The Cosmology and Astrophysics of Axion-like Particles

Andrew Powell

St. Cross College

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

Trinity 2016

Abstract

In this thesis I study astrophysical and cosmological effects of axion-like particles (ALPs). ALPs are pseudo-scalar particles, which are generally very weakly-interacting, with a coupling $\frac{a}{M} E \cdot B$ to electromagnetism. They are predicted by many theories which extend the standard model (SM) of particle physics, most notably string theory. String theory compactifications also predict many scalar fields called moduli which describe the size and shape of the extra, compact dimensions.

In string theory models generically the moduli fields are responsible for reheating the universe after inflation. Being gravitationally-coupled, they will also decay to any other particles or sectors of the theory, including any light ALPs, of which there are usually many. The ALPs produced by moduli decay will contribute to dark radiation, additional relativistic energy density. The amount of dark radiation is tightly constrained by observations, this bounds the branching fraction of moduli decays into ALPs, which constrains the string theory model itself. I calculate the amount of dark radiation produced in a model with one light modulus, solely responsible for reheating, called the Large Volume Scenario. I study a minimal version of this model with one ALP and a visible sector comprised of the minimal supersymmetric SM. The dominant visible sector decay mode is to two Higgses, I include radiative corrections to this decay and find that ALP dark radiation is over-produced in this minimal version of the model, effectively ruling it out.

The production of ALPs from moduli decay at reheating seems to be a generic feature of string theory models. These ALPs would exist today as a homogeneous cosmic ALP background (CAB). The coupling of ALPs to electromagnetism allows ALPs to convert to photons and vice versa in a magnetic field, leading to potential observable astrophysical signals of this CAB. Observations have shown an excess in soft X-ray emission from many galaxy clusters. I use detailed simulations of galaxy cluster magnetic fields to show that a CAB can explain these observations by conversion of ALPs into X-ray photons. I simulate ALP–photon conversion in four galaxy clusters and compare to soft X-ray observations. I show the excesses (or lack thereof) can be fit consistently across the clusters for a CAB with ALP–photon inverse coupling of $M = 6 - 12 \times 10^{12}$ GeV, if the CAB spectrum has energy ~ 200 eV.

I also study the possibility of using galaxy clusters to search for and constrain the ALP coupling to photons using cluster X-ray emission. Conversion of X-ray photons into ALPs will cause spectral distortions to the thermal X-ray spectrum emitted by galaxy clusters. I show that the non-observation of these distortions is able to produce the strongest constraints to date on the ALP–photon inverse coupling, $M \gtrsim 7 \times 10^{11}$ GeV.

Acknowledgements

I would like to start by thanking my supervisor Joseph Conlon. I am grateful for his support, encouragement and help during my DPhil. I am also very grateful for the many interesting projects and topics he has introduced me to, and for the tips and advice for writing this thesis. I would also like to thank my examiners Pedro Ferreira and Mark Goodsell for taking the time to read this thesis, and prodiving insightful and interesting questions and comments about the work.

I would also like to thank my collaborators Stephen Angus, Marcus Berg, Francesca Day, Giovanni Grilli di Cortona, Ulrich Haisch, Edward Hardy, Nicholas Jennings, David Marsh, Markus Rummel, Sven Krippendorf, and Lukas Witkowski, who were a pleasure to work with. I would also like to thank the many other people, lecturers and fellow students who have helped me with my research, and have made the last four years fly by. An incomplete list includes Eirik Svanes, David Kraljic, James Scargill, James Scoville, Jim Talbert, Alexander Karlberg, Pedro Alvarez, Robert Lasenby, Subir Sarkar, and many others. I would especially like to thank Ed Hardy for taking the time to proof-read this thesis.

Most of all I would like to thank my wife, Steph, whose love, support and encouragement has helped me no end throughout my DPhil, and during our seven wonderful years together. I also thank my family and all my other friends not mentioned above.

Finally I am very grateful to the STFC and the ERC for providing me with the funding for my studies, and for allowing me to travel to many interesting conferences, schools, and workshops.

Statement of Originality

This thesis is based on original research and contains no material that has already been accepted, or is concurrently being submitted, for any degree or diploma or certificate or other qualification in this university or elsewhere. To the best of my knowledge and belief this thesis contains no material previously published or written by another person, except where due reference is made in the text.

Andrew Powell 2016

Contents

1 Introduction and Motivation **1**
1.1 Outline . . . 4

2 String Moduli and Cosmology **7**
2.1 Dark Radiation . . . 8
2.2 String Moduli . . . 12
2.3 Moduli in Cosmology . . . 14

3 Axion-like Particles **18**
3.1 The QCD Axion . . . 18
3.2 String Axions and Axion-like Particles . . . 21
3.3 Searching for ALPs . . . 23
3.3.1 ALP–photon Mixing . . . 24
3.3.2 Astrophysics . . . 26
3.3.3 Terrestrial Experiments . . . 29

