 3, it had not been optimized in the initial round of simulations. In particular, further tuning was required to improve uniformity in depth and cadence, as well as to reduce noise and improve systematic performance at different elevation splits.

The original scan had its observing time split equally between the three elevations. In optimizing, a wide range of different observing splits were considered, with time spent at the lowest elevation (40°) ranging between 30–75%. The four most promising variations were studied in detail—these had elevation splits of 60:20:20, 67:17:17, 68:23:10, and 69:17:14, at $40/50/60^\circ$, respectively, and were compared against the one elevation scan (Type 1.2.0) and the one elevation scan with varying speed (Type 2).

The same FoMs as listed in section 8.4.1 were used for optimization, with a focus on retaining good f_{sky} values and cross-linking, while improving the uniformity (of both hits and cadence). An additional metric, the ‘elevation noise’ was analysed:

$$\text{noise} = \frac{A}{\sin(\theta_{\text{el}})} + C, \quad (8.3)$$

where θ_{el} is measured in radians, and the coefficients A and C depend on the frequencies measured by the OTs, and are given in table D.1. This equation is found by the modelling of

noise at a range of elevations, which is then normalized to the noise at 50° for each frequency band.

Quantitative measures of the performances of the four best-performing three-elevation (Type 1.2.1) scans, as well as the one-elevation (Type 1.2.0) and varying speed (Type 2) scans, are shown in table 8.4. It is clear from the results that the compromise for vastly improved cross-linking will be a smaller scan (which can be seen in the lower f_{sky} values—less sky is available for observation at the higher elevations) with lower average cadences. The optimized three-elevation scan chosen has a time split of 60:20:20, at 40° , 50° , 60° , respectively.

Table 8.4

Scan Strategy		f_{sky} Values		
Az. speed	Elevations	$f_{\text{sky}}^{\text{raw}}$	$f_{\text{sky}}^{\text{NN}}$	$f_{\text{sky}}^{\text{SS}}$
constant	1	0.607	0.571	0.502
constant	3 (60:20:20)	0.586	0.548	0.470
constant	3 (67:17:17)	0.586	0.555	0.489
constant	3 (68:23:10)	0.586	0.560	0.500
constant	3 (69:17:14)	0.586	0.561	0.503
varying	1	0.607	0.577	0.538

(a)

% of Observed Sky Retained						
Cut-Off Value, $P_{\times, \text{max}}$						
0.7	0.75	0.8	0.85	0.9	0.95	0.99
30.1	37.2	42.8	49.8	59.6	71.4	87.6
46.7	50.4	56.8	65.6	75.3	82.4	99.8
45.4	48.9	54.9	63.6	73.8	83.5	99.5
44.1	47.2	53.4	62.9	73.1	82.4	99.8
44.4	47.7	53.7	63.0	73.3	83.2	99.8
31.4	37.7	42.9	49.7	59.4	71.2	87.4

(b)

Continued on next page

Cadence				Pixel Hits	
1 day mean	$\pm 1\sigma$	3 day mean	$\pm 1\sigma$	mean	$\pm 1\sigma$
219.5	32.0	214.6	33.4	20367	5111
169.3	31.6	124.8	40.3	20948	5492
164.3	29.1	131.9	34.8	20976	4936
171.8	29.5	137.6	34.0	20992	4556
174.2	29.4	142.5	30.4	20976	4468
218.9	32.6	213.9	34.4	20445	4683

(c)

Elevation Noise					
LF1	LF2	MF1	MF2	UHF1	UHF2
1.0000	1.0000	1.0000	1.0000	1.0000	1.0000
0.9923	0.9799	0.9801	0.9788	0.9647	0.9570
0.9936	0.9833	0.9834	0.9823	0.9706	0.9641
0.9944	0.9949	0.9948	0.9801	0.9600	0.9685
0.9942	0.9945	0.9940	0.9796	0.9595	0.9675
1.0000	1.0000	1.0000	1.0000	1.0000	1.0000

(d)

Table 8.4: Comparison of the four leading three-elevation (Type 1.2.1) scan variations, against the constant (Type 1.2.0) and varying (Type 2) speed one-elevation scans.

(a) The f_{sky} values.

(b) The % of observed sky which remains after a cross-linking cut, for a range of $P_{\times \text{max}}$. Note that $P_{\times \text{max}} = 0.7$ retains most regions observed with two orthogonal scans [217].

(c) The means and standard deviations of cadence and hits.

(d) The elevation noise, normalised to 50° , for different frequency bands.

Figures 8.10 and 8.11 compare the optimized three-elevation scan (Type 1.2.1) with the one-elevation (Type 1.2.0), varying speed (Type 2), and large tiles (Type 1.1.0) scans. The improvement of the optimized scan over Types 1.2.0 and 2 with respect to cross-linking can be seen in the significantly larger proportions at low $P_{\times}$, and while the cadence is slightly poorer than Types 1.2.0, it is still a significant improvement over the large tiles (Type 1.1.0) scan.




Figure 8.10: Maps showing the cross-linking statistic $P_{\times}$ at each pixel.




Figure 8.11: Maps showing the cadence at each pixel.

8.4.4 Improved Simulation Run

An improved set of simulations was then run for the remaining scan types (and with the optimized 1.2.1 scan), this time with the inclusion of coarse atmospheric noise, and at a higher resolution of $N_{\text{side}} = 4096$. As before, only the 150 GHz band was simulated. The white noise and atmospheric noise maps were simulated separately to enable scaling between different frequency bands, since each scales differently. These simulations spanned a whole observing season, but were made up of individual observing days spread evenly across the run. This guarantees that the seasonal effects of Solar avoidance and atmospheric noise are captured in the data set,[¶] while reducing the computational cost.

The simulated data were then reduced with an ACT-style, [ML](#) mapmaker to make the residual noise as realistic as possible. To limit the cost of the simulation, the conjugate gradient iteration in the mapmaker was terminated after 100 iterations. This was enough to reduce the residual noise by a factor ~ 105 , sufficient for this study.

Other improvements to the simulations were made:

- A reduction in the focal plane decimation factor, from 16 to 8 (i.e. 1/8 pixels in the focal plane simulated, up from 1/16),
- An increase in the sampling rate, from 10 Hz to 100 Hz,
- An improvement in the resolution, from 3.4 arcmin to 0.5 arcmin,

In simulating the noise for the oscillating elevation scan (Type 3), it was found that the scan introduced a strong elevation-modulated signal which was complicated and difficult to handle in the mapmaker. This prohibited further analysis of this scan type. It was not clear whether this complexity was real, or an artefact of the simulations, and on-sky tests are therefore needed to understand the noise properties of this promising scan type.

The maps from the improved simulation run can be seen in [Appendix D.2](#). The different scan types follow the same patterns as before (which are described in [table 8.3](#)). With respect to the cadence specifically, the optimized Type 1.2.1 scan observes the same sky pixel on ~ 50 fewer observing days than the Type 1.2.0 scan but still on ~ 170 days. That is 100 more days than the large tile scan (with the small tile scan having even poor cadence than large tiles). Both of the Type 1.1 scans have sufficiently poor cadence properties to hinder transient science, and are also recommended to be removed from selection.

[¶]Weather parameters (temperature, wind, precipitable water vapour ([PWV](#))) were fixed to historical median values for the season and time of day rather than sampling from historical distributions.

8.5 Scientific Performance

8.5.1 Transients, Galactic Science, SZ

Cadence is the most important factor for transient science, performances of which can be determined by the average cadences, as well as the spread of cadence across the map. The scan types are preferentially ranked:

- Type 3 (High Cadence, Varying Elevation)
- Type 2 (High Cadence, Varying Speed)
- Type 1.2.0 (High Cadence, One Elevation)
- Type 1.2.1 (High Cadence, Three Elevations)
- Type 1.1.1. (Small Tiles)
- Type 1.1.0 (Large Tiles)
- Type 1.0 (Classical)

Galactic science requires adequate observing time spent looking at the Galaxy, and as such disfavours any approaches which de-prioritize Galactic observations—here, that is just the Classical scan.

Analysis was performed to obtain the CO line, an important tracer of Galactic structure and foreground contaminant of the [CMB](#), for each scanning strategy. A ranking was produced in order of the detectable polarization of the CO line emission at 3σ . This was done for a few areas in the Galactic plane with known CO emission, and as a 3σ upper limit outside the Galactic plane. A CO map noise power spectrum was also produced.

No major difference was found between the majority of the scanning strategies, save for the Classical scan which was disfavoured. The three-elevation scan (non-optimized, as this analysis was carried out for the initial round of simulations) performs best in terms of the polarization fraction, although has a slightly worse noise level on the CO power spectrum.

Minimising variation in the noise across the map is generally beneficial to Sunyaev-Zeldovich effect ([SZ](#)) science, as well as maximizing the footprint overlap with Dark Energy Spectroscopic Instrument ([DESI](#)), which is at $\text{Dec} > -10^\circ$. The small tiles scan (Type 1.1.1) has the most uniform coverage and was therefore preferred, while the slightly lower uniformity of high cadence strategies (due to their Sun avoidance patterns) down-weights those types. No major changes in overlap with [DESI](#) and different scanning strategies were noted.

8.5.2 Lensing

As discussed in section 2.4.4, lensing is a significant foreground for CMB polarization, and its impact on the analysis was considered through enhanced map-domain noise simulations. Two key metrics were derived: (1) the amplitude of the noise-only lensing power spectrum, and (2) the chi-square values from noise-only null tests.

The first metric is closely linked to the number of well-cross-linked pixels, which reflects the quality of the scan strategy. As expected, the results follow the cross-linking analysis, with the optimized three-elevation scan (Type 1.2.1) being the preferred option.

The second metric ranks scan strategies based on their deviation from zero. Again, the three-elevation scan emerges as the preferred strategy.

The Type 3 scan was not analysed due to difficulties in fitting its noise properties. However, based on the metrics above, it is expected to outperform Type 1.2.1.

8.5.3 Power Spectrum

Analysis on the effects of the different scanning strategies on the power spectrum relied on two metrics: the variance of the power spectrum at a realistic fixed cross-linking cut; and the difference between power spectrum variance computed from simulations, and using typical analytic approximations. Map-based simulations of the expected noise were used for estimations of performance against these metrics, based on the initial noise simulations.

Against fixed cross-linking cuts of 0.96 (a typical cut) and 0.7 (an aggressive cut), most scan strategies performed similarly, but the one-elevation (Type 1.1.0) and varying speed (Type 2) scans perform particularly badly. This is due to their large areas of the sky with poor cross-linking, which leads to more sky being lost after the fixed cuts.

Computing the covariance of the power spectrum analytically is only possible under assumptions, but requires orders of magnitude less computation time. Scanning strategy can influence these assumptions by e.g. varying the noise power spectrum over the sky. As such, any of the remaining scans except the varying speed (Type 2) scans would be acceptable.

The Type 3 scan was not analysed due to the complex noise properties, but again, based on these metrics, we expect it to perform best.

8.6 Summary and Discussion

From the seven initial scan types, a few were eliminated as options: the Classical scan was heavily disfavoured by the [AWGs](#), and precluded Galactic science; and the tiled scans (Types 1.1.0 and 1.1.1) were not strongly preferred, and are very detrimental to transient science. The oscillation nod scan (Type 3) showed promise, despite the highest technical risk, and warrants on-sky testing to better understand its noise properties to assess feasibility.

The nominal strategy recommended to (and accepted by) the collaboration was the Type 1.2.1 optimised 3-elevation high cadence scan. It performs better than all other scans in lensing and power spectra analyses, mostly due to large, even sky coverage with good cross-linking (although different map-making techniques which are less dependent on cross-linking may find different results). It has low additional technical risk over existing strategies—it is a wider scan than previously demonstrated on-sky, but does retain constant elevation and azimuthal speed.

The Type 1.2.1 scan does have reduced uniformity compared to Types 1.2.0 and 2, which could make [SZ](#) and Galactic science analyses slightly more difficult. It has a slightly reduced cadence, impacting transient science: 10% of the sources will be hit less frequently than in the other scans, however it still provides a large improvement over the cadence from [ACT](#). These small hits to [SZ](#), Galactic, and transient science cases are felt to be outweighed by the improved performance in lensing and power spectrum analyses.

A further [LAT](#) scan strategy review will be run once [ISO](#) data has been taken, to set the scan strategy for near-term observations.

A scan strategy which benefits the lensing analysis is crucial in making good measurements of the lensing foreground of the primordial [GWB](#) for analysis with the Small Aperture Telescopes ([SATs](#)), and one which maximises detail in the temperature power spectrum at high ℓ will increase [SO](#)'s sensitivity to the damping effects of the primordial [GWB](#). Refer back to section [5.3.3](#) for details on how the [CMB](#) can be used to detect the primordial components of the stochastic gravitational wave background ([SGWB](#)).

Chapter 9

Small Aperture Telescope Scan Strategy

This chapter details a campaign of work on the SAT scan strategy, to which I made significant contributions as part of the low- ℓ B-modes (BB) analysis group in SO, in collaboration with Sara Simon, Ken Ganga, Reijo Keskitalo, Aashrita Mangu, David Alonso, and Akito Kusaka. Significant portions of the text are based on reference [220], a set of review slides prepared for internal SO discussion, and which will ultimately form the basis for a paper. I made substantial contributions to the campaign and review slides in both analysis and writing. My specific contributions to the analysis were in evaluating the sensitivities of each proposed scan variation to the tensor-to-scalar ratio, including making improvements to the BB forecasting pipeline by adding purification.

9.1 Introduction and Science Goals

The Large Aperture Telescope (LAT) and Small Aperture Telescope (SAT) scan strategies differ significantly due to their distinct instrumental designs and scientific goals. The LAT, with its high angular resolution, is optimized for small-scale measurements of the cosmic microwave background (CMB), which requires wide sky coverage and good cross-linking to reduce noise and systematics. As a result, its scan strategy emphasizes stability, uniformity, and large-area mapping. In contrast, the SATs are designed for sensitive measurements of large-scale B-mode polarization, particularly targeting the tensor-to-scalar ratio r . This requires deep integration on relatively small sky patches, with minimal atmospheric contamination.

The baseline strategy considered in the Simons Observatory (SO) white paper [79] is the

‘opportunistic scheduler’ detailed in reference [90], which divides the sky into a Southern and Northern patch in low foreground areas, each split into ten tiles—see figure 9.1 for the footprint of the baseline scan. This scanning strategy was designed to maximize efficiency by constantly varying the elevation, azimuth throw and boresight direction. The primary limiting systematic for the SATs is typically ground pickup, and to help integrate this down, it would be better to move to a strategy with fixed azimuth ranges and elevation—this will allow us to understand the scan synchronous signal better. There remains scope to add more configurations to the schedule later, if data look good, to gain additional efficiency.

Other limitations to the original scanning strategy also follow from frequent changes in elevation: it requires rebiasing the detectors after every elevation change, which takes ~ 20 minutes. Assuming elevation changes once per hour leads to $\sim 25\%$ of observing time being spent rebiasing. However, time lost to cryogenic settling after changing elevations has been found to be negligible.