4 ALP Dark Radiation Production in Large Volume Scenarios **32**
4.1 Dark Radiation . . . 33
4.2 The Large Volume Scenario . . . 34
4.3 Tree Level Dark Radiation Prediction . . . 36
4.3.1 Decay Rates . . . 36
4.3.2 Dark Radiation Prediction . . . 38
4.4 Analytic Results . . . 39
4.4.1 Running of Volume Modulus Higgs Coupling . . . 40
4.5 Numerical Results . . . 42
4.5.1 Solution of RG Equations . . . 42
4.5.2 SM and MSUGRA Parameter Dependencies . . . 43
4.5.3 Predictions for the Effective Excess Number of Neutrinos . . . 45
4.6 Conclusions . . . 47

5 A Cosmic ALP Background and the Galaxy Cluster Soft X-ray Excess **50**
5.1 Introduction . . . 51
5.2 A Cosmic ALP Background . . . 53
5.3 The Cluster Soft X-ray Excess . . . 56
5.3.1 History of Soft Excess Observations . . . 56
5.3.2 Specific Clusters . . . 59
5.3.3 Astrophysical Models of the Soft Excess . . . 62
5.4 Cluster Magnetic Fields . . . 66
5.4.1 Magnetic Field Observations . . . 66
5.4.2 Magnetic Field Model . . . 69
5.4.3 Magnetic Field Generation . . . 77
5.5 ALP–photon Propagation . . . 77

5.5.1 Homogeneous Solution 78
5.5.2 Inhomogeneous Magnetic Fields 80
5.6 Consistency Checks 82
5.7 Coma 83
5.7.1 Luminosity Calculation 84
5.7.2 Results 84
5.7.3 Summary of Coma Results 100
5.8 A665, A2199 and A2255 101
5.8.1 Luminosity and Fractional Excess 101
5.8.2 Results 103
5.8.3 Comparison and Summary 113
5.9 Conclusions 117

6 Constraining ALPs using Galaxy Clusters **121**
6.1 Current Bounds 121
6.2 Photon-ALP Conversion in Galaxy Clusters 122
6.3 Cluster Spectral Distortions from ALPs 124
6.4 Further Properties of the Distorted Spectrum 128
6.5 Conclusions 130

7 Summary and Conclusions **133**

References **139**

1 Introduction and Motivation

With the discovery of the Higgs boson, the Standard Model (SM) of particle physics is now complete. The SM is a theory of the electromagnetic, weak, and strong interactions and the three generations of quarks and leptons on which they act. The interactions are mediated by the bosons, the photon, gluon, and the W/Z. The Higgs mechanism solves the puzzle of how the W/Z bosons, the carriers of the weak interaction, get their masses. It is also responsible for giving the leptons and quarks their masses.

However we know that the SM does not describe every interaction between fundamental particles. The theory does not include gravity, which is described at low energies by the General Theory of Relativity (GR). The SM is a quantum theory, however treating GR in this way leads to an 'effective field theory' valid only for energies well below the Planck scale, $M_{pl} = 1.22 \times 10^{19}$ GeV. At energies close to the Planck scale the theory breaks down. Thus it is thought that some new quantum gravity theory must complete GR close to this scale.

We have other reasons for believing the SM is not the end of the story for particle physics. Astronomical and cosmological observations have shown that the majority of matter in the universe is not made up of particles from the SM. Instead, dark matter

accounts for more than 80% of the matter in the universe. Even more strangely, 68% of the energy density of the universe appears to reside as so-called dark energy. This is a fluid whose energy density stays constant with the expansion of the universe, thus the importance of which grows with time. It is a puzzle why this dark energy is only now starting to dominate the universe, as is the seemingly arbitrary value of the energy density, and the fact that its energy density is so much smaller than the natural particle physics scales.

In addition, the SM alone cannot explain the particle anti-particle asymmetry observed in the universe. Neutrino masses, and why they are so much smaller than the other fundamental particles' masses, are also unexplained. There is no explanation for why the Higgs mechanism breaks electroweak symmetry at scales much lower than the Planck scale, $m_{weak} \ll M_{pl}$. The Higgs mass should receive radiative corrections proportional to any new mass scale that accompanies the solutions to the above problems. Thus we should expect the electroweak symmetry to be broken at a much higher scale, this is called the electroweak hierarchy problem. The most popular solution to this is supersymmetry, a new symmetry between fermions and bosons. Each particle then has a 'superpartner' with spin differing by half a unit. This symmetry causes the problematic corrections to the Higgs mass to cancel, explaining the lightness of the Higgs boson.

More abstractly, there are issues such as explaining why there are three generations of quarks and leptons, or why the Yukawa couplings of the fermions vary by so many orders of magnitude, or even why we observe four space-time dimensions. In sum, there are many reasons why there is a need to investigate theories beyond the Standard Model (BSM).

The best understood quantum gravity theory is string theory. In string theory the fundamental degrees of freedom are supersymmetric one-dimensional objects called strings. These strings can only be consistently embedded in a ten-dimensional space-time. Particles arise as excitations of this string, one of these is the graviton, the massless spin-2 particle which mediates the gravitational force. At very high energies

(typically) $M_{string} \gg m_{weak}$ the true one-dimensional nature of particles is resolvable, solving the problems which plague GR close to the Planck scale. However this feature of string theory is also one of its biggest challenges. Physics predicted by string theory is inaccessible to particle colliders and thus not directly testable.

A major effort in the last few decades has gone into string phenomenology—connecting string-scale physics to the physics of the low energy world we live in. String theory requires ten-dimensional space-time, but we only observe four dimensions, so the six extra spatial dimensions must only be resolvable at very high energies, or very small distances, close to the string scale. This 'compactification' of the extra dimensions leads to a rich low energy four-dimensional effective field theory. This must include the matter and force content of the SM, but will also typically include many new scalar fields called moduli and axion-like particles. Moduli fields arise from compactification of the ten-dimensional graviton, whereas axion-like particles come from the compactification of other massless fields in the ten-dimensional spectrum. The moduli can have very important consequences in cosmology. Their existence can change the early evolution of the universe, adding periods of matter domination. It is then the decay of moduli which produces the hot bath of SM particles from which the usual cosmological evolution follows. The details of the decay are important since the moduli can also decay to particles such as axion-like particles. Part of this thesis will be concerned with the phenomenological implications of such decays to axion-like particles in the early universe, and in that sense we will be probing a class of string theory models.

Axion-like particles (ALPs) are pseudo-scalar particles. They are generally very light, very weakly interacting and have a coupling to electromagnetism. They are abundant in string theory and are also predicted in many other BSM models. The QCD axion, after which ALPs are named, is a pseudo-Goldstone boson which is hypothesised to solve another problem with the SM, the lack of CP violation in QCD. The QCD Lagrangian contains a P (parity) and CP (charge-parity) violating term $\sim \theta\, G_{\mu\nu}\tilde{G}^{\mu\nu}$, where G is the gluon field strength tensor. Its existence leads to a

neutron electric dipole moment (EDM) proportional to the dimensionless parameter θ in the Lagrangian. Such a neutron EDM has not been observed, requiring $\theta < 10^{-10}$. This is unnaturally small compared to its expected value $\theta \sim \mathcal{O}(1)$. The situation is even worse as this term receives an additional contribution from any complex phases in the quark mass matrix. Thus the two contributions to θ must each be very small, or cancel to at least one part in 10^{10}. This is the strong CP problem.

The QCD axion solves this problem since it has the same coupling to QCD gauge fields as θ. QCD induces a potential for the axion, such that at the minimum of this potential the axion has a vacuum expectation value $a = -\theta$, and thus the effective CP violating parameter is set to zero, solving the strong CP problem. ALPs are in the same class of particles, but they have no coupling to QCD. The ALPs we study are very light and couple to photons through the term $aF_{\mu\nu}\tilde{F}^{\mu\nu}$. This interaction allows ALPs to be detected. We will use this coupling to both search for the ALPs potentially produced in moduli decays in the early universe, and to constrain the properties of ALPs.

1.1 Outline

The outline of this thesis is as follows. In Chapter 2 I review the standard cosmological model, before reviewing the prediction of moduli fields in string theory compactifications. I outline how the existence of moduli can drastically change the evolution of the early universe, with the universe potentially going through several periods of moduli energy domination. It is the moduli fields then that are responsible for reheating the universe after inflation. I discuss one cosmological observable which can be used to test string theory models of the early universe, the number of relativistic degrees of freedom or 'dark radiation'. The amount of dark radiation is observable both in the cosmic microwave background (CMB) and in the primordial abundances of light elements.

In Chapter 3 I will motivate the study of axion-like particles by firstly introducing the QCD axion as the solution to the strong CP problem. String theory models

often feature many very light, very weakly-coupled particles similar to the QCD axion, known as axion-like particles (ALPs). ALPs can be observed by their coupling to electromagnetism, which allows them to convert into photons and vice versa in external magnetic fields. I review this process, before discussing the constraints on this coupling from astrophysical observations and terrestrial experiments.

The first non-introductory chapter will be Chapter 4. In this chapter I look at dark radiation production in a class of string theory models called the Large Volume Scenario (LVS). In these models there is one lightest modulus that is responsible for reheating. I calculate the branching ratio of this modulus to ALPs, including relevant radiative corrections to the visible sector decay rate. We will see that even in the minimal version of this model, with one ALP, dark radiation is over-produced. Thus ruling out the minimal version of the LVS.

Nevertheless it is clear that ALP production by modulus decay is a fairly generic prediction in string theory models. In Chapter 5 I will discuss a possible way to observe these light ALPs produced during moduli decay. These ALPs will form a homogeneous background and the conversion of this in astrophysical magnetic fields leads to potential signals. X-ray observations have revealed a 'soft X-ray excess' in many galaxy clusters. If this cosmic ALP background (CAB) has X-ray energies now (set by the parent modulus' mass), then this excess can be explained by ALP–photon conversion. I simulate the conversion of ALPs into photons in the magnetic fields of four galaxy clusters and compare to observations. I also look in detail at the predicted morphology in these clusters, and show that a CAB reproduces the observed trend of an excess which increases with cluster radii.