To develop the new scanning strategy, it was necessary to evaluate which observing elevation, azimuth throw, and boresight directions perform best against a series of figures of merit (FoMs): initial FoMs included the f_{sky} fractions and elevation noise metrics defined in section 8.4.1.*

9.2 Technical Considerations

The SAT instrument constraints may be more conservative than the designed performance, due to noise and vibrations introduced by the scans, and issues with cable wrap degradation. During initial science observations (ISO) a number of implementations will be tested to better understand the performance.

The minimum observing elevation, set by the shielding and ground pickup, is 50° . There is no instrumental maximum elevation, but the available sky coverage shrinks with elevation, and in practice we do not expect to observe above 70° . In addition, the half-wave plate (HWP) sags with elevation angle (see figure 17 from reference [221]). However, this is only at low observing elevations, and it does not have a worrying impact above the existing 50° minimum elevation discussed above.

The cooling power of the pulse-tube cooler (PTC) varies with tilt angle (see figure 8 in reference [222]). On the SAT, the PTC sits at 27.5° from the zenith, if the window is pointing at the zenith, which limits the boresight rotation to $\pm 70^\circ$.

*The half-wave plates (HWPs) for the SATs perform such that cross-linking is not a relevant metric for performance.

A 45° minimum Sun and Moon avoidance is required, since the SAT baffle is illuminated if sources are any closer.

9.3 Proposed Strategies

A large number of strategies and observation directions were initially considered, with convergence on the most simple (‘rising/setting’) scans for systematic mitigation: these scans have fixed azimuth ranges, a single observing elevation, and limited numbers of boresight directions (the azimuth at which the boresight of the SAT is pointing), all of which limit the systematics involved.

Different orders of polynomial filtering for the ground pickup were initially explored—that is, filtering out noise, beginning with a constant term and continuing up to the specified polynomial. As a low-amplitude signal, the CMB can only be measured after filtering, and the higher the polynomial order, the more noise is removed. However, some CMB signal is removed alongside the noise, with low- ℓ terms becoming unavailable.

Scan strategies were explored with a range of parameters, which fulfil the technical constraints discussed in section 9.2:

- Observing elevation [50°, 55°, 60°],
- Azimuth throw [20°, 30°, 40°],
- Ground pickup polynomial order [1, 2, 3],
- Scan speeds [1, 2] °/s,
- Boresight directions: discrete number of 1-4 azimuth ranges [(S), (E), (SW-SE), (W-E), (SW-N-SE), and (SW-SE-NW-NE)].

Of these six initial sets of boresight directions, three were ruled out (discussed in section 9.5.5), to leave the three most promising cases:

SW-SE (2 Azimuth)

This strategy has two central boresight directions: $az = 120^\circ$ & 240° (where North is at $az = 0^\circ$). The first priority is the Southern CMB field. When that isn’t permitted (e.g. due to Sun or Moon avoidance) the Galaxy is observed at the same azimuth and elevation. Only the Southern




Figure 9.1: The original 2018 baseline scanning strategy for the SATs, with a Mollweide projection in Equatorial co-ordinates (as are all subsequent maps, unless stated otherwise).

(a) The footprint overlaid on a map of Galactic dust, with the footprints of SPIDER and BICEP.

(b) A hit map, showing the variation in depth of observations across the observed sky. Figure credit for both: SO Collaboration.

field is observed (see figure 9.2), with both rising and setting scans—these provide a very strong null test, since the same areas are observed from both directions.




Figure 9.2: A hit map of the SW-SE scan strategy (at an elevation of 50° , with an azimuth throw of 40°), showing the variation in depth of observations, as well as the overall footprint of the scan.

SW-N-SE (3 Azimuth)

This strategy has three central boresight directions: $az = 0^\circ, 120^\circ, \& 240^\circ$. The first priority remains the Southern field, with the Northern field as the next priority, followed by the Galaxy. Southern and Northern fields are observed, with both fields in lower foreground regions than previous scans (compare the dusty regions shown in figure 9.1a with the footprint of this scan

in figure 9.3), with the Southern field as the lower-foreground of the two.




Figure 9.3: A hit map of the SW-N-SE scan strategy (at an elevation of 50° , with an azimuth throw of 40°).

SW-NE-NW-SE (4 Azimuth)

This strategy has four central boresight directions: $az = 75^\circ, 120^\circ, 240^\circ, \& 285^\circ$. Again, the first priority is the Southern field, with the Northern field as the next priority, followed by the Galaxy. Southern field and Northern fields are observed, with footprints which are more similar to the baseline SAT fields from [90]. This scan variation also has rising and setting scans for its whole footprint.




Figure 9.4: A hit map of the SW-NE-NW-SE scan strategy (at an elevation of 50° , with an azimuth throw of 40°).

9.4 Methodology

Simulated scan runs were carried out to generate hit maps (incorporating the Planck Galactic mask, expected to be used for real data) and simplified transfer functions. Initial FoMs assessed were the f_{sky} values (described in section 8.4.1), as well as the observing efficiencies. Simulations were run to be representative of a full year’s observing, including a break between 15/01–15/03. As such, the maximum observing efficiency (the proportion of time spent observing) and the CMB observing efficiency (the proportion of time spent observing the CMB) is 83%.

Hit maps were generated, which included Sun and Moon avoidances of 45° , but with no scan synchronous signal, or its removal. However, characterization of the scan synchronous signal has been explored during ISO, and is discussed in [223].

9.4.1 Approximating the Transfer Function

The transfer function characterizes the spherical harmonic modes which can be extracted from observations. Simplified transfer functions were produced by calculating the recovered variance. One hour of timeline data was simulated for each scan variation. The variance for each spherical harmonic $Y_{\ell m}$ was calculated pre- and post-filtering. A transfer function over ℓ was then built by summing the variance over m for each ℓ , done for both pre- and post-filtered data.

This is only a rough approximation, and notably does not include scan synchronous signal. Work on this is ongoing in the calibration, systematics and sensitivity (CSS) working group.

It was found that the polynomial order used for filtering only made an impact during post processing, and hence was eliminated as a parameter in the forecasting; decisions on the order will be made following ISO.[†]

9.4.2 Low- ℓ analysis

The primary science FoM for the SATs is $\sigma(r)$, the predicted constraint on the tensor-to-scalar ratio r . Forecasting $\sigma(r)$ requires realistic modelling of how scan strategy, instrumental response, and noise inhomogeneity affect the measurement of the B -mode power spectrum. In particular, non-uniform sky coverage due to varying integration time and filtering must be accounted for.

The pipeline for evaluating $\sigma(r)$ takes the f_{sky} values defined in equation 8.2, and transfer functions T_ℓ described above as inputs. Note that the transfer functions could only be calculated

[†]There is also some manoeuvrability between stronger filtering (which reduces the number of modes available for analysis) and wider throws (which increase the number). This will be explored during ISO.




Figure 9.5: The (un-normalized) transfer function for the 50° elevation, 40° throw, SW-SE scan. The blue and orange lines show the transfer function pre- and post-filtering, respectively. Figure credit: Ken Ganga, SO Collaboration.

for individual boresight directions, with mirrored directions having identical transfer functions (e.g. SW and SE). As such, only some scans were provided with transfer functions matching all of the boresight directions involved. In cases where there were multiple non-mirrored directions, the transfer function which matched the Southern directions was chosen (since time spent scanning in these directions dominates).

The $\sigma(r)$ Pipeline

The pipeline divides the observed sky into patches, assuming that each has approximately homogeneous noise. In patch p , the noise power spectrum is given by:

$$N_\ell^p = \frac{\sigma_h^2 \Omega_p}{N_{\text{hits},p}} \equiv \bar{N}_\ell \frac{\langle N_{\text{hits}}^2 \rangle}{\langle N_{\text{hits}} \rangle N_{\text{hits},p}}, \quad (9.1)$$

where σ_h^2 is the per-hit noise variance, Ω_p is the patch area, and $N_{\text{hits},p}$ is the number of observations in patch p . We can then define the average noise level:

$$\bar{N}_\ell = \frac{\sigma_h^2 \Omega_p}{\langle N_{\text{hits}} \rangle} \frac{\langle N_{\text{hits}} \rangle^2}{\langle N_{\text{hits}}^2 \rangle}. \quad (9.2)$$

To compute the full-sky power spectrum, we average over patches using weights based on the inverse noise variance:

$$\hat{C}_\ell = \sum_p w_p \hat{C}_\ell^p, \quad \text{where} \quad w_p = \frac{N_{\text{hits},p}^2}{\sum_{p'} N_{\text{hits},p'}^2}. \quad (9.3)$$

The variance of the power spectrum estimator is then:

$$\begin{aligned}\text{Var}(\hat{C}_\ell) &= \sum_p w_p^2 \text{Var}(\hat{C}_\ell^p) \\ &= \frac{2}{(2\ell + 1)\Delta\ell} \sum_p w_p^2 (S_\ell + N_\ell^p)^2,\end{aligned}\tag{9.4}$$

where S_ℓ is the signal, and N_ℓ the noise, power spectrum. Substituting the expressions for w_p and N_ℓ^p leads to:

$$\text{Var}(\hat{C}_\ell) = \frac{2}{(2\ell + 1)\Delta\ell} \left[\frac{S_\ell^2}{f_{\text{sky}}^{\text{SS}}} + 2 \frac{S_\ell \bar{N}_\ell}{f_{\text{sky}}^{\text{SN}}} + \frac{\bar{N}_\ell^2}{f_{\text{sky}}^{\text{NN}}} \right],\tag{9.5}$$

where $f_{\text{sky}}^{\text{SS}}$, $f_{\text{sky}}^{\text{SN}}$, and $f_{\text{sky}}^{\text{NN}}$ are those given in equation 8.2.

To propagate the impact of scan strategy and time-domain filtering, noise power spectra are constructed using the transfer functions detailed in section 9.4.1:

$$N_\ell = \frac{N_{\text{white}}}{T_\ell}.\tag{9.6}$$

These N_ℓ curves are then passed to the **BBPower** forecasting tool, which evaluates expected $\sigma(r)$ through a likelihood-based approach. The pipeline models the observed B -mode power spectrum as a sum of cosmological signal, foreground contributions, and instrumental noise (N_ℓ). The pipeline uses that C_ℓ^{BB} spectra, to build a Gaussian likelihood.

This pipeline provides a forecast of $\sigma(r)$ for a given scan, giving an estimate of the impact of mode loss as encapsulated in the transfer function. It assumes only mask-based E/B purification to mitigate leakage in the power spectrum estimation. A matrix-based purification method—which has subsequently finished development—would address leakage at the map level, and is expected to further improve accuracy at low ℓ .

Note that the pipeline does not consider the footprint shape (although the overall footprint size is encoded in the $f_{\text{sky}}^{\text{raw}}$ value). Qualitatively, a narrower footprint can be assumed to have a negative impact, since it would require much stronger apodization—increased smoothing of map edges leads to increased mode mixing and leakage.

The forecasting pipeline was run with Fisher matrix sampling, with a synchrotron and dust foreground (see section 2.4.5), for predicted **SO** five-year noise levels (for each of the LF, MF, and UHF frequency bands). It was run with a HEALPix resolution of $N_{\text{side}} = 512$ (i.e. ≈ 6.8 arcmin), and with mode cuts of $30 \leq \ell \leq 300$ for a fiducial model of $r = 0$. The value for A_{lens} was set at 1—i.e. no delensing. However, a recent study into the delensing capacity of **SO** gives a more optimistic value of ~ 0.35 [63]. This would significantly improve constraints on r .

9.5 Results

9.5.1 f_{sky} Values

f_{sky} values for each of the scan variations considered can be found in table 9.1. The area of sky observed varies widely (see the range between the lowest and highest $f_{\text{sky}}^{\text{raw}}$ values) with smaller but still notable variations in the three other f_{sky} values.

Depending on the regime in which the CMB is being observed, different f_{sky} values become the most relevant FoM. For example, E -mode data is signal driven, and so $f_{\text{sky}}^{\text{SS}}$ dominates analysis, but for B -mode observations, it is largely $f_{\text{sky}}^{\text{NN}}$.

Table 9.1

Elevation	Azimuth Throw	Boresight Direction	$f_{\text{sky}}^{\text{raw}}$	$f_{\text{sky}}^{\text{SS}}$	$f_{\text{sky}}^{\text{SN}}$	$f_{\text{sky}}^{\text{NN}}$
50	30	S	0.118	0.066	0.075	0.082
50	30	SW-SE	0.317	0.125	0.160	0.188
50	30	W-E	0.385	0.188	0.226	0.255
50	30	E	0.385	0.204	0.236	0.261
50	30	SW-NE-NW-SE	0.495	0.158	0.220	0.250
50	30	SW-N-SE	0.543	0.139	0.216	0.288
50	40	S	0.126	0.071	0.079	0.086
50	40	SW-SE	0.349	0.131	0.150	0.203
50	40	W-E	0.416	0.197	0.240	0.273
50	40	E	0.416	0.210	0.248	0.278
50	40	SW-NE-NW-SE	0.512	0.169	0.236	0.290
50	40	SW-N-SE	0.550	0.148	0.231	0.309
55	30	S	0.131	0.075	0.087	0.095
55	30	SW-SE	0.318	0.131	0.165	0.192
55	30	W-E	0.382	0.186	0.223	0.252
55	30	E	0.382	0.201	0.232	0.257
55	30	SW-NE-NW-SE	0.468	0.163	0.220	0.267
55	30	SW-N-SE	0.515	0.148	0.225	0.296
55	40	S	0.139	0.080	0.091	0.099
55	40	SW-SE	0.344	0.135	0.173	0.204
55	40	W-E	0.410	0.192	0.234	0.266

Continued on next page

55	40	E	0.410	0.205	0.242	0.271
55	40	SW-NE-NW-SE	0.489	0.169	0.232	0.282
55	40	SW-N-SE	0.520	0.153	0.235	0.311
60	30	S	0.143	0.084	0.098	0.108
60	30	SW-SE	0.318	0.141	0.173	0.199
60	30	E	0.378	0.199	0.231	0.256
60	30	W-E	0.378	0.184	0.221	0.250
60	30	SW-NE-NW-SE	0.453	0.171	0.223	0.266
60	30	SW-N-SE	0.482	0.159	0.236	0.304
60	40	S	0.148	0.090	0.102	0.111
60	40	SW-SE	0.340	0.142	0.179	0.208
60	40	E	0.403	0.201	0.237	0.266
60	40	W-E	0.403	0.189	0.230	0.262
60	40	SW-NE-NW-SE	0.465	0.177	0.233	0.278
60	40	SW-N-SE	0.484	0.164	0.244	0.315

Table 9.1: The f_{sky} values for each of the SAT scan variations considered.

9.5.2 Elevation

Observing elevations of $[50^\circ, 55^\circ, 60^\circ]$ were evaluated. A small preference for 50° was found in $\sigma(r)$, as can be seen in table 9.2. This was likely due to its smaller f_{sky} values—as can be seen in table 9.1, the 50° elevation scans have consistently low f_{sky} values, with averages of 0.74, and 0.87 of the value of the 60° elevation scans’ $f_{\text{sky}}^{\text{raw}}$, and $f_{\text{sky}}^{\text{NN}}$, respectively.

It was also noted that the 50° elevation scan had slightly better observing efficiency, with average total efficiency of 74%, and average CMB efficiency of 57%, compared to the 60° elevation scan’s 71% and 55%—see table E.1 for the observing efficiencies across all scan variations tested.

The sky patch at 50° is also slightly cleaner in foregrounds, and the SAT instrument was designed to be optimal at 50° . Per the elevation noise model given by equation 8.3, there is slightly more atmospheric noise at 50° . However, the half-wave plate (HWP) is expected to mitigate this to some extent.

In light of these results, 50° was selected as the ideal observing elevation.



Elevation	Azimuth Throw	Boresight Direction	$f_{\text{sky}}^{\text{SS}}$	$f_{\text{sky}}^{\text{SN}}$	$f_{\text{sky}}^{\text{NN}}$	$\sigma(r) \times 10^{-3}$	
2018 optimistic			0.114	0.152	0.184	1.883	
50	40	SW-SE	0.131	0.150	0.203	1.911	
55	40	SW-SE	0.135	0.173	0.204	1.937	
60	40	SW-SE	0.142	0.179	0.208	1.980	
50	40	SW-NE-NW-SE	0.169	0.236	0.290	2.015	
55	40	SW-NE-NW-SE	0.169	0.232	0.282	2.037	
50	40	SW-N-SE	0.148	0.231	0.309	2.065	
60	40	SW-NE-NW-SE	0.177	0.233	0.278	2.074	
55	40	SW-N-SE	0.153	0.235	0.311	2.099	
60	40	SW-N-SE	0.164	0.244	0.315	2.148	
2018 pessimistic			0.114	0.152	0.184	2.412	

Table 9.2: The performances of scans at different elevations, ranked by their forecasted $\sigma(r)$. The ‘2018 pessimistic’ scan corresponds to the baseline scan.

9.5.3 Azimuth Throw

20°, 30°, and 40° azimuth throws on the mount were considered. Initial analysis showed that 20° scans performed significantly worse than 30° and 40° scans, and the 20° throw was dropped from further analysis.

A slight preference for a throw of 40° was found with respect to $\sigma(r)$, as shown in table 9.3, and therefore the 40° throw was selected.

This can be explained by the increased footprint size that a longer throw generates, which captures more low- ℓ modes. This does come at the expense of survey depth, but mode loss has a higher impact than survey depth following filtering.



Elevation	Azimuth Throw	Boresight Direction	$f_{\text{sky}}^{\text{SS}}$	$f_{\text{sky}}^{\text{SN}}$	$f_{\text{sky}}^{\text{NN}}$	$\sigma(r) \times 10^{-3}$	
2018 optimistic			0.114	0.152	0.184	1.883	
50	40	SW-SE	0.131	0.150	0.203	1.911	
50	30	SW-N-SE	0.139	0.216	0.288	1.951	
50	30	SW-SE	0.125	0.160	0.188	1.978	
50	40	SW-NE-NW-SE	0.169	0.236	0.290	2.015	
50	40	SW-N-SE	0.148	0.231	0.309	2.065	
50	30	SW-NE-NW-SE	0.158	0.220	0.250	2.092	
2018 pessimistic			0.114	0.152	0.184	2.412	

Table 9.3: The performances of scans with different throws, ranked by their forecasted $\sigma(r)$.

9.5.4 Scan Speed

There is a trade-off between the lower atmospheric noise (including atmospheric polarization) of a high scan speed, and the higher efficiency of a lower scan speed (which has shorter

turnarounds, and a higher proportion of usable observations).

Scan speeds of 1 and 2 deg/s were evaluated, with no impact on $\sigma(r)$, as can be seen in table 9.4. The decision of the final scan speed was left for evaluation during ISO.



Elevation	Azimuth Throw	Boresight Direction	$\sigma(r) \times 10^{-3}$ speed: 1 deg/s		$\sigma(r) \times 10^{-3}$ speed: 2 deg/s	
						
50	40	S	1.857		1.857	
50	40	SW-SE	1.911		1.911	
50	40	W-E	1.968		1.968	
50	40	E	1.969		1.969	
50	40	SW-NE-NW-SE	2.015		2.015	
50	40	SW-N-SE	2.065		2.065	

Table 9.4: The performances of scans at different speeds, ranked by their forecasted $\sigma(r)$.

9.5.5 Boresight Directions

Six sets of boresight directions were evaluated, with the ranking of their effects on $\sigma(r)$ shown in table 9.5. Three of these were ruled out: the S- and E-only scans, with only one boresight direction permitted, had much higher proportions of Galactic observations; and the E-W scan sat at a Dec with significantly higher dust foregrounds. This left a two (SW-SE), three (SW-N-SE), and four azimuth (SW-SE-NW-NE) scan to decide between.




Elevation	Azimuth Throw	Boresight Direction	$f_{\text{sky}}^{\text{SS}}$	$f_{\text{sky}}^{\text{SN}}$	$f_{\text{sky}}^{\text{NN}}$	$\sigma(r) \times 10^{-3}$	
							
50	40	S	0.071	0.079	0.086	1.857	
2018 optimistic			0.114	0.152	0.184	1.883	
50	40	SW-SE	0.131	0.150	0.203	1.911	
50	40	W-E	0.197	0.240	0.273	1.968	
50	40	E	0.210	0.248	0.278	1.969	
50	40	SW-NE-NW-SE	0.169	0.236	0.290	2.015	
50	40	SW-N-SE	0.148	0.231	0.309	2.065	
2018 pessimistic			0.114	0.152	0.184	2.412	

Table 9.5: The performances of scans with different boresight directions, ranked by their forecasted $\sigma(r)$.

Between the remaining three scans, the two azimuth scan performs best in terms of $\sigma(r)$. Each scan had identical total observing efficiencies of 74%, with a range in CMB observing efficiencies between 53–62% (see table E.1).

However, it was noted that the schedule optimizer used for each of the scans had some limitations in prioritizing the Southern field—improved logic of the scheduler would have better

prioritization of the Southern field in the three (SW-N-SE) and four (SW-NE-NW-SE) azimuth scans, which would improve their $\sigma(r)$ to be at a similar level.

Moving to an option which also observes the Northern sky retains the benefits of the two azimuth (SW-SE) scan, but also fills observing time more productively when the main field isn't available. The choice is then in how the Northern sky is observed.

The four azimuth (SW-NE-NW-SE) scan slightly outperforms the three azimuth (SW-N-SE) scan in $\sigma(r)$, which has its own additional complications—pointing due North raises issues in removing scan synchronous signal (not included in this study), since the scan modes are in an identical direction to sky rotation.

In light of all of these considerations, the four azimuth (SW-NE-NW-SE) scan was selected.

9.6 Summary and Discussion

The recommended strategy for the SATs was a ‘rising/setting’ scan with an observing elevation of 50° , azimuth throw (on mount) of 40° , and four boresight directions (SW-NE-NW-SE) with improved logic. The field prioritization behaves as follows:

1. Define the Southern patch when the boresight is inside the galactic mask—the Southern yellow box in the original scans, see figure 9.6. Define the Northern patch as the northern yellow box.
2. Observe the Southern patch whenever it is available.
3. If the boresight goes outside the Southern patch area or it is unavailable due to avoidances, check if the Northern patch is available and observe it.
4. If the Northern patch is not available, continue observing the galaxy in SE/SW until it is available.
5. If the SW/SE galaxy, Southern patch, and Northern patch are all unavailable, observe the galaxy in NW/NE.

ISO feedback will be used to further inform the science observation strategy, and the strategy will continue to be optimized with regular calibration and maintenance. This begins with testing for instrumental effects by comparing noise and scan synchronous signal at different scan speeds, elevations, boresight angles, and accelerations. Nominal scan observations were to be run, switching between two scan speeds approximately hourly, with boresight rotation every two days for at least one month.




Figure 9.6: The proposed new baseline strategy for the SATs.

	Baseline 2018 Scan	50° (original)	50° (updated)	60° (original)	60° (updated)
$\sigma(r) \times 10^{-3}$	2.412	2.015	1.911	2.074	1.980
Improvement	N/A	1.20	1.26	1.16	1.22

Table 9.6: The factors of improvement of the proposed scans, compared to the 2018 baseline scan. Note that $\sigma(r)$ for the ‘updated’ scans match those of the equivalent SW-SE scans—this is because the improved logic of the SW-NE-NW-SE scans is expected to bring their constraints in line with the SW-SE scans.

Initial plans were to test mount speeds of $[0.8, 1.5]$, but following ISO testing these were reduced to $[0.5, 0.8]$ deg/s. Likewise, the mount acceleration tests were reduced from $[1, 2, 3]$ to $[0.25]$ deg/s².

Observing elevations of $[50^\circ, 60^\circ]$ were to be tested, but this has been fixed at 60° , due to the hardware on SATp3 which is more susceptible to ground pickup. Boresight angles of $[-45^\circ, 0^\circ, 45^\circ]$ are being evaluated. The settling times for elevation and boresight changes are being measured to inform the cadence of boresight rotation, and potential elevation changes.

The Sun avoidance range remains at 45° , but the Moon avoidance range has been scrapped (which now matches the LAT), so that the Moon can be used as a point source for detector calibration. Other things that will need to be assessed in ISO to further inform scan choices include: polarization; ground pickup (e.g., local peaks, etc.); pixel sampling (vis-a-vis throw vs scan duration, for example); satellites (e.g., Starlink & others).

The baseline scan strategy for near-term observations will be decided once three months of data have been collected.

The improvement to $\sigma(r)$ purely on the basis of scanning strategy can be seen in table 9.6. The proposed scanning strategy has a forecasted improvement over the baseline constraint by a factor of ≈ 1.26 , while the altered strategy which has arisen from ISO testing gives an improvement of ≈ 1.20 . Other improvements to $\sigma(r)$ may come from improved delensing,

compared to the 2018 baseline delensing analysis—see reference [63] for further details.

Improving the constraint on r directly enhances SO’s sensitivity to the primordial gravitational wave background, which imprints a distinct B -mode polarization pattern in the CMB (see section 5.3.3 for details). A lower $\sigma(r)$ reduces the threshold for a statistically significant detection of r , increasing the likelihood of confirming a primordial signal if it exists near current upper limits. This will enable tighter tests of, e.g., inflationary models, particularly those predicting energy scales near the grand unification scale [224]. As such, optimizing survey strategy to minimize $\sigma(r)$ plays an important role in being able to achieve SO’s science goals.

Part IV

Final Remarks

Chapter 10

Conclusions and Future Directions

10.1 Introduction

The aim of this thesis has been to forecast the future sensitivities of experiments to the stochastic gravitational wave background (SGWB)—the background of low-level GWs generated by a broad range of astrophysical and cosmological sources. The SGWB has only recently been observationally confirmed with pulsar timing array (PTA) data; experiments looking at other frequency bands have so far only yielded upper limits. The coming decades will see the deployment of a new generation of gravitational wave (GW) and cosmic microwave background (CMB) experiments, capable of probing the SGWB with greatly improved sensitivity.

Chapters 6 and 7 consider direct GW detection, with a general forecasting framework using Fisher matrix techniques. The framework incorporates evolving noise curves, combinations of experiments, the inclusion of observational priors, and the ability to assess sensitivity to both simple and complex spectral shapes. The results given in Chapter 7 show the sensitivities of upcoming GW experiments (individually, in combination, and with the inclusion of the NANOGrav 15-year results as an observational prior) to a range of fiducial SGWB spectra.

Separately, as a member of the Simons Observatory (SO) collaboration, I contributed directly to the experimental design process, playing a key role in forecasting the scientific implications of various scan strategies for the Large Aperture Telescope (LAT) and Small Aperture Telescopes (SATs), with a particular focus on how different strategies could improve SO’s constraining power with respect to the primordial GWB. The methodology, technical considerations, and results from these two campaigns of work, are covered in Chapters 8 and 9.

Taken together, the work presented in this thesis bridges the operational design of next-generation CMB experiments with the theoretical goals of SGWB cosmology.

10.2 Key Findings and Contributions

This thesis has presented an investigation into future prospects for observing the [SGWB](#) through the forecasting of upcoming [GW](#) and [CMB](#) experiments. The work can be split into two components: forecasting the sensitivities of upcoming [GW](#) detectors to different proposed [SGWB](#) spectra; and scan strategy optimization for [SO](#).

10.2.1 Forecasting Sensitivities of GW Experiments

The forecasting results presented in Chapter 7 show that upcoming [GW](#) experiments will be able to detect or constrain [SGWB](#) signals across a wide range of frequencies, although the observing gap between [PTA](#) and interferometer frequencies will persist. The sensitivities of experiments improve not only with observing time, but by being taken in combination with other experiments. Our overall sensitivities therefore improve as more and better experiments come online. It was shown that the sensitivities of experiments do not noticeably improve if their predecessor remains online, but often improvements of greater than the naive minimum $\sqrt{2}$ can be conferred from an additional (similarly sensitive) experiment in the same frequency range. The inclusion of current observations from [NANOGrav](#) as an observational prior was considered, and shown to not add noticeable improvements to constraining power at their current level of precision.

A novel visualization of the constraining power of experiments to the slope and amplitude of simple power laws (as ‘bowties’) was presented, and shown to be a useful tool in many circumstances—its limitations in over- and under-estimating constraining power below the sensitivity curve were discussed in Chapter 7, but it was shown to be a good fit to the set of error bars across the frequency range above the Ω_{nwt} curve.

A similar visualization was used for the preliminary results shown at the end of Chapter 7, which forecast the ability of a given experiment to constrain the deviation (or lack thereof) of an observed spectrum from a simple power law. This can be taken as an indication of which more complex spectra (i.e., broken power laws) an experiment might be able to make useful constraints on, to inform future work with these spectral shapes.

10.2.2 Scan Strategy Optimization for SO

The second component of this thesis was the optimization of scan strategies for [SO](#)’s [LAT](#) and [SAT](#)s, with a particular focus on optimizing for sensitivity to the [SGWB](#) through the [CMB](#)’s

B-mode polarization.

For both the [LAT](#) and [SAT](#)s, multiple scan variations were considered, with optimization of promising candidates subsequently carried out. Proposed scan strategies were produced, with quantifiable improvements to the 2018 baseline strategies in terms of sky coverage, uniformity, and $\sigma(r)$ etc., to be tested during initial science observations ([ISO](#)).

10.3 Scientific Implications

The results presented in this thesis highlight the increasing likelihood that [SGWB](#) observations will play a pivotal role in probing the physics of the early Universe. Next-generation [GW](#) experiments will continue to increase in sensitivity to the [SGWB](#), and signals from cosmological sources may well be detected over the coming decades.

Next-generation [CMB](#) experiments, and [SO](#) in particular, will reach significantly better sensitivities to the primordial gravitational wave background ([GWB](#)) through their measurements of *B*-mode polarization. The demonstration that high cadence scan strategies can improve on classical tiling strategies not only for transient science, but also for polarization measurements, will help guide scan strategy decisions for future [CMB](#) experiments.