In Chapter 6 I again focus on ALP–photon conversion in galaxy clusters, this time looking at the reverse process. I look at potential distortions to the X-ray spectrum emitted by galaxy clusters that would be caused by photons converting to ALPs. This study is independent of any specific BSM model like string theory, as it does not rely on the production of ALPs in the early universe. I show that the lack of observation of this effect is able to provide the best limits on the ALP–photon coupling for low

mass ALPs. Finally, I will conclude in Chapter 7.

This thesis is based on the following papers:

- S. Angus, J. P. Conlon, U. Haisch and A. J. Powell, *Loop corrections to ΔN_{eff} in large volume models*, JHEP **12** (2013) 061, arXiv:1305.4128 [1],
- S. Angus, J. P. Conlon, M. C. D. Marsh, A. J. Powell, L. T. Witkowski, *Soft X-ray Excess in the Coma Cluster from a Cosmic ALP Background*, JCAP **1409** (2014), no. 09 026, arXiv:1312.3947 [2],
- A. J. Powell, *A Cosmic ALP Background and the Cluster Soft X-ray Excess in A665, A2199 and A2255*, JCAP **1509** (2015), no. 09 017, arXiv:1411.4172 [3],
- J. P. Conlon, M. C. D. Marsh, A. J. Powell, *Galaxy Cluster Thermal X-ray Spectra Constrain Axion-Like Particles*, arXiv:1509.0674 (To appear in Phys. Rev. D) [4].

Chapter 4 is based on the paper [1], written in collaboration with Stephen Angus, Joseph P. Conlon and Ulrich Haisch. My main contribution to this work was the computation of the anomalous dimension of the modulus' dominant coupling to the visible sector, Z. The wording of this chapter has been changed to reflect the current state of observations of dark radiation.

Chapter 5 is based on the papers [2, 3], the first of which was written in collaboration with Stephen Angus, Joseph P. Conlon, David Marsh and Lukas Witkowski. My contribution to the Coma cluster analysis [2] was implementing and running the simulation of the magnetic field, and the numerical calculation of the conversion probabilities. The follow up study of the A665, A2199, and A2255 galaxy clusters [3] is all my own work.

Chapter 6 is based on the paper [4], written in collaboration with Joseph P. Conlon and David Marsh. My contribution was calculating the conversion probabilities, and analysing these in order to get bounds on the ALP–photon inverse coupling M.

Other work not included in this thesis includes:

- J. P. Conlon, A. J. Powell, *A 3.55 keV line from DM $\to a \to \gamma$: predictions for cool-core and non-cool-core clusters*, JCAP **1501** (2015), no. 01 019, arXiv:1406.5518 [5],
- M. Berg, J. P. Conlon, F. Day, N. Jennings, S. Krippendorf, A. J. Powell and M. Rummel, *Searches for Axion-Like Particles with NGC1275: Observation of Spectral Modulations*, arXiv:1605.01043 [6],
- G. Grilli di Cortona, E. Hardy and A. J. Powell, *Dirac vs Majorana gauginos at a 100 TeV collider*, JHEP **1608** (2016), 014, arXiv:1606.07090 [7].

2 String Moduli and Cosmology

The universe is observed to be expanding. It naturally follows that the universe was much denser and hotter in the past. Thus the universe was once filled with a hot thermal bath of relativistic particles. The expansion of the universe caused the bath to cool, interaction rates to fall, and particles to become non-relativistic, until finally the universe became matter-dominated. The universe then became transparent to photons and the cosmic microwave background (CMB) was formed. Gravitational interactions then became important leading to the galaxies and galaxy clusters we see today.

Whilst we may naturally conclude that the further we look back in the universe's history, the hotter the thermal bath was, there is however no observational evidence for a universe hotter than a few MeV. This is the temperature at which light elements form (big bang nucleosynthesis, or BBN). Thus it is not necessary that the universe was always dominated by the hot thermal SM bath pre-BBN.

A period of inflation much earlier in the universe's history is thought to set the initial conditions for the subsequent evolution: homogeneity, low curvature, the seeds of structure formation, and super-horizon correlations. Such a period has certainly gained a lot of experimental evidence in favour of it [8], however the energy scale

of, and mechanism inducing inflation is unknown. It seems likely that inflation was driven by one or more scalar fields (the inflaton), whose potential energy dominated the energy density of the universe, leading to an exponential expansion, lasting around 60 e-folds, for a review, see [9].

After inflation ends the energy density stored in the inflaton must be transferred into the SM fields, since inflation dilutes the SM particles to almost zero energy density. Decays of the inflaton must generate the hot SM thermal bath. This is known as reheating, and as we have seen, the reheating temperature is only constrained to be $\gtrsim$ MeV. In this chapter I will briefly review the thermal history of the universe, putting emphasis on one cosmological observable—dark radiation—which we shall use to test string theory models later. I then show that string theory generically posits the existence of a multitude of scalar fields, moduli, whose properties will lead to a much different history of the universe prior to BBN then in the standard picture. Oscillating moduli fields can dominate the energy density of the universe leading to periods of non-relativistic matter domination. It is then the decay of these moduli not the inflaton that reheats the universe, and this idea will be of central importance to Chapters 4 and 5.