The combination of direct and indirect methods of observation provides complementary constraints on the [SGWB](#) at different frequencies, from the [CMB](#) at very low frequencies, through [PTA](#) frequencies, and to space- and ground-based detector frequencies.

10.4 Future Directions

The forecasting for [GW](#) experiments presented in this thesis was carried out for simple fiducial spectra, which can be reasonable representations of predicted cosmological spectra in some cases, but a wider range of spectra can be accurately modelled with more complicated functions (e.g. double broken power laws ([DBPLs](#))). Work on forecasting sensitivities to these more complicated shapes is ongoing, with preliminary results for the intermediate step of forecasting sensitivities to general deviations from simple power laws presented at the end of Chapter 7. Beyond the forecasting of sensitivities to more complicated spectra (e.g. [DBPLs](#)) to assess a wider range of fiducial models, sensitivities to the physical parameters which govern these spectral shapes for each source could be evaluated.

The results presented in Chapter 7 do include astrophysical foregrounds, but their separation from observed [SGWB](#) is not explicitly considered. The forecasts were also done purely using

projected noise curves, and the reality of observations will bring additional uncertainties—while the forecasts presented here offer insight into future sensitivities, refinements will be necessary as observations are made. When the new combined [PTA](#) data is released from the International Pulsar Timing Array ([IPTA](#)), it can be tested as an observational prior to evaluate whether it brings any additional constraining power to current [GW](#) experiments. Given that the [IPTA](#) results are expected to be significantly more accurate than the individual results from e.g. North American Nanohertz Observatory for Gravitational Waves ([NANOGrav](#)), [IPTA](#) may confer a non-negligible improvement.

I intend to package the forecasting pipeline into a publicly accessible open-source tool, to allow the broader cosmological community to access it for easy forecasting of sensitivities to cosmological models of interest. This could be used by theorists to assess model detectability against projected experimental sensitivities, as well as by experimentalists to assess the implications of new experiments to the existing combined sensitivity of current and/or upcoming experiments.

The forecasted sensitivities of [SO](#) from the scan strategy campaigns has so far been limited to simulations—testing during [ISO](#) and subsequent observing runs will better characterize the implications of different strategies for the science goals. The scan strategy recommendations made to [SO](#), once tested on-sky, will provide guidance to future [CMB](#) experiments as they decide on their own scanning strategies.

10.5 Final Remarks

This thesis has explored the emerging observational landscape of the [SGWB](#), developing theoretical tools and optimizing experimental strategies to investigate how and when signals at different frequencies may be detected. By contributing to both the forecasting of [GW](#) detector sensitivities and the optimization of [CMB](#) scan strategies, this work bridges the gap between theory and experimental design.

In the coming decades, we hope that, much as the [CMB](#) has already proved to be, the [SGWB](#) will present itself as a new cornerstone of observational cosmology, with a unique insight into the earliest stages of the Universe.

Bibliography

- [1] E. Gleave, H. Shirling, and F. Yusaf, “Dark Matter and Dark Energy.” (Lecture notes), 2019.
- [2] E. Gleave, “The Cosmic Microwave Background.” (Literature Review), 2019.
- [3] T. Baker, P. G. Ferreira, and C. Skordis, “The Parametrized Post-Friedmann Framework for Theories of Modified Gravity: Concepts, Formalism and Examples,” *Physical Review D*, vol. 87, no. 2, p. 024015, 2013.
- [4] J.-H. He, L. Guzzo, B. Li, and C. M. Baugh, “No evidence for modifications of gravity from galaxy motions on cosmological scales,” *Nature Astronomy*, vol. 2, pp. 967–972, Sept. 2018.
- [5] S. Boran, S. Desai, E.-O. Kahya, and R.-P. Woodard, “GW170817 falsifies dark matter emulators,” *Physical Review D*, vol. 97, Feb. 2018.
- [6] B. Li, K. Koyama, and G.-B. Zhao, “Modified Gravity: Progresses and Outlook of Theories, Numerical Techniques and Observational Tests,” *Contemporary Physics*, vol. 61, no. 1, pp. 1–59, 2020.
- [7] A. Einstein, “Näherungsweise Integration der Feldgleichungen der Gravitation,” *Sitzungsberichte der Königlich Preussischen Akademie der Wissenschaften (Berlin)*, pp. 688–696, 1916.
- [8] R. A. Hulse and J. H. Taylor, “Discovery of a pulsar in a binary system.,” , vol. 195, pp. L51–L53, Jan. 1975.
- [9] J. H. Taylor and J. M. Weisberg, “A new test of general relativity - Gravitational radiation and the binary pulsar PSR 1913+16,” *Astrophysical Journal*, vol. 253, pp. 908–920, Feb. 1982.

- [10] LIGO Scientific Collaboration, “Advanced LIGO,” *Classical and Quantum Gravity*, vol. 32, p. 074001, Apr. 2015.
- [11] G. M. Harry, “Advanced LIGO: The next generation of gravitational wave detectors,” *Class. Quant. Grav.*, vol. 27, p. 084006, 2010.
- [12] I. Chakrabarty, “Gravitational Waves: An Introduction,” 1999.
- [13] N. Christensen, “Stochastic gravitational wave backgrounds,” *Reports on Progress in Physics*, vol. 82, p. 016903, Nov. 2018.
- [14] N. J. Cornish and A. Sesana, “Pulsar timing array analysis for black hole backgrounds,” *Classical and Quantum Gravity*, vol. 30, p. 224005, Nov. 2013.
- [15] C. L. Fryer and K. C. B. New, “Gravitational Waves from Gravitational Collapse,” *Living Reviews in Relativity*, vol. 14, Jan. 2011.
- [16] V. Skliris, M. R. K. Norman, and P. J. Sutton, “Real-Time Detection of Unmodelled Gravitational-Wave Transients Using Convolutional Neural Networks,” *arXiv e-prints*, 2024.
- [17] V. M. Slipher, “Nebulae,” *Proceedings of the American Philosophical Society*, vol. 56, pp. 403–409, Jan. 1917.
- [18] E. Hubble, “A Relation between Distance and Radial Velocity among Extra-Galactic Nebulae,” *Proceedings of the National Academy of Science*, vol. 15, pp. 168–173, Mar. 1929.
- [19] M. Pascale, B. Frye, and J. P. et al., “SN H0pe: The First Measurement of H_0 from a Multiply Imaged Type Ia Supernova, Discovered by JWST,” *The Astrophysical Journal*, 2025.
- [20] A. G. Riess, G. S. Anand, W. Yuan, S. Casertano, and A. Dolphin, “Crowded No More: The Accuracy of the Hubble Constant Tested with High Resolution Observations of Cepheids by JWST,” *Astrophysical Journal Letters*, vol. 956, p. L18, 2023.
- [21] A. G. Riess, W. Yuan, L. M. Macri, D. Scolnic, D. Brout, *et al.*, “A Comprehensive Measurement of the Local Value of the Hubble Constant with 1 km/s/Mpc Uncertainty from the Hubble Space Telescope and the SH0ES Team,” *Astrophysical Journal Letters*, vol. 934, p. L7, 2022.

- [22] N. Aghanim, Y. Akrami, M. Ashdown, J. Aumont, C. Baccigalupi, *et al.*, “Planck 2018 results: VI. Cosmological parameters,” *Astronomy & Astrophysics*, vol. 641, p. A6, Sept. 2020.
- [23] LIGO Scientific Collaboration & Virgo Collaboration, “Standard-siren Cosmology Using Gravitational Waves from Binary Black Holes,” in *Theory Testing in Gravitational-Wave Astrophysics*, Springer, 2021.
- [24] A. Palmese, C. R. Bom, S. Mucesh, and W. G. Hartley, “A Standard Siren Measurement of the Hubble Constant Using Gravitational-wave Events from the First Three LIGO/Virgo Observing Runs and the DESI Legacy Survey,” *Astrophysical Journal*, vol. 943, p. 56, Jan. 2023.
- [25] A. Palmese, R. Kaur, A. Hajela, R. Margutti, A. McDowell, and A. MacFadyen, “Standard siren measurement of the Hubble constant using GW170817 and the latest observations of the electromagnetic counterpart afterglow,” *Physical Review D*, vol. 109, p. 063508, Mar. 2024.
- [26] J. Binney and S. Tremaine, *Galactic Dynamics*. Princeton University Press, 2nd ed., 2008.
- [27] A. J. Cuesta, F. Prada, A. Klypin, and M. Moles, “The virialized mass of dark matter haloes,” *Monthly Notices of the Royal Astronomical Society*, vol. 389, pp. 385–397, Sept. 2008.
- [28] J. F. Navarro, C. S. Frenk, and S. D. M. White, “The Structure of Cold Dark Matter Halos,” *Astrophysical Journal*, vol. 462, p. 563, May 1996.
- [29] K. Jedamzik and M. Pospelov, “Big Bang Nucleosynthesis and Particle Dark Matter,” *Physical Review D*, vol. 79, p. 123529, 2009.
- [30] G. Bertone, D. Hooper, and J. Silk, “Particle Dark Matter: Evidence, Candidates and Constraints,” *Physics Reports.*, vol. 405, pp. 279–390, 2005.
- [31] A. Arbey and F. Mahmoudi, “Dark matter and the early universe: A review,” *Progress in Particle and Nuclear Physics*, vol. 119, p. 103865, July 2021.
- [32] Y. Wang, X.-D. Li, and X. Zhang, “Cosmological constraints on variable warm dark matter,” *Physics Letters B*, vol. 720, pp. 271–275, 2013.

- [33] F. Hervás Peters, A. Schneider, J. Bucko, S. K. Giri, and G. Parimbelli, “Constraining hot dark matter sub-species with weak lensing and the cosmic microwave background radiation,” *Astronomy & Astrophysics*, vol. 687, p. A161, July 2024.
- [34] M. Chevallier and D. Polarski, “Accelerating Universes with Scaling Dark Matter,” *International Journal of Modern Physics D*, vol. 10, pp. 213–224, 2001.
- [35] E. V. Linder, “Exploring the Expansion History of the Universe,” *Physical Review Letters*, vol. 90, p. 091301, 2003.
- [36] R. J. Scherrer, “Mapping the Chevallier–Polarski–Linder parametrization onto physical dark energy models,” *Physical Review D*, vol. 92, no. 12, p. 124018, 2015.
- [37] DESI Collaboration, “DESI 2024 VI: Cosmological Constraints from the Measurements of Baryon Acoustic Oscillations,” *arXiv e-prints*, 2024.
- [38] C.-G. Park, J. de Cruz Pérez, and B. Ratra, “Using non-DESI data to confirm and strengthen the DESI 2024 spatially-flat w_0w_a CDM cosmological parameterization result,” *Physical Review D*, vol. 110, p. 123533, Dec. 2024.
- [39] J. Silk, *On the Shores of the Unknown: A Short History of the Universe*. Cambridge University Press, 2005.
- [40] E. W. Kolb and M. S. Turner, “Grand unified theories and the origin of the baryon asymmetry,” *Annual Review of Nuclear and Particle Science*, vol. 33, pp. 645–696, 1983.
- [41] D. Bödeker and W. Buchmüller, “Baryogenesis from the weak scale to the grand unification scale,” *Physics Reports*, vol. 913, pp. 1–113, 2020.
- [42] A. H. Guth, “Inflationary universe: A possible solution to the horizon and flatness problems,” *Physical Review D*, vol. 23, no. 2, pp. 347–356, 1981.
- [43] V. Salzano and et al., “Relativistic Cosmology with an Introduction to Inflation,” *Universe*, vol. 7, no. 8, p. 276, 2020.
- [44] B. D. Fields, “The Primordial Lithium Problem,” *Annual Review of Nuclear and Particle Science*, vol. 61, pp. 47–68, Nov. 2011.
- [45] L. Koopmans, J. Pritchard, G. Mellema, J. Aguirre, K. Ahn, *et al.*, “The Cosmic Dawn and Epoch of Reionization with SKA: Focus on 21 cm tomography,” *Proceedings of Science*, 2015.

- [46] D. R. DeBoer, A. R. Parsons, *et al.*, “Hydrogen Epoch of Reionization Array (HERA),” *Publications of the Astronomical Society of the Pacific*, 2017.
- [47] E. Abdalla, G. F. Abellán, A. Aboubrahim, A. Agnello, Ö. Akarsu, *et al.*, “Cosmology intertwined: A review of the particle physics, astrophysics, and cosmology associated with the cosmological tensions and anomalies,” *Journal of High Energy Astrophysics*, vol. 34, pp. 49–211, June 2022.
- [48] C. Heymans, T. Tröster, M. Asgari, C. Blake, H. Hildebrandt, *et al.*, “KiDS-1000 Cosmology: Multi-probe weak gravitational lensing and spectroscopic galaxy clustering constraints,” , vol. 646, p. A140, Feb. 2021.
- [49] G. Agazie, A. Anumalapudi, A. M. Archibald, Z. Arzoumanian, P. T. Baker, *et al.*, “The NANOGrav 15 yr Data Set: Evidence for a Gravitational-wave Background,” *The Astrophysical Journal Letters*, vol. 951, p. L8, June 2023.
- [50] G. Agazie, M. F. Alam, A. Anumalapudi, A. M. Archibald, Z. Arzoumanian, *et al.*, “The NANOGrav 15 yr Data Set: Observations and Timing of 68 Millisecond Pulsars,” *The Astrophysical Journal Letters*, vol. 951, p. L9, June 2023.
- [51] J. Antoniadis, P. Arumugam, S. Arumugam, S. Babak, M. Bagchi, *et al.*, “The second data release from the European Pulsar Timing Array: III. Search for gravitational wave signals,” *Astronomy & Astrophysics*, vol. 678, p. A50, Oct. 2023.
- [52] D. J. Reardon, A. Zic, R. M. Shannon, G. B. Hobbs, M. Bailes, *et al.*, “Search for an Isotropic Gravitational-wave Background with the Parkes Pulsar Timing Array,” *The Astrophysical Journal Letters*, vol. 951, p. L6, June 2023.
- [53] H. Xu, S. Chen, Y. Guo, J. Jiang, B. Wang, *et al.*, “Searching for the Nano-Hertz Stochastic Gravitational Wave Background with the Chinese Pulsar Timing Array Data Release I,” *Research in Astronomy & Astrophysics*, vol. 23, p. 075024, June 2023.
- [54] A. A. Penzias and R. W. Wilson, “A Measurement of Excess Antenna Temperature at 4080 Mc/s.,” *Astrophysical Journal*, vol. 142, pp. 419–421, July 1965.
- [55] A. P. S. Yadav and B. D. Wandelt, “Primordial Non-Gaussianity in the Cosmic Microwave Background,” *Advances in Astronomy*, vol. 2010, p. 565248, 2010.

- [56] R. K. Sachs and A. M. Wolfe, “Perturbations of a Cosmological Model and Angular Variations of the Microwave Background,” *Astrophysical Journal*, vol. 147, p. 73, Jan. 1967.
- [57] J. Silk, “Cosmic Black-Body Radiation and Galaxy Formation,” *Astrophysical Journal*, vol. 151, p. 459, Feb. 1968.
- [58] M. J. Rees and D. W. Sciama, “Large-scale Density Inhomogeneities in the Universe,” *Nature*, vol. 217, pp. 511–516, Feb. 1968.
- [59] BICEP/Keck Collaboration, “Improved Constraints on Primordial Gravitational Waves using Planck, WMAP, and BICEP/Keck Observations through the 2018 Observing Season,” *Physical Review Letters*, vol. 127, Oct. 2021.
- [60] Planck Collaboration, “Planck 2018 results: I. Overview and the cosmological legacy of Planck,” *Astronomy & Astrophysics*, vol. 641, p. A1, Sept. 2020.
- [61] W. Hu and S. Dodelson, “Cosmic microwave background anisotropies,” *Annual Review of Astronomy & Astrophysics*, vol. 40, p. 171–216, Sept. 2002.
- [62] M. Zaldarriaga and U. Seljak, “All-sky analysis of polarization in the microwave background,” *Physical Review D*, vol. 55, p. 1830–1840, Feb. 1997.
- [63] E. Hertig, K. Wolz, T. Namikawa, A. B. Lizancos, S. Azzoni, *et al.*, “The Simons Observatory: Combining cross-spectral foreground cleaning with multitracer B-mode delensing for improved constraints on inflation,” *Physical Review D*, vol. 110, Aug. 2024.
- [64] A. Baleato Lizancos, A. Challinor, and J. Carron, “Impact of internal-delensing biases on searches for primordial B-modes of CMB polarisation,” *Journal of Cosmology and Astroparticle Physics*, vol. 2021, p. 016, Mar. 2021.
- [65] R. A. Alpher, *Origin and Relative Abundance of the Chemical Elements*. Doctoral dissertation, The George Washington University, Washington, D.C., May 1948.
- [66] R. A. Alpher, H. Bethe, and G. Gamow, “The Origin of Chemical Elements,” *Physical Review*, vol. 73, pp. 803–804, Apr. 1948.
- [67] J. Mather, E. Cheng, D. Cottingham, R. Eplee, D. Fixsen, T. Hewagama, R. Isaacman, K. Jensen, S. Meyer, and P. Noerdlinger, “Measurement of the cosmic microwave background spectrum by the COBE FIRAS instrument,” *The Astrophysical Journal*, vol. 420, pp. 439–444, 01 1994.

- [68] G. F. Smoot, C. L. Bennett, A. Kogut, E. L. Wright, J. Aymon, *et al.*, “Structure in the COBE Differential Microwave Radiometer First-Year Maps,” *Astrophysical Journal, Letters*, vol. 396, p. L1, Sept. 1992.
- [69] C. L. Bennett, M. Halpern, G. Hinshaw, N. Jarosik, A. Kogut, *et al.*, “First-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Preliminary Maps and Basic Results,” *Astrophysical Journal, Supplement*, vol. 148, pp. 1–27, Sept. 2003.
- [70] J. M. Kovac, E. M. Leitch, C. Pryke, J. E. Carlstrom, N. W. Halverson, and W. L. Holzapfel, “Detection of Polarization in the Cosmic Microwave Background using DASI,” *Nature*, vol. 420, p. 772–787, Dec. 2002.
- [71] BICEP/Keck Collaboration, “BICEP/Keck XV: The BICEP3 Cosmic Microwave Background Polarimeter and the First Three-year Data Set,” *Astrophysical Journal*, vol. 927, p. 77, Mar. 2022.
- [72] R. Keisler, C. L. Reichardt, K. A. Aird, B. A. Benson, L. E. Bleem, *et al.*, “A Measurement of the Damping Tail of the Cosmic Microwave Background Power Spectrum with the South Pole Telescope,” *Astrophysical Journal*, vol. 743, p. 28, Dec. 2011.
- [73] S. Padin, J. K. Cartwright, B. S. Mason, T. J. Pearson, A. C. S. Readhead, M. C. Shepherd, J. Sievers, P. S. Udomprasert, W. L. Holzapfel, S. T. Myers, J. E. Carlstrom, E. M. Leitch, M. Joy, L. Bronfman, and J. May, “First Intrinsic Anisotropy Observations with the Cosmic Background Imager,” *The Astrophysical Journal*, vol. 549, p. L1–L5, Mar. 2001.
- [74] CLASS Collaboration, “Two-year Cosmology Large Angular Scale Surveyor (CLASS) Observations: 40 GHz Telescope Pointing, Beam Profile, Window Function, and Polarization Performance,” *Astrophysical Journal*, vol. 891, p. 134, Mar. 2020.
- [75] Polarbear Collaboration, “A Measurement of the Cosmic Microwave Background B-mode Polarization Power Spectrum at Sub-degree Scales with POLARBEAR,” *Astrophysical Journal*, vol. 794, p. 171, Oct. 2014.
- [76] S. Aiola, E. Calabrese, L. Maurin, S. Naess, B. L. Schmitt, *et al.*, “The Atacama Cosmology Telescope: DR4 maps and cosmological parameters,” *Journal of Cosmology and Astroparticle Physics*, vol. 2020, p. 047, Dec. 2020.

- [77] Spider Collaboration, “A Constraint on Primordial B-modes from the First Flight of the SPIDER Balloon-borne Telescope,” *Astrophysical Journal*, vol. 927, p. 174, Mar. 2022.
- [78] C. B. Netterfield, P. A. R. Ade, J. J. Bock, J. R. Bond, J. Borrill, *et al.*, “A Measurement by BOOMERANG of Multiple Peaks in the Angular Power Spectrum of the Cosmic Microwave Background,” *Astrophysical Journal*, vol. 571, pp. 604–614, June 2002.
- [79] P. Ade, J. Aguirre, Z. Ahmed, S. Aiola, A. Ali, *et al.*, “The Simons Observatory: science goals and forecasts,” *Journal of Cosmology and Astroparticle Physics*, vol. 2019, pp. 056–056, Feb. 2019.
- [80] K. N. Abazajian, P. Adshead, Z. Ahmed, S. W. Allen, D. Alonso, *et al.*, “CMB-S4 Science Book, First Edition,” *arXiv e-prints*, Oct. 2016.
- [81] LiteBIRD Collaboration, “Probing cosmic inflation with the LiteBIRD cosmic microwave background polarization survey,” *Progress of Theoretical and Experimental Physics*, vol. 2023, p. 042F01, Apr. 2023.
- [82] LiteBIRD Collaboration, “LiteBIRD: A Satellite for the Studies of B-Mode Polarization and Inflation from Cosmic Background Radiation Detection,” *Journal of Low Temperature Physics*, vol. 194, pp. 443–452, Mar. 2019.
- [83] Simons Observatory, 2025. <https://simonsobservatory.org/>.
- [84] Simons Foundation, 2025. <https://www.simonsfoundation.org/flatiron/center-for-computational-astrophysics/cosmology-x-data-science/simons-observatory/>.
- [85] S. W. Henderson, R. Allison, J. Austermann, T. Baildon, N. Battaglia, *et al.*, “Advanced ACTPol Cryogenic Detector Arrays and Readout,” *Journal of Low Temperature Physics*, vol. 184, p. 772779, Mar. 2016.
- [86] C. M. Posada, P. A. R. Ade, A. J. Anderson, J. Avva, Z. Ahmed, *et al.*, “Large arrays of dual-polarized multichroic TES detectors for CMB measurements with the SPT-3G receiver,” in *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy VIII* (W. S. Holland and J. Zmuidzinas, eds.), vol. 9914, pp. 290 – 300, International Society for Optics and Photonics, The Society of Photo-Optical Instrumentation Engineers, 2016.
- [87] J. A. Sobrin, P. Ade, Z. Ahmed, A. Anderson, J. Avva, *et al.*, “Design and characterization of the SPT-3G receiver,” in *Millimeter, Submillimeter, and Far-Infrared Detectors and*

- Instrumentation for Astronomy IX* (J. Zmuidzinas and J.-R. Gao, eds.), p. 51, The Society of Photo-Optical Instrumentation Engineers, Aug. 2018.
- [88] C. Dragone, “Offset multireflector antennas with perfect pattern symmetry and polarization discrimination,” *The Bell System Technical Journal*, vol. 57, no. 7, pp. 2663–2684, 1978.
 - [89] S. C. Parshley, M. D. Niemack, R. Hills, S. R. Dicker, R. Dunner, *et al.*, “The optical design of the six-meter CCAT-prime and Simons Observatory telescopes,” in *Ground-based and Airborne Telescopes VII*, p. 145, The Society of Photo-Optical Instrumentation Engineers, July 2018.
 - [90] N. Goeckner-Wald, R. Keskitalo, Y. Chinone, P. A. Gallardo, A. Kusaka, *et al.*, “Designs for next generation CMB survey strategies from Chile,” in *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy IX*, p. 136, The Society of Photo-Optical Instrumentation Engineers, July 2018.
 - [91] SO Communications Team, “Project Update—March 2025,” fact sheet, The Simons Observatory, Mar. 2025. https://simonsobservatory.org/wp-content/uploads/2025/03/SO_Fact_Sheet_Mar_2025.PNG.png.
 - [92] A. Achúcarro, M. Biagetti, M. Braglia, G. Cabass, R. Caldwell, *et al.*, “Inflation: Theory and Observations,” *arXiv e-prints*, 2022.
 - [93] M. Tristram, A. J. Banday, K. M. Górski, R. Keskitalo, C. R. Lawrence, K. J. Andersen, R. B. Barreiro, J. Borrill, H. K. Eriksen, R. Fernandez-Cobos, *et al.*, “Planck constraints on the tensor-to-scalar ratio,” *Astronomy & Astrophysics*, vol. 647, 2021.
 - [94] S. Shaikh, S. Mukherjee, A. Rotti, and T. Souradeep, “Constraining stochastic gravitational wave background from weak lensing of CMB B-modes,” *Journal of Cosmology and Astroparticle Physics*, vol. 2016, p. 029–029, Sept. 2016.
 - [95] Planck Collaboration, “Planck 2018 results. IX. Constraints on primordial non-Gaussianity,” *Astronomy & Astrophysics*, vol. 641, p. A9, Sept. 2020.
 - [96] V. F. Shvartsman, “Density of relict particles with zero rest mass in the universe,” *Pisma Zh. Eksp. Teor. Fiz.*, vol. 9, pp. 315–317, 1969.
 - [97] G. Steigman, D. N. Schramm, and J. E. Gunn, “Cosmological Limits to the Number of Massive Leptons,” *Phys. Lett. B*, vol. 66, pp. 202–204, 1977.

- [98] G. Mangano, G. Miele, S. Pastor, and M. Peloso, “A precision calculation of the effective number of cosmological neutrinos,” *Physics Letters B*, vol. 534, p. 8–16, May 2002.
- [99] S. Bashinsky and U. Seljak, “Signatures of relativistic neutrinos in CMB anisotropy and matter clustering,” *Physical Review D*, vol. 69, Apr. 2004.
- [100] KATRIN Collaboration, “Direct neutrino-mass measurement with sub-electronvolt sensitivity,” *Nature Physics*, vol. 18, no. 2, pp. 160–166, 2022.
- [101] N. Aghanim, Y. Akrami, M. Ashdown, J. Aumont, C. Baccigalupi, *et al.*, “Planck 2018 results: VIII. Gravitational lensing,” *Astronomy & Astrophysics*, vol. 641, p. A8, Sept. 2020.
- [102] DESI Collaboration, “DESI 2024 VII: Cosmological Constraints from the Full-Shape Modeling of Clustering Measurements,” *arXiv e-prints*, 2024.
- [103] Planck Collaboration, “Planck intermediate results: XLVII Planck constraints on reionization history,” *Astronomy & Astrophysics*, vol. 596, p. A108, Dec. 2016.
- [104] The Simons Observatory Collaboration, “The Simons Observatory: Science Goals and Forecasts for the Enhanced Large Aperture Telescope,” *arXiv e-prints*, 2025.
- [105] E. Gleave and A. Jaffe, “How to Constrain the Stochastic Gravitational Wave Background with Multi-Frequency Detections,” *arXiv e-prints*, 2025.
- [106] G. Agazie, A. Anumalapudi, A. M. Archibald, P. T. Baker, B. Bécsy, *et al.*, “The NANOGrav 15 yr Data Set: Constraints on Supermassive Black Hole Binaries from the Gravitational-wave Background,” *The Astrophysical Journal Letters*, vol. 952, p. L37, Aug. 2023.
- [107] M. C. Guzzetti, N. Bartolo, M. Liguori, and S. Matarrese, “Gravitational Waves from Inflation,” *La Rivista del Nuovo Cimento*, vol. 39, p. 399–495, Aug. 2016.
- [108] G. Domenech, “Scalar Induced Gravitational Waves Review,” *Universe*, vol. 7, no. 11, 2021.
- [109] C. Yuan and Q.-G. Huang, “A topic review on probing primordial black hole dark matter with scalar induced gravitational waves,” *arXiv e-prints*, 2021.

- [110] M. Hindmarsh, S. J. Huber, K. Rummukainen, and D. J. Weir, “Gravitational Waves from the Sound of a First Order Phase Transition,” *Physical Review Letters*, vol. 112, Jan. 2014.
- [111] C. Caprini, M. Hindmarsh, S. Huber, T. Konstandin, J. Kozaczuk, G. Nardini, J. M. No, A. Petiteau, P. Schwaller, G. Servant, and D. J. Weir, “Science with the space-based interferometer eLISA. II: gravitational waves from cosmological phase transitions,” *Journal of Cosmology and Astroparticle Physics*, vol. 2016, p. 001, Apr. 2016.
- [112] C. Caprini, M. Chala, G. C. Dorsch, M. Hindmarsh, S. J. Huber, T. Konstandin, J. Kozaczuk, G. Nardini, J. M. No, K. Rummukainen, P. Schwaller, G. Servant, A. Tranberg, and D. J. Weir, “Detecting gravitational waves from cosmological phase transitions with LISA: an update,” *Journal of Cosmology and Astroparticle Physics*, vol. 2020, pp. 024–024, Mar. 2020.
- [113] T. W. B. Kibble, “Topology of cosmic domains and strings,” *Journal of Physics A: Mathematical and General*, vol. 9, pp. 1387–1398, Aug. 1976.
- [114] B. Li, P. R. Shapiro, and T. Rindler-Daller, “Bose-Einstein-condensed scalar field dark matter and the gravitational wave background from inflation: New cosmological constraints and its detectability by LIGO,” *Physical Review D*, vol. 96, Sept. 2017.
- [115] G. Lambiase, L. Mastrototaro, and L. Visinelli, “Gravitational waves and neutrino oscillations in Chern-Simons axion gravity,” *Journal of Cosmology and Astroparticle Physics*, vol. 2023, p. 011, Jan. 2023.
- [116] D. Chowdhury, A. Hait, S. Mohanty, and S. Prakash, “Ultra-light dark matter explanation of NANOGrav observations,” *arXiv e-prints*, 2024.
- [117] M. Maggiore, “Gravitational wave experiments and early universe cosmology,” *Physics Reports*, vol. 331, p. 283–367, July 2000.
- [118] C. J. Moore, R. H. Cole, and C. P. L. Berry, “Gravitational-wave sensitivity curves,” *Classical and Quantum Gravity*, vol. 32, p. 015014, Dec. 2014.
- [119] E. Thrane and J. D. Romano, “Sensitivity curves for searches for gravitational-wave backgrounds,” *Physical Review D*, vol. 88, Dec. 2013.
- [120] M. Bonetti and A. Sesana, “Gravitational wave background from extreme mass ratio inspirals,” *Physical Review D*, vol. 102, Nov. 2020.