2.1 Dark Radiation

Let us start post-inflation and after reheating. At this point the universe is in a hot thermal equilibrium, dominated by relativistic particles, with temperatures at least above a few MeV. The energy density of a population of relativistic particles redshifts as $\rho_{\rm rad} \sim a^{-4}$, where a is the scale factor, and falls with temperature as,

$$\rho_{\rm rad} = \begin{cases} \frac{\pi^2}{30}\, g\, T^4 & \text{(bosons)}, \\ \frac{7}{8}\,\frac{\pi^2}{30}\, g\, T^4 & \text{(fermions)}, \end{cases} \tag{2.1}$$

where the difference is due to Bose/Fermi statistics, and g counts the degrees of freedom of the particle. In the early universe, if we assume chemical and thermal equilibrium, then all particles have the same temperature, and the total energy density

of the universe is $\rho_{\rm rad} = \frac{\pi^2}{30} g_\star T^4$, where

$$g_\star = \sum_{bosons} g_i + \frac{7}{8} \sum_{fermions} g_i, \tag{2.2}$$

with g_i the number of internal degrees of freedom for each particle species.

As the universe cools below the mass of a particle species, it becomes non-relativistic. The thermal bath is now no longer able to produce the particle, so the interactions such as $X + \overline{X} \rightleftharpoons Y + \overline{Y}$ go out of equilibrium, and the particle annihilates away, until the annihilation rate falls below the Hubble scale. Subsequently the typical time for a particle to annihilate is longer than the age of the universe at that time, so the number of that particle is essentially fixed, i.e. it 'freezes in'. The energy density in the annihilating particles is transferred to the thermal bath of standard model particles. Since the entropy per comoving volume, s, remains a constant with expansion, and $s \propto g\, a^3 T^3$, the temperature of the thermal bath evolves as $T \propto g_\star^{-1/3} a^{-1}$.

At temperatures of around a few MeV, the weak interaction rate becomes smaller than the current expansion rate. The cross-section for scattering by a massive gauge boson is $\sigma \propto \left(\frac{\alpha}{M_X^2} E\right)^2 \sim G^2 T^2$, where $G \sim \frac{\alpha}{M_X^2}$, thus the interaction rate is $\Gamma = \langle n \sigma v \rangle \sim G^2 T^5$. This interaction freezes out when $H \sim \Gamma$, and since $H^2 \sim \frac{T^4}{M_{pl}^2}$, this implies that for temperatures

$$T \lesssim \left(\frac{M_X}{100\ \text{GeV}}\right)^{4/3}\ \text{MeV}, \tag{2.3}$$

the interaction is no longer strong enough to keep particles in thermal equilibrium.

This has two important consequences: i) neutrino freeze-out and ii) big bang nucleosynthesis (BBN). Since neutrinos only interact via the weak interaction, as soon as the weak interaction freezes out, neutrinos decouple from the thermal bath. Soon afterwards electrons and positrons become non-relativistic and annihilate away. The energy is transferred into the thermal bath, which by now only consists of photons, but not into the neutrinos since they have decoupled. Thus the photon temperature gets a $\left(\frac{g_\star(\text{before})}{g_\star\ (\text{after})}\right)^{1/3} = \left(\frac{11/2}{2}\right)^{1/3}$ boost relative to the neutrinos. The total energy in

the universe is now $\rho = \frac{\pi^2}{30} g_\star T^4$, with T the photon temperature, and now

$$\begin{aligned} g_\star &= \sum_{bosons} g_i \left(\frac{T_i}{T}\right)^4 + \frac{7}{8} \sum_{fermions} g_i \left(\frac{T_i}{T}\right)^4, \\ g_\star &= 2 + \frac{7}{4} N_\nu \left(\frac{4}{11}\right)^{4/3}, \end{aligned} \tag{2.4}$$

with T_i the temperature of particle i.

Here N_ν counts the number of neutrino species with temperature $T_\nu = \left(\frac{4}{11}\right)^{1/3} T$. The standard model contains three neutrinos, however the SM prediction for N_ν is 3.046. This is because the decoupling of neutrinos is not instantaneous as we have assumed. In addition, electron neutrinos have an extra charged current interaction with the electrons in the thermal bath, and thus decouple slightly later than the muon and tau neutrinos, which only interact with the electrons via the neutral current interaction. These two effects result in the neutrinos getting partially heated by the electron-positron annihilation. This effect is conventionally absorbed into the parameter N_ν.

The parameter N_ν is very sensitive to new physics. Any new, light relativistic degrees of freedom produced either thermally or non-thermally in the early universe will contribute to $g_\star$ and will thus increase any measurement of N_ν. This extra relativistic energy density is called dark radiation (see e.g. [10] for a review). Including dark radiation $\rho_{\rm dark}$, the total energy density in radiation below the electron mass is,

$$\begin{aligned} \rho_{\rm rad} &= \rho_\gamma + \rho_\nu + \rho_{\rm dark}, \\ \rho_{\rm rad} &= \rho_\gamma \left(1 + \frac{7}{8} N_\nu \left(\frac{4}{11}\right)^{4/3}\right) + \rho_{\rm dark}, \\ \rho_{\rm rad} &= \rho_\gamma \left(1 + \frac{7}{8} N_{\rm eff} \left(\frac{4}{11}\right)^{4/3}\right). \end{aligned} \tag{2.5}$$

In the last line we have used the conventional parametrisation of absorbing dark radiation into the parameter $N_{\rm eff} = N_\nu + \Delta N_{\rm eff}$, and thus measurements of $N_{\rm eff}$ tell us about the additional amount of radiation density in the universe.