- [121] L. Z. Kelley, L. Blecha, L. Hernquist, A. Sesana, and S. R. Taylor, “The gravitational wave background from massive black hole binaries in Illustris: spectral features and time to detection with pulsar timing arrays,” *Monthly Notices of the Royal Astronomical Society*, vol. 471, pp. 4508–4526, Jun. 2017.
- [122] M. Enoki and M. Nagashima, “The Effect of Orbital Eccentricity on Gravitational Wave Background Radiation from Supermassive Black Hole Binaries,” *Progress of Theoretical Physics*, vol. 117, pp. 241–256, Feb. 2007.
- [123] S. Vagnozzi, “Implications of the NANOGrav results for inflation,” *Monthly Notices of the Royal Astronomical Society: Letters*, vol. 502, pp. L11–L15, Dec. 2020.
- [124] M. J. Rees, “Black Hole Models for Active Galactic Nuclei,” *Annual Review of Astronomy & Astrophysics*, vol. 22, pp. 471–506, Jan. 1984.
- [125] S. W. Hawking, “Particle creation by black holes,” *Communications in Mathematical Physics*, vol. 43, pp. 199–220, Aug. 1975.
- [126] J. C. Niemeyer and K. Jedamzik, “Dynamics of primordial black hole formation,” *Physical Review D*, vol. 59, May 1999.
- [127] Vaskonen, Ville and Veermäe, Hardi, “Did NANOGrav See a Signal from Primordial Black Hole Formation?,” *Physical Review Letters*, vol. 126, Feb. 2021.
- [128] M. Sasaki, T. Suyama, T. Tanaka, and S. Yokoyama, “Primordial black holes—perspectives in gravitational wave astronomy,” *Classical and Quantum Gravity*, vol. 35, p. 063001, Feb. 2018.
- [129] V. Atal, A. Sanglas, and N. Triantafyllou, “NANOGrav signal as mergers of Stupendously Large Primordial Black Holes,” *arXiv e-prints*, 2020.
- [130] S. Kuroyanagi, T. Takahashi, and S. Yokoyama, “Blue-tilted inflationary tensor spectrum and reheating in the light of NANOGrav results,” *Journal of Cosmology and Astroparticle Physics*, vol. 2021, p. 071, Jan. 2021.
- [131] M. Sakellariadou, “Cosmic Strings and Cosmic Superstrings,” *Nuclear Physics B - Proceedings Supplements*, vol. 192-193, pp. 68–90, 2009.
- [132] J. Blanco-Pillado, K. Olum, and X. Siemens, “New limits on cosmic strings from gravitational wave observation,” *Physics Letters B*, vol. 778, Sep. 2017.

- [133] J. Aasi, J. Abadie, B. P. Abbott, R. Abbott, T. Abbott, *et al.*, “Constraints on Cosmic Strings from the LIGO-Virgo Gravitational-Wave Detectors,” *Physical Review Letters*, vol. 112, Apr. 2014.
- [134] P. A. R. Ade, N. Aghanim, C. Armitage-Caplan, M. Arnaud, M. Ashdown, *et al.*, “Planck 2013 results. XXV. Searches for cosmic strings and other topological defects,” *Astronomy & Astrophysics*, vol. 571, p. A25, Oct. 2014.
- [135] S. Henrot-Versillé, F. Robinet, N. Leroy, S. Plaszczyński, N. Arnaud, *et al.*, “Improved constraint on the primordial gravitational-wave density using recent cosmological data and its impact on cosmic string models,” *Classical and Quantum Gravity*, vol. 32, no. 4, p. 045003, 2015.
- [136] B. P. Abbott, R. Abbott, T. D. Abbott, F. Acernese, K. Ackley, *et al.*, “Constraints on cosmic strings using data from the first Advanced LIGO observing run,” *Physical Review D*, vol. 97, May 2018.
- [137] F. P. Huang and X. Zhang, “Probing the gauge symmetry breaking of the early universe in 3-3-1 models and beyond by gravitational waves,” *Physics Letters B*, vol. 788, pp. 288–294, Jan. 2019.
- [138] A. Mazumdar and G. White, “Cosmic phase transitions: their applications and experimental signatures,” *arXiv e-prints*, 2019.
- [139] A. D. Linde, “Decay of the False Vacuum at Finite Temperature,” *Nuclear Physics B*, vol. 216, pp. 421–445, May 1983.
- [140] E. Gleave, “Gravitational Waves and Cosmological Phase Transitions.” (Masters Thesis, KCL), 2020.
- [141] A. Kosowsky and M. S. Turner, “Gravitational radiation from colliding vacuum bubbles: Envelope approximation to many-bubble collisions,” *Physical Review D*, vol. 47, May 1993.
- [142] H. Child and J. Giblin, “Gravitational radiation from first-order phase transitions,” *Journal of Cosmology and Astroparticle Physics*, vol. 2012, 2012.
- [143] C. Caprini, R. Durrer, and G. Servant, “Gravitational wave generation from bubble collisions in first-order phase transitions: An analytic approach,” *Physical Review D*, vol. 77, 2008.

- [144] D. Cutting, M. Hindmarsh, and D. J. Weir, “Vorticity, kinetic energy, and suppressed gravitational wave production in strong first order phase transitions,” *arXiv e-prints*, 2019.
- [145] C. Caprini, R. Durrer, and G. Servant, “The stochastic gravitational wave background from turbulence and magnetic fields generated by a first-order phase transition,” *Journal of Cosmology and Astroparticle Physics*, vol. 2009, Dec. 2009.
- [146] T. Konstandin, “Gravitational radiation from a bulk flow model,” *Journal of Cosmology and Astroparticle Physics*, vol. 2018, pp. 047–047, Mar. 2018.
- [147] T. Damour and A. Vilenkin, “Gravitational radiation from cosmic (super)strings: Bursts, stochastic background, and observational windows,” *Physical Review D*, vol. 71, Mar. 2005.
- [148] E. J. Copeland and T. W. B. Kibble, “Cosmic strings and superstrings,” *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 466, pp. 623–657, Jan. 2010.
- [149] S. Kuroyanagi, K. Miyamoto, T. Sekiguchi, K. Takahashi, and J. Silk, “Forecast constraints on cosmic string parameters from gravitational wave direct detection experiments,” *Physical Review D*, vol. 86, p. 023503, Jul. 2012.
- [150] Y. Matsui, K. Horiguchi, D. Nitta, and S. Kuroyanagi, “Gravitational wave spectrum from kinks on infinite cosmic strings with Y-junctions,” *Journal of Cosmology and Astroparticle Physics*, vol. 2020, p. 039–039, Nov. 2020.
- [151] A. Avgoustidis and E. P. S. Shellard, “Effect of reconnection probability on cosmic (super)string network density,” *Physical Review D*, vol. 73, p. 041301, Feb. 2006.
- [152] X. Siemens, V. Mandic, and J. Creighton, “Gravitational-Wave Stochastic Background from Cosmic Strings,” *Physical Review Letters*, vol. 98, Mar. 2007.
- [153] R. Peccei and H. Quinn, “CP Conservation in the Presence of Pseudoparticles,” *Physical Review Letters*, vol. 38, pp. 1440–1443, Jun. 1977.
- [154] S. Weinberg, “A New Light Boson?,” *Physical Review Letters*, vol. 40, pp. 223–226, Jan. 1978.

- [155] L. D. Duffy and K. v. Bibber, “Axions as dark matter particles,” *New Journal of Physics*, vol. 11, p. 105008, Oct. 2009.
- [156] A. Ringwald, “Axions and Axion-Like Particles,” *arXiv e-prints*, 2014.
- [157] B. Salehian, M. A. Gorji, S. Mukohyama, and H. Firouzjahi, “Analytic study of dark photon and gravitational wave production from axion,” *Journal of High Energy Physics*, vol. 2021, May 2021.
- [158] J. Alvey, U. Bhardwaj, V. Domcke, M. Pieroni, and C. Weniger, “Simulation-based inference for stochastic gravitational wave background data analysis,” *arXiv e-prints*, 2024.
- [159] D. Poletti, “Measuring the primordial gravitational wave background in the presence of other stochastic signals,” *Journal of Cosmology and Astroparticle Physics*, vol. 2021, p. 052, May 2021.
- [160] V. Mandic, E. Thrane, S. Giampanis, and T. Regimbau, “Parameter Estimation in Searches for the Stochastic Gravitational-Wave Background,” *Physical Review Letters*, vol. 109, Oct. 2012.
- [161] M. Kachelriess and M. P. Nødtvedt, “Lunar response to gravitational waves,” *arXiv e-prints*, 2024.
- [162] D. Blas and A. C. Jenkins, “Bridging the μHz Gap in the Gravitational-Wave Landscape with Binary Resonances,” *Physical Review Letters*, vol. 128, Mar. 2022.
- [163] W. Martens, M. Khan, and J.-B. Bayle, “LISAmix: Improving the Low-Frequency Gravitational-Wave Sensitivity by Two Orders of Magnitude,” *arXiv e-prints*, 2023.
- [164] P. J. Carter, T. R. Marsh, D. Steeghs, P. J. Groot, G. Nelemans, D. Levitan, A. Rau, C. M. Copperwheat, T. Kupfer, and G. H. A. Roelofs, “A search for the hidden population of AM CVn binaries in the Sloan Digital Sky Survey,” *Monthly Notices of the Royal Astronomical Society*, vol. 429, p. 2143–2160, Dec. 2012.
- [165] R. Abbott, T. D. Abbott, S. Abraham, F. Acernese, K. Ackley, *et al.*, “Population Properties of Compact Objects from the Second LIGO–Virgo Gravitational-Wave Transient Catalog,” *The Astrophysical Journal Letters*, vol. 913, p. L7, May 2021.
- [166] The LIGO Scientific Collaboration, the Virgo Collaboration, and the KAGRA Collaboration, “Observation of Gravitational Waves from Two Neutron Star–Black Hole Coalescences,” *The Astrophysical Journal Letters*, vol. 915, p. L5, Jun. 2021.

- [167] A. A. Michelson and E. W. Morley, “On the Relative Motion of the Earth and the Luminiferous Ether,” *American Journal of Science*, vol. 34, no. 203, pp. 333–345, 1887.
- [168] T. M. Tauris, V. M. Kaspi, R. P. Breton, A. T. Deller, E. F. Keane, *et al.*, “Understanding the Neutron Star Population with the SKA,” *arXiv e-prints*, 2014.
- [169] T. L.-V.-K. O. S. W. Groups, “The LSC-Virgo-KAGRA Observational Science White Paper (2024 Edition),” Tech. Rep. LIGO-T2300406-v1,VIR-1043A-23,JGW-T2315390-v0, LIGO Scientific Collaboration, Virgo Collaboration, KAGRA Collaboration, 2024.
- [170] M. Evans, R. X. Adhikari, C. Afle, S. W. Ballmer, S. Biscoveanu, *et al.*, “A Horizon Study for Cosmic Explorer: Science, Observatories, and Community,” *arXiv e-prints*, 2021.
- [171] M. Maggiore, C. V. D. Broeck, N. Bartolo, E. Belgacem, D. Bertacca, *et al.*, “Science case for the Einstein telescope,” *Journal of Cosmology and Astroparticle Physics*, vol. 2020, p. 050–050, Mar. 2020.
- [172] ET steering committee *et al.*, “Einstein Telescope: Science Case, Design Study and Feasibility Report,” Tech. Rep. ET-0028A-20, Einstein Telescope, 2020.
- [173] LISA Science Study Team, “LISA Science Requirements Document,” Tech. Rep. ESA-L3-EST-SCI-RS-001, LISA, 2018.
- [174] S. Kawamura, M. Ando, N. Seto, S. Sato, M. Musha, *et al.*, “Current status of space gravitational wave antenna DECIGO and B-DECIGO,” *arXiv e-prints*, 2020.
- [175] G. Harry, “Talk for the 2009 GWADW about the Big Bang Observer,” Tech. Rep. LIGO-G0900426, Big Bang Observer, 2009.
- [176] K. Yagi and N. Seto, “Detector configuration of DECIGO/BBO and identification of cosmological neutron-star binaries,” *Physical Review D*, vol. 83, Feb. 2011.
- [177] A. Brazier, S. Chatterjee, T. Cohen, J. M. Cordes, M. E. DeCesar, *et al.*, “The NANOGrav Program for Gravitational Waves and Fundamental Physics,” *arXiv e-prints*, 2019.
- [178] R. N. Manchester, “The International Pulsar Timing Array,” *Classical and Quantum Gravity*, vol. 30, p. 224010, Nov. 2013.

- [179] R. Smits, M. Kramer, B. Stappers, D. R. Lorimer, J. Cordes, and A. Faulkner, “Pulsar searches and timing with the square kilometre array,” *Astronomy & Astrophysics*, vol. 493, p. 1161–1170, Nov. 2008.
- [180] M. Punturo, M. Abernathy, F. Acernese, B. Allen, N. Andersson, *et al.*, “The Einstein Telescope: a third-generation gravitational wave observatory,” *Classical and Quantum Gravity*, vol. 27, p. 194002, Sep. 2010.
- [181] B. P. Abbott, R. Abbott, T. D. Abbott, M. R. Abernathy, K. Ackley, *et al.*, “Exploring the sensitivity of next generation gravitational wave detectors,” *Classical and Quantum Gravity*, vol. 34, p. 044001, Jan. 2017.
- [182] D. Lorimer and M. Kramer, “Handbook of Pulsar Astronomy,” *Handbook of pulsar astronomy, by D.R. Lorimer and M. Kramer. Cambridge observing handbooks for research astronomers, Vol. 4. Cambridge, UK: Cambridge University Press, 2004*, Nov. 2004.
- [183] R. W. Hellings and G. S. Downs, “Upper limits on the isotropic gravitational radiation background from pulsar timing analysis.,” *The Astrophysical Journal Letters*, vol. 265, pp. L39–L42, Feb. 1983.
- [184] M. T. Miles, R. M. Shannon, D. J. Reardon, M. Bailes, D. J. Champion, M. Geyer, P. Gitika, K. Grunthal, M. J. Keith, M. Kramer, A. D. Kulkarni, R. S. Nathan, A. Parthasarathy, N. K. Porayko, J. Singha, G. Theureau, F. Abbate, S. Buchner, A. D. Cameron, F. Camilo, B. E. Moreschi, G. Shaifullah, M. Shamohammadi, and V. V. Krishnan, “The MeerKAT Pulsar Timing Array: the 4.5-yr data release and the noise and stochastic signals of the millisecond pulsar population,” *Monthly Notices of the Royal Astronomical Society*, vol. 536, p. 1467–1488, Dec. 2024.
- [185] R. Abbott, T. Abbott, S. Abraham, F. Acernese, K. Ackley, D., D. Blair, R. Blair, F. Bobba, and N. o. Bode, “Upper limits on the isotropic gravitational-wave background from Advanced LIGO and Advanced Virgo’s third observing run,” *Physical Review D*, vol. 104, July 2021.
- [186] M. Kamionkowski and E. D. Kovetz, “The Quest for B Modes from Inflationary Gravitational Waves,” *Annual Review of Astronomy & Astrophysics*, vol. 54, no. 1, pp. 227–269, 2016.

- [187] T. Namikawa, S. Saga, D. Yamauchi, and A. Taruya, “CMB Constraints on the Stochastic Gravitational-Wave Background at Mpc scales,” *Physical Review D*, vol. 100, no. 2, p. 021303, 2019.
- [188] N. Bartolo, C. Caprini, V. Domcke, D. G. Figueroa, J. Garcia-Bellido, *et al.*, “Science with the space-based interferometer LISA. IV: probing inflation with gravitational waves,” *Journal of Cosmology and Astroparticle Physics*, vol. 2016, pp. 026–026, Dec. 2016.
- [189] T. Regimbau, M. Evans, N. Christensen, E. Katsavounidis, B. Sathyaprakash, and S. Vitale, “Digging Deeper: Observing Primordial Gravitational Waves below the Binary-Black-Hole-Produced Stochastic Background,” *Physical Review Letters*, vol. 118, p. 151105, Apr. 2017.
- [190] C. Xin, C. M. F. Mingarelli, and J. S. Hazboun, “Multimessenger pulsar timing array constraints on supermassive black hole binaries traced by periodic light curves,” *arXiv e-prints*, 2021.
- [191] E. Gleave and A. Jaffe, “Collective Experimental Sensitivities to Further Stochastic Gravitational Wave Background Sources,” (In Prep.).
- [192] N. Cornish and T. Robson, “Galactic binary science with the new LISA design,” *Journal of Physics: Conference Series*, vol. 840, p. 012024, May 2017.
- [193] P. Ajith, S. Babak, Y. Chen, M. Hewitson, B. Krishnan, A. M. Sintes, J. T. Whelan, B. Brügmann, P. Diener, N. Dorband, J. Gonzalez, M. Hannam, S. Husa, D. Pollney, L. Rezzolla, L. Santamaría, U. Sperhake, and J. Thornburg, “Template bank for gravitational waveforms from coalescing binary black holes: Nonspinning binaries,” *Physical Review D*, vol. 77, p. 104017, May 2008.
- [194] X.-J. Zhu, E. Howell, T. Regimbau, D. Blair, and Z.-H. Zhu, “Stochastic Gravitational Wave Background from Coalescing Binary Black Holes,” *The Astrophysical Journal*, vol. 739, p. 86, Sept. 2011.
- [195] Y.-C. Bi, Y.-M. Wu, Z.-C. Chen, and Q.-G. Huang, “Implications for the supermassive black hole binaries from the NANOGrav 15-year data set,” *Science China Physics, Mechanics & Astronomy*, vol. 66, Nov. 2023.
- [196] N. Steinle, H. Middleton, C. J. Moore, S. Chen, A. Klein, G. Pratten, R. Buscicchio,

- E. Finch, and A. Vecchio, “Implications of pulsar timing array observations for LISA detections of massive black hole binaries,” *arXiv e-prints*, 2023.
- [197] C. Gowling and M. Hindmarsh, “Observational prospects for phase transitions at LISA: Fisher matrix analysis,” *Journal of Cosmology and Astroparticle Physics*, vol. 2021, p. 039, Oct. 2021.
- [198] M. Pieroni and E. Barausse, “Foreground cleaning and template-free stochastic background extraction for LISA,” *Journal of Cosmology and Astroparticle Physics*, vol. 2020, p. 021–021, July 2020.
- [199] J. Carron, “On the assumption of Gaussianity for cosmological two-point statistics and parameter dependent covariance matrices,” *Astronomy & Astrophysics*, vol. 551, p. A88, Feb. 2013.
- [200] V. De Luca, G. Franciolini, and A. Riotto, “NANOGrav Data Hints at Primordial Black Holes as Dark Matter,” *Physical Review Letters*, vol. 126, p. 041303, Jan. 2021.
- [201] G. Boileau, N. Christensen, C. Gowling, M. Hindmarsh, and R. Meyer, “Prospects for LISA to detect a gravitational-wave background from first order phase transitions,” *Journal of Cosmology and Astroparticle Physics*, vol. 2023, p. 056, Feb. 2023.
- [202] C. Gowling, M. Hindmarsh, D. C. Hooper, and J. Torrado, “Reconstructing physical parameters from template gravitational wave spectra at LISA: first order phase transitions,” *Journal of Cosmology and Astroparticle Physics*, vol. 2023, p. 061, Apr. 2023.
- [203] C. Cecchini, G. Franciolini, and M. Pieroni, “Forecasting Constraints on SIGW with Future Pulsar Timing Array Observations,” *arXiv e-prints*, 2025.
- [204] G.-C. Wang, H.-B. Jin, and X. Zhang, “Detecting stochastic gravitational wave background from cosmic strings with next-generation detector networks: Component separation based on a multi-source astrophysical foreground noise model,” *arXiv e-prints*, 2025.
- [205] I. Caporali, G. Capurri, W. D. Pozzo, A. Ricciardone, and L. V. Dall’Armi, “Impact of correlated noise on the reconstruction of the stochastic gravitational wave background with Einstein Telescope,” *arXiv e-prints*, 2025.
- [206] S. Babak, M. Falxa, G. Franciolini, and M. Pieroni, “Forecasting the sensitivity of Pulsar Timing Arrays to gravitational wave backgrounds,” *arXiv e-prints*, 2024.

- [207] A. R. Pol, A. Neronov, C. Caprini, T. Boyer, and D. Semikoz, “LISA and γ -ray telescopes as multi-messenger probes of a first-order cosmological phase transition,” *arXiv e-prints*, 2023.
- [208] C. Caprini, D. G. Figueroa, R. Flauger, G. Nardini, M. Peloso, M. Pieroni, A. Ricciardone, and G. Tasinato, “Reconstructing the spectral shape of a stochastic gravitational wave background with LISA,” *Journal of Cosmology and Astroparticle Physics*, vol. 2019, p. 017–017, Nov. 2019.
- [209] C. Tan and K. Wang, “Forecasting constraints on the no-hair theorem from the stochastic gravitational wave background,” *arXiv e-prints*, 2024.
- [210] A. Torres-Orjuela, “Joint gravitational wave detection by TianQin and LISA,” *arXiv e-prints*, 2024.
- [211] A. Marriott-Best, D. Chowdhury, A. Ghoshal, and G. Tasinato, “Exploring cosmological gravitational wave backgrounds through the synergy of LISA and ET,” *arXiv e-prints*, 2024.
- [212] B.-R. Wang, J. Li, and H. Wang, “Probing the gravitational wave background from cosmic strings with Alternative LISA-TAIJI network,” *arXiv e-prints*, 2022.
- [213] M. Braglia, J. García-Bellido, and S. Kuroyanagi, “Testing Primordial Black Holes with multi-band observations of the stochastic gravitational wave background,” *Journal of Cosmology and Astroparticle Physics*, vol. 2021, p. 012, Dec. 2021.
- [214] S. Simon *et al.*, “Review Into LAT Scanning Strategy for Simons Observatory,” (In Prep.).
- [215] T. Louis, E. Grace, M. Hasselfield, M. Lungu, L. Maurin, *et al.*, “The Atacama Cosmology Telescope: two-season ACTPol spectra and parameters,” *Journal of Cosmology and Astroparticle Physics*, vol. 2017, p. 031–031, June 2017.
- [216] S. Simon *et al.*, “Scan Strategy Requirements,” tech. rep., The Simons Observatory, (in prep.). SO-DATA-DOC-002B.
- [217] S. K. Choi, M. Hasselfield, S.-P. P. Ho, B. Koopman, M. Lungu, *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*, vol. 2020, p. 045–045, Dec. 2020.

- [218] S. Simon, M. Silva-Feaver, J. Godmundsson, H. McCarrick, and M. Lungu, “Simons Observatory (SO) Instrument Model.” SO-DATA-ICD-004F, 2023.
- [219] E. Keihänen, R. Keskitalo, H. Kurki-Suonio, T. Poutanen, and A.-S. Sirviö, “Making cosmic microwave background temperature and polarization maps with MADAM,” *Astronomy & Astrophysics*, vol. 510, p. A57, Feb. 2010.
- [220] S. Simon, K. Ganga, E. Gleave, R. Keskitalo, A. Mangu, D. Alonso, and A. Kusaka, “Simons Observatory SAT Scan Review.” (Internal review slides), 2023.
- [221] K. Yamada, B. Bixler, Y. Sakurai, P. C. Ashton, J. Sugiyama, *et al.*, “The Simons Observatory: Cryogenic half wave plate rotation mechanism for the small aperture telescopes,” *Review of Scientific Instruments*, vol. 95, p. 024504, Feb. 2024.
- [222] T. Tsan, N. Galitzki, A. M. Ali, K. Arnold, G. Coppi, *et al.*, “The effects of inclination on a two stage pulse tube cryocooler for use with a ground based observatory,” *Cryogenics*, vol. 117, p. 103323, Jul. 2021.
- [223] G. Capurri, A. Lapi, and C. Baccigalupi, “Detectability of the Cross-Correlation between CMB Lensing and Stochastic GW Background from Compact Object Mergers,” *Universe*, vol. 8, p. 160, Mar. 2022.
- [224] W. Buchmuller, V. Domcke, H. Murayama, and K. Schmitz, “Probing the scale of grand unification with gravitational waves,” *Physics Letters B*, vol. 809, p. 135764, Oct. 2020.
- [225] B. P. Abbott, R. Abbott, T. D. Abbott, S. Abraham, F. Acernese, *et al.*, “Prospects for observing and localizing gravitational-wave transients with Advanced LIGO, Advanced Virgo and KAGRA,” *Living Reviews in Relativity*, vol. 23, Sept. 2020.
- [226] E. D. Hall, “Cosmic Explorer: A Next-Generation Ground-Based Gravitational-Wave Observatory,” *Galaxies*, vol. 10, no. 4, 2022.
- [227] T. Robson, N. J. Cornish, and C. Liu, “The construction and use of LISA sensitivity curves,” *Classical and Quantum Gravity*, vol. 36, p. 105011, Apr. 2019.
- [228] R. W. Hellings and G. S. Downs, “Upper limits on the isotropic gravitational radiation background from pulsar timing analysis,” *The Astrophysical Journal Letters*, vol. 265, pp. L39–L42, Feb. 1983.

- [229] K. Schmitz, “New sensitivity curves for gravitational-wave signals from cosmological phase transitions,” *Journal of High Energy Physics*, vol. 2021, Jan. 2021.
- [230] A. I. Renzini, B. Goncharov, A. C. Jenkins, and P. M. Meyers, “Stochastic Gravitational-Wave Backgrounds: Current Detection Efforts and Future Prospects,” *Galaxies*, vol. 10, no. 1, 2022.
- [231] LIGO Scientific Collaboration and Virgo Collaboration, “Searching for stochastic gravitational waves using data from the two colocated LIGO Hanford detectors,” *Physical Review D*, vol. 91, p. 022003, Jan. 2015.
- [232] K. Kohri and T. Terada, “Solar-mass primordial black holes explain NANOGrav hint of gravitational waves,” *Physics Letters B*, vol. 813, p. 136040, Feb. 2021.
- [233] Y. Gouttenoire, *Gravitational Waves from Cosmic Strings*, pp. 419–499. Cham: Springer International Publishing, 2022.
- [234] D. I. Dunsky, L. J. Hall, and K. Harigaya, “A Heavy QCD Axion and the Mirror World,” *arXiv e-prints*, 2023.
- [235] I. P. Ivanov, “Building and testing models with extended Higgs sectors,” *Progress in Particle and Nuclear Physics*, vol. 95, p. 160–208, July 2017.
- [236] N. Barnaby, E. Pajer, and M. Peloso, “Gauge field production in axion inflation: Consequences for monodromy, non-Gaussianity in the CMB, and gravitational waves at interferometers,” *Physical Review D*, vol. 85, p. 023525, Jan. 2012.

Appendix A

Foreground Models for the SGWB

In this thesis the astrophysical foregrounds are modelled with the function

$$\Omega_{\text{fg}}(f) = A f^{2/3} \left[\exp \left(- (f/f_{\text{ref}})^B \right) + D \sin^2 (\pi f / f_{\text{ref}}) \right] \times [1 + \tanh (E (f_* - f))], \quad (\text{A.1})$$

with the coefficients for each binary population given in table A.1. Equation A.1 is an adaptation of a WD-WD model developed by Cornish and Robson [192], here converted to Ω from S :

$$\Omega_{\text{fg}}(f) = A f^{2/3} \exp (-f^\alpha + \beta \sin(\kappa f)) \times [1 + \tanh(\gamma(f_k - f))] . \quad (\text{A.2})$$

I developed the new function to give a more general fit across the different binary populations. Both functions behave as a power law ($\Omega \propto f^{2/3}$) at low f , with the index $2/3$ matching theoretical predictions (see section 4.2). At high f (i.e. $f \gtrsim f_*$ or $f \gtrsim f_k$) the behaviour of both functions again matches; here the tanh component suppresses the function, with a transition sharpness dictated by E or γ .

The two functions differ in their behaviour at mid-frequencies. The original function's exponential damping term $\exp(-f^\alpha)$ is replaced with a more tunable term incorporating a reference frequency: $\exp([f/f_{\text{ref}}]^B)$. This is particularly important in generalizing the function, as the original function only had to work well for one frequency range (that for WD binaries) whereas the new one must span a much wider range.

The other difference comes in the placement of the oscillatory term, where the new function moves it outside of the exponent. The new function has oscillations which are tuned with f_{ref} such that the $\sin^2$ term peaks only once in the relevant frequency range, and with D to create the right amplitude to shape the overall function's peak. This term could have been left within the exponent, but the non-linear and asymmetric behaviour of the oscillations is harder to fit

across the different populations.

Pop.	A [$\text{Hz}^{-2/3}$]	B	D	E [Hz^{-1}]	f_* [Hz]	f_{ref} [Hz]	Model
WD-WD	9.48×10^{-9}	2.25	0.00449	8527	0.00346	0.00449	[192, 164]
BH-BH	6.85×10^{-11}	2.69	0.00148	0.00139	42.5	793	[223, 165, 166]
BH-NS	7.34×10^{-12}	4.08	2.96×10^8	0.00139	1571	2771	[223, 165, 166]
NS-NS	6.89×10^{-12}	3.83	0.00242	0.000648	6481	5961	[223, 165, 166]

Table A.1: Coefficients for the foreground model used in this work, for the different populations of binaries. These were found with a least-squares fitting to the population model spectra given in the right-hand column: WD-WD from the Sloan Digital Sky Survey (SDSS) [164], and BH-BH, BH-NS, NS-NS from LVK data [165, 166].