Dark radiation can be any new particle that remains relativistic until BBN and CMB decoupling times. It can either be in thermal contact with or decoupled from

the SM thermal bath. Light states produced from decays at reheating are a natural candidate for dark radiation, and we will discuss such a candidate, axion-like particles, in Chapter 3. Dark radiation is thus very interesting theoretically, as an absence would imply the particle responsible for reheating had no light hidden sector decay modes. Note that despite its parametrisation in terms of number of neutrinos, it can have nothing to do with these particles, and the number $\Delta N_{\rm eff}$ need not be an integer.

BBN takes place when the weak interaction rate keeping protons and neutrons in equilibrium, $p + e \rightleftharpoons n + \nu_e$ falls below the Hubble rate. Given that the particles were in equilibrium, the ratio of their number densities is $n_n/n_p \sim \mathrm{e}^{-(m_n - m_p)/T} \sim 1/6$, which is reduced to $n_n/n_p \sim 1/7$ due to neutron decay before light elements are formed. Helium can now be formed, and it continues to be formed until all the neutrons are inside nuclei. A small amount of lithium is also formed before the reaction rates for further nucleosynthesis also freeze-out. Since neutron-proton equilibrium is maintained until $\Gamma \sim H \sim g_\star^{1/2} T^2$, the number density n_n/n_p is very sensitive to additional light relativistic degrees of freedom present that will contribute to $g_\star$. Measurements of the amount of primordial Helium for instance allow a test for the SM prediction of N_ν in the early universe. For a review of BBN and its use in studying BSM physics, see for example [11].

Later on in the universe's history, since radiation redshifts as $\rho \sim a^{-4}$ and matter redshifts as $\rho \sim a^{-3}$, we pass the period of matter-radiation equality and now the universe becomes matter-dominated. Soon after, the rates for interactions such as $p + e \rightleftharpoons H$ fall out of equilibrium, and the electrons in the universe become captured on nuclei forming electrically neutral atoms. After this the mean free path for photons is increased, and at some point becomes larger than the Hubble distance. This is known as the last scattering surface or the time of decoupling. Photons now free stream, and can be observed at current times as the CMB. Now gravitational effects between the non-relativistic particles become important, and structure formation begins.

2.2 String Moduli

One of the main predictions of string theory is that space-time is ten-dimensional. The six unobserved dimensions must be compact, and very small in order to evade observational constraints. The compactification of ten-dimensional fields leads to a multitude of scalar fields in the observable four dimensions. I review this here.

The appearance of scalar fields happens for any number of compact dimensions. Let us start by considering the simple example of a massless scalar field in five dimensions, Φ. The field has the action $\int d^5x\ \partial_M \Phi^\dagger \partial^M \Phi$, where $M = 0, 1, ..., 4$. Now suppose the fifth dimension is compactified into a circle, such that we identify the points $x^4 = 2\pi n R$ for $n = 0, 1, ...$, with R the radius of the compact dimension. We can then decompose the field Φ into a part in the four space-time coordinates x^μ and modes in the compact dimension,

$$\Phi(x^M) = \sum_n \phi_n(x^\mu) e^{inx^4/R}. \tag{2.6}$$

Under this decomposition the action splits as $\int d^4x \int_0^{2\pi R} dx^4 \left(\partial_\mu \Phi \partial^\mu \Phi + \partial_4 \Phi \partial^4 \Phi\right)$, and upon integrating over the compact dimension, we find the four-dimensional action for the fields ϕ_n

$$\int d^4x \sum_n \left(\partial_\mu \phi_n^\dagger \partial^\mu \phi_n + \frac{n^2}{R^2} \phi_n^\dagger \phi_n \right). \tag{2.7}$$

The compactification has then produced an infinite tower of scalar fields in four dimensions, one of which is massless, the rest separated in mass by $1/R$.

The above was true because the five-dimensional space could be written $\mathcal{M}_5 = M_4 \times S^1$ where M_4 is four-dimensional Minkowski space, and S^1 is the circle. In this case the five-dimensional Laplacian splits as $\Delta_5 = \Box + \partial_4 \partial^4$ where $\Box = \partial_\mu \partial^\mu$ as usual. The mode which satisfied $\partial_4 \partial^4 \Phi = 0$ gave a massless particle in the four-dimensional spectrum, since the five-dimensional equation of motion for this mode reduces to $\Delta_5 \Phi = \Box \Phi = 0$, which is the equation of motion of a massless field in four dimensions.