Appendix B

GW Experiment Noise Curves

Given below are the nominal sensitivities Ω_{nominal} of the experiments considered in this work. These are given in terms of $S_{\text{nominal}}(f)$, and can be converted to Ω_{nominal} through equation 6.3. Where simple analytic approximations to the instrumental noise have been used, they are given below. Otherwise, references with more complicated fits are listed.

In general, the strain S_{noise} can be calculated through

$$S_{\text{noise}}(f) = \frac{D_{\text{noise}}(f)}{\mathcal{R}(f)}, \quad (\text{B.1})$$

where D_{noise} is the detector noise function, and $\mathcal{R}$ is the detector response function.

B.1 Ground-Based Interferometers

- **LIGO-Virgo-KAGRA (LVK):** The strain for a series of detectors involves the calculation of the combined network response function $\mathcal{R}_{\text{HLVK}}$:

$$\mathcal{R}_{\text{HLVK}}^2(f) = \Gamma_{\text{HL}}^2(f) + \Gamma_{\text{HV}}^2(f) + \Gamma_{\text{HK}}^2(f) + \Gamma_{\text{LV}}^2(f) + \Gamma_{\text{LK}}^2(f) + \Gamma_{\text{VK}}^2(f), \quad (\text{B.2})$$

where each of the overlap reduction functions $\Gamma_{i,j}^2(f)$ denotes the sensitivity of a pair of detectors i and j to the [SGWB](#). Here, the detectors are [LIGO](#) (Hanford), [LIGO](#) (Livingston), [Virgo](#) and [KAGRA](#).

Obtaining individual expression for each $\Gamma_{i,j}$, in addition to a shared D , is complicated, and instead a direct S curve from the [LVK](#) collaboration was used—this can be found in reference [\[225\]](#).

- **Cosmic Explorer (CE):** There are not currently simple analytic models of [CE](#)’s pre-

dicted D and $\mathcal{R}$, and again a direct S curve was used, which can be found in reference [226].

- **Einstein Telescope (ET)**: Likewise for ET. Information on the strain curve used in the work in this thesis can be found in reference [226].

B.2 Space-Based Interferometers

- **Laser Interferometer Space Antenna (LISA)**: The strain is given by

$$S = \frac{40}{3} \times \left[\frac{S_1(f)}{(2\pi f)^4} + S_2(f) \right] \times \left(1 + \left(\frac{3f}{4f_t} \right)^2 \right), \quad (\text{B.3})$$

where

$$f_t = (3 \times 10^8 \text{ m s}^{-1}) / (2\pi(2.5 \times 10^8 \text{ m})) \approx 0.2 \text{ Hz}, \quad (\text{B.4})$$

is the characteristic transfer frequency, determined by the arm length $L = 2.5 \times 10^8 \text{ m}$.

$$S_1 = 5.76 \times 10^{-48} \text{ s}^{-4} \text{ Hz}^{-1} \times (1 + (0.4 \text{ mHz}/f)^2) \text{ Hz}^{-1} \quad (\text{B.5})$$

is the single test mass acceleration noise, which gives the power spectral density (PSD) of non-gravitational forces which act on the test mass. And

$$S_2 = 3.6 \times 10^{-41} \text{ Hz}^{-1} \quad (\text{B.6})$$

is the single-link optical metrology noise, which quantifies the PSD of distance measurements noise (per arm) from the laser system. [227, 197, 173].

- **Deci-hertz Interferometer Gravitational wave Observatory (DECIGO)**: The strain is given by

$$S = \left[7.05 \times 10^{-48} \times \left(1 + \left(\frac{f}{7.36 \text{ Hz}} \right)^2 \right) \right] + \left[\frac{4.8 \times 10^{-51}}{1 + \left(\frac{f}{7.36 \text{ Hz}} \right)^2} \times \left(\frac{f}{1 \text{ Hz}} \right)^{-4} \right] + \left[5.33 \times 10^{-52} \times \left(\frac{f}{1 \text{ Hz}} \right)^{-4} \right] \text{ Hz}^{-1}, \quad (\text{B.7})$$

where the first term shows the shot noise, the second the radiation pressure noise, and the third the acceleration noise [176].

- **Big Bang Observer (BBO)**: The strain is given by

$$S = \left[2.00 \times 10^{-49} \left(\frac{f}{1^{-4}} \right)^2 + 4.58 \times 10^{-49} + 1.26 \times 10^{-51} \left(\frac{f}{1^{-4}} \right)^{-4} \right] \text{ Hz}^{-1}, \quad (\text{B.8})$$

where the three terms likewise represent the shot, radiation pressure, and acceleration noises [176].

B.3 Pulsar Timing Arrays

The strain of a pulsar timing array is given by

$$S = \sqrt{\frac{2}{N(N-1)}} \times \frac{D}{\zeta_{\text{rms}} R}, \quad (\text{B.9})$$

where N is the number of pulsars observed, $D = 2t\sigma^2$ is the intrinsic (white) timing noise per pulsar, t is the average cadence of observations, σ is the standard deviation of the detector noise, $\zeta_{\text{rms}} = 0.147$ is the rms value for the Hellings-Downs factor [228] between two pulsars, and $R = 1/(12\pi^2 f^2)$ is the detector response function [229].

- **International Pulsar Timing Array (IPTA)**: The strain is given by equation B.9, with $N = 20$, $t = 1.208 \times 10^6 \text{ s}$, $\sigma = 1.0 \times 10^{-7} \text{ s}$ [229].
- **Square Kilometre Array (SKA)**: The strain is given by equation B.9, with $N = 50$, $t = 6.04 \times 10^5 \text{ s}$, $\sigma = 3.0 \times 10^{-8} \text{ s}$ [229].

Appendix C

SGWB Fiducial Models

Table C.1

Model	Details	f_1	f_2	Ω_2	m	Ref.
Inflation:						
inflation1	SFSR, $r = 0.1$	-9.94	2.42	1.76×10^{-16}	0.00	[230]
inflation2	SFSR, $r = 0.2$	-9.94	3.07	2.64×10^{-16}	-0.02	[231]
Primordial Black Holes:						
PBHs1*	solar mass PBHs	—	—	—	—	[232]
PBHs2	PBHs which would satisfy all of dark matter	-9.02	-1.25	2.47×10^{-10}	-0.01	[200]
Cosmic Strings:						
strings1a	$G\mu \approx 10^{-9}$	-8.84	0.78	6.45×10^{-10}	-0.07	[132]
strings1b	$G\mu \approx 10^{-11}$	-7.11	4.30	6.17×10^{-11}	-0.07	[132]
strings1c	$G\mu \approx 10^{-13}$	-4.76	4.15	7.55×10^{-12}	-0.08	[132]
strings1d	$G\mu \approx 10^{-15}$	-2.92	4.50	9.24×10^{-13}	-0.09	[132]
strings1e	$G\mu \approx 10^{-17}$	-0.87	3.82	9.63×10^{-1}	-0.12	[132]
strings1f	$G\mu \approx 10^{-19}$	-0.87	3.82	9.63×10^{-1}	-0.12	[132]
strings1g	$G\mu \approx 10^{-21}$	-0.87	3.82	9.63×10^{-1}	-0.12	[132]
strings1h	$G\mu \approx 10^{-23}$	-0.87	3.82	9.63×10^{-1}	-0.12	[132]
strings2	secondary population of small loops	0.53	2.13	9.68×10^{-8}	0.43	[233]
strings3i	$G\mu \approx 10^{-7}$ (seg. 1)	-5.78	-0.56	7.53×10^{-9}	-0.38	[152]

Continued on next page

strings3ii	$G\mu \approx 10^{-7}$ (seg. 2)	0.00	3.76	5.02×10^{-9}	-0.01	[152]
First-Order Phase Transitions:						
FOPTs1a	electroweak (≈ 200 GeV)	-9.99	-7.13	4.68×10^{-7}	2.83	[111]
FOPTs1b	QCD (≈ 200 MeV)	-3.64	-1.19	1.00×10^{-16}	-3.85	[111]
FOPTs2a	QCD, $\beta/H = 10^2$	-2.07	0.41	1.40×10^{-18}	-1.53	[234]
FOPTs2b	QCD, $\beta/H = 10^3$	-1.52	0.20	1.01×10^{-18}	-1.38	[234]
FOPTs2c	QCD, $\beta/H = 10^4$	-1.44	-0.61	2.90×10^{-18}	-0.47	[234]
FOPTs3ai	$T_n = 10^7$ GeV (seg. 1)	-0.09	0.76	2.31×10^{-9}	2.85	[235]
FOPTs3aii	$T_n = 10^7$ GeV (seg. 2)	1.30	3.63	1.30×10^{-11}	-0.99	[235]
FOPTs3bi	$T_n = 10^8$ GeV (seg. 1)	0.71	1.68	1.99×10^{-9}	2.80	[235]
FOPTs3bii	$T_n = 10^8$ GeV (seg. 2)	2.29	4.70	1.21×10^{-11}	-1.00	[235]
FOPTs3ci	$T_n = 10^3$ GeV (seg. 1)	-3.60	-3.08	1.37×10^{-9}	2.63	[235]
FOPTs3cii	$T_n = 10^3$ GeV (seg. 2)	-2.60	-0.47	1.21×10^{-11}	-1.01	[235]
Exotic Phenomena:						
axions1i	Axion-based (seg. 1)	inflation	-3.68	-1.33	2.04×10^{-11}	1.18 [236]
axions1ii	Axion-based (seg. 2)	inflation	-0.12	2.49	1.22×10^{-8}	0.47 [236]

Table C.1: Details of the fiducial models presented in figure 4.1, and used as fiducial models in Chapter 7. The values listed for $f_1 = \log_{10}(f_{\min}/\text{Hz})$, $f_2 = \log_{10}(f_{\max}/\text{Hz})$ correspond to the frequency range over which these models can be usefully approximated by a simple power law, where applicable. $\Omega_2 = \Omega(f_{\max})$ is then the power at $f_{\max}$, and m the power law index. All values are approximate. Some fiducial models are split into multiple power law ‘segments’ (seg.s).

* Shown in figure 4.1 for reference, but cannot be fit to a simple power law.

Appendix D

Additional LAT Scan Strategy Results

D.1 Initial Simulation Run Results

D.1.1 Cross-linking Maps




Figure D.1: Maps of the cross-linking statistic P_x (see equation 8.1), with separate colouring for pixels between discrete cut-off values. Figure credit: SO Collaboration.

D.2 Improved Simulation Run Results

D.2.1 Depth and Noise




Figure D.2: Maps of the depth, and both noise simulations, here with CAR rather than HEALpix pixelization. Figure credit: SO Collaboration.

D.2.2 Elevation Noise Coefficients

Optics Tube	A	C
LF1	0.06	0.92
LF2	0.16	0.79
MF1	0.16	0.8
MF2	0.17	0.78
UHF1	0.29	0.62
UHF2	0.36	0.53

Table D.1: Coefficient values for the Large Aperture Telescope (LAT) elevation noise, as used in equation 8.3.

Appendix E

Additional SAT Scan Strategy Results

E.1 Figures of Merit

Elevation	Azimuth Throw	Boresight Direction	% Time Observing	% Time Observing CMB
50	30	SW-SE	78	57
50	30	SW-N-SE	79	66
50	30	SW-NE-NW-SE	79	63
50	40	SW-SE	77	56
50	40	SW-N-SE	78	65
50	40	SW-NE-NW-SE	78	62
55	30	SW-SE	77	56
55	30	SW-N-SE	77	65
55	30	SW-NE-NW-SE	77	62
55	40	SW-SE	76	55
55	40	SW-N-SE	76	64
55	40	SW-NE-NW-SE	77	61
60	30	SW-SE	75	55
60	30	SW-N-SE	75	63
60	30	SW-NE-NW-SE	76	61
60	40	SW-SE	74	54
60	40	SW-N-SE	74	62
60	40	SW-NE-NW-SE	75	60
2018 baseline			51	50

Table E.1: The total and CMB observing efficiencies for the three most promising boresight direction scan types, where the maximum efficiency (taking into account a two-month observing break) is 83%.

Elevation	Azimuth Throw	Boresight Direction	$f_{\text{sky}}^{\text{SS}}$	$f_{\text{sky}}^{\text{SN}}$	$f_{\text{sky}}^{\text{NN}}$	$\sigma(r) \times 10^{-3}$	
50	40	S	0.071	0.079	0.086	1.857	
55	40	S	0.080	0.091	0.099	1.867	
2018 optimistic			0.114	0.152	0.184	1.883	
60	40	S	0.090	0.102	0.111	1.892	
50	40	SW-SE	0.131	0.150	0.203	1.911	
50	30	S	0.066	0.075	0.082	1.937	
55	40	SW-SE	0.135	0.173	0.204	1.937	
50	30	SW-N-SE	0.139	0.216	0.288	1.951	
55	30	S	0.075	0.087	0.095	1.961	
50	40	W-E	0.197	0.240	0.273	1.968	
50	40	E	0.210	0.248	0.278	1.969	
50	30	SW-SE	0.125	0.160	0.188	1.978	
60	40	SW-SE	0.142	0.179	0.208	1.980	
55	40	E	0.205	0.242	0.271	1.994	
55	40	W-E	0.192	0.234	0.266	1.994	
50	40	SW-NE-NW-SE	0.169	0.236	0.290	2.015	
55	30	SW-SE	0.131	0.165	0.192	2.029	
60	30	S	0.084	0.098	0.108	2.030	
60	40	E	0.201	0.237	0.266	2.037	
55	40	SW-NE-NW-SE	0.169	0.232	0.282	2.037	
60	40	W-E	0.189	0.230	0.262	2.038	
50	30	E	0.204	0.236	0.261	2.042	
50	30	W-E	0.188	0.226	0.255	2.042	
50	40	SW-N-SE	0.148	0.231	0.309	2.065	
60	40	SW-NE-NW-SE	0.177	0.233	0.278	2.074	
50	30	SW-NE-NW-SE	0.158	0.220	0.250	2.092	
55	30	E	0.201	0.232	0.257	2.095	
55	40	SW-N-SE	0.153	0.235	0.311	2.099	
60	30	SW-SE	0.141	0.173	0.199	2.119	
55	30	SW-NE-NW-SE	0.163	0.220	0.267	2.140	
55	30	W-E	0.186	0.223	0.252	2.144	
60	40	SW-N-SE	0.164	0.244	0.315	2.148	
60	30	W-E	0.184	0.221	0.250	2.189	
60	30	E	0.199	0.231	0.256	2.192	
55	30	SW-N-SE	0.148	0.225	0.296	2.209	
60	30	SW-NE-NW-SE	0.171	0.223	0.266	2.231	
60	30	SW-N-SE	0.159	0.236	0.304	2.231	
2018 pessimistic			0.114	0.152	0.184	2.412	

Table E.2: The forecasted constraints on r for each of the scan variations considered, ranked by their $\sigma(r)$.