If the metric is itself a dynamical field, then compactification gives a massless

scalar field in the four-dimensional theory. The five-dimensional metric splits into a four-dimensional metric $g_{\mu\nu}$, along with a vector field $g_{\mu 4} \sim A_\mu$ and a scalar, $g_{44} \sim \phi$, where all three are a function of the x^M. Then just like above we can decompose these fields into Fourier modes of the circular compact dimension, with the $n = 0$ mode giving a massless field, and an infinite tower of massive modes for $n > 0$. The metric of the five-dimensional space can be written

$$\mathrm{d}s^2 = g_{\mu\nu}\mathrm{d}x^\mu x^\nu + e^{2\Phi}\left(\mathrm{d}x^4 + A_\mu \mathrm{d}x^\mu\right), \tag{2.8}$$

then the vacuum expectation value of the $n = 0$ mode, ϕ, of field Φ determines the radius of the compact dimension, since

$$\int \mathrm{d}^5x \, \det|g^{(5)}| \longrightarrow 2\pi R \int \mathrm{d}^4x \, e^{\phi} \det|g^{(4)}|, \tag{2.9}$$

when integrated over the circle. This is a simple example of a modulus, whose vacuum expectation value characterises the size of the compact space.

In the simple example above the fields were split into Fourier modes of the circle, and the mode for which $\partial_4\partial^4\phi = 0$ gave a massless field. In higher dimensional spaces there are many more modes for which $\Delta_{\text{compact}}\phi = 0$, and thus there are many more massless fields. In string theory the compact space is usually a Calabi-Yau manifold, and the moduli space of such manifolds is very complex [12]. There can be hundreds of such modes, leading to hundreds of moduli in the four-dimensional spectrum. Just like in the simple example above, the vacuum expectation values will determine the size and shape of the compact space.[1]

In addition, there are several other massless fields in the ten-dimensional string spectrum. The dimensional reduction of these fields leads to more massless scalar fields in four-dimensions just like in the simple case of the massless scalar field in five dimensions. These fields are the string axions, or axion-like particles. We will discuss these fields more in the following chapter, Chapter 3.

[1]The moduli space is larger because one can continuously deform the manifold, whilst keeping the topology the same, in many different ways. In the example of the circle the only deformation corresponded to changing the radius.

Massless moduli transmit long-range forces [13], and are thus ruled out since constraints from fifth-force experiments imply moduli should have masses $\gtrsim \mathcal{O}(\text{meV})$ [14]. In addition we would like to have some control over the moduli, since they effectively characterise the size and shape of the extra dimensions. A massless field has no potential allowing the field to take any value, for instance allowing the compact dimensions to become infinite in size. Thus in an actual compactification, these fields should not remain massless. Giving a moduli field a mass is known as moduli stabilisation. We will not discuss moduli stabilisation here, it has been reviewed in many places, and instead refer the reader to [15,16] for reviews. Stabilisation of all moduli indeed seems possible, and it is possible to have moduli with masses hierarchically smaller than the string scale [17], making them of phenomenological interest. We will discuss such a model in Chapter 4.

2.3 Moduli in Cosmology

If moduli exist, then they can drastically change the evolution of the universe. There can be periods where the moduli dominate the energy density of the universe post-inflation. The moduli fields must then decay to the SM producing the hot thermal bath, instead of the inflaton. Here we review moduli domination and what can be learnt from their decays.

Moduli have gravitational couplings to all sectors of the theory, in addition the values of the moduli determine the coupling constants, thus if the energy density during inflation is sufficiently high, the minimum of the moduli potential can be shifted relative to the 'true' post-inflationary minimum [18]. This can happen through couplings like $\frac{\phi^2}{M_{pl}^2}V_{infl}$. This will mean that post-inflation, the modulus starts out with a vacuum expectation value which can be $\mathcal{O}(M_{pl})$ away from its post-inflationary minimum.

The equation of motion for the evolution of a scalar field in an expanding universe is

$$\ddot{\phi} + 3H(t)\dot{\phi} + V_{,\phi} = 0, \tag{2.10}$$

where $V_{,\phi} = \frac{\partial V}{\partial \phi}$. If we assume the post-inflationary potential has the form $V(\phi) = m_\phi^2 \phi^2$ then when $H \gg m_\phi$, the equation above describes an over-damped harmonic oscillator. Thus ϕ is stuck at the value it took during inflation.

This equation of motion above can be re-written

$$\begin{aligned} \frac{d}{dt}\left(\frac{1}{2}\dot{\phi}^2 + V(\phi)\right) &= -3H\dot{\phi}^2, \\ \frac{d\rho}{dt} &= -3H(\rho + p), \end{aligned} \tag{2.11}$$

where $\rho = \frac{1}{2}\dot{\phi}^2 + V(\phi)$ is the energy density, and $p = \frac{1}{2}\dot{\phi}^2 - V(\phi)$ is the pressure. When H drops below m_ϕ, the scalar field is no longer over-damped, and is free to oscillate around its minimum. On timescales $t \ll H^{-1}$, the energy density is a constant, $\rho = V_{max} = \frac{1}{2}m^2\phi_{max}^2$, and so $\phi = \phi_{max}\cos(mt)$. This results in a pressure

$$\begin{aligned} p = V_{max} - 2V(\phi) &= \frac{1}{2}m^2\phi_{max}^2(1 - 2\cos^2(\omega t)), \\ &= -\frac{1}{2}m^2\phi_{max}^2\cos(2mt). \end{aligned} \tag{2.12}$$

Thus on short timescales the pressure oscillates, but averaged over timescales $t > m^{-1}$ the pressure is rapidly oscillating and averages to zero $\langle p \rangle = 0$. Then on timescales $t \sim H^{-1}$ the energy density evolves, according to equations 2.11, as

$$\rho \propto a^{-3}. \tag{2.13}$$

Thus, an oscillating scalar field has an energy density which evolves in exactly the same way as non-relativistic matter [19].

Any radiation produced by the decay of the inflaton post-inflation will quickly redshift away compared to this oscillating scalar field. The decay products of the inflaton will be relativistic and thus redshift as a^{-4}, whereas the oscillating moduli field redshifts as a^{-3}. Thus the moduli field will come to dominate the energy density of the universe, and the details of the inflaton's decay modes become unimportant. Due to its gravitational couplings to the other fields, the modulus will be unstable, and will decay. This decay will produce relativistic particles, whose energy density will again redshift away quickly. If there are other moduli fields with longer lifetimes, then these

moduli fields will dominate the energy density of the universe later. Thus it is possible for the universe to undergo several successive stages of moduli (matter) domination followed by radiation-domination. This is perfectly consistent with observations as long as the final modulus to decay, does so into SM particles, and produces a thermal bath with a temperature of at least a few MeV, the temperature at which BBN begins. Thus, in string models it is generally expected that the moduli fields are responsible for reheating. Reheating is then almost completely decoupled from the details of the inflationary model, and so reheating can be discussed independently, without specifying an inflationary model.

Moduli typically have decay rates $\Gamma \sim \frac{1}{48\pi}\frac{m_\phi^3}{M_{pl}^2}$, and decay when $\Gamma \sim H$. If we assume for simplicity that the modulus decays only into SM particles, instantaneously, then the post-decay energy density of the universe is $\rho = M_{pl}^2\Gamma^2 \sim T_{rh}^4$, with T_{rh} the reheat temperature, given by

$$T_{rh} \sim \sqrt{\Gamma M_{pl}} \sim \mathcal{O}(1)\ \text{GeV}\left(\frac{m_\phi}{3\cdot 10^6\ \text{GeV}}\right)^{3/2}. \tag{2.14}$$

Importantly in deriving this formula we have assumed that the SM thermal bath quickly thermalises post decay. Whether this assumption holds was studied for example in [20], it was found that thermalisation is dominated by $2 \to 3$ inelastic scattering processes, with the thermalisation timescale

$$\tau_{therm} \sim \left(\frac{m_\Phi}{\alpha^3 T_{rh}^2}\right). \tag{2.15}$$

Number changing processes are necessary to obtain thermal number densities, and these processes are also very efficient at cooling the modulus decay products, allowing kinetic and chemical equilibrium to be quickly established. For our purposes, this thermalisation timescale will always be shorter than the Hubble time, thus we will continue to assume instantaneous thermalisation.

Moduli with masses less than ~ 50 TeV will decay during or after BBN has taken place, ruining the successful predictions for light element abundances. The measured primordial abundances put a tight constraint on the allowed variation of the entropy

density and baryon-to-photon ratio from BBN to now [11]. If a modulus decays and produces a large number of photons, this would lead to an unacceptable decrease in the baryon-to-photon ratio. Not only this, but a large number of energetic photons produced by modulus decay can photo-dissociate the light elements that have been created. If the moduli are stable on cosmological lifetimes, then their energy density will vastly exceed the currently observed total energy density. These problems together are called the 'Cosmological Moduli Problem' (CMP) [18, 21, 22], thus string moduli must have masses $m_\phi \gtrsim 50$ TeV.

Dark radiation is another constraint on moduli. It constrains the decay modes of the last modulus to decay. In addition to the decay to the SM particles, the modulus will decay to any light particles that are present in the theory, these include for example ALPs. These particles will contribute to dark radiation. Thus, bounds on dark radiation directly test the decay modes of the particle responsible for reheating. As we will see in Chapter 3, there are typically many light axion-like particles in string models, see e.g. [23], and thus a large amount of dark radiation can be produced in these models. We return to this constraint in Chapter 4 where we will look at predictions for dark radiation in a specific string compactification scenario.

3 Axion-like Particles

In this chapter I introduce the main topic of this thesis, axion-like particles (ALPs). To do so it is illustrative to first discuss the QCD axion. The QCD axion is a hypothesised particle which explains the lack of CP violation in the strong interactions. I then motivate ALPs from string theory, before discussing attempts to find them using their coupling to photons. I also summarise relevant bounds on this coupling from astrophysics and terrestrial experiments.

3.1 The QCD Axion

The QCD Lagrangian contains the parity (P) and charge-parity (CP) violating term

$$\mathcal{L} = \theta \frac{\alpha_s}{8\pi} G_a^{\mu\nu} \tilde{G}_{a,\mu\nu}, \tag{3.1}$$

where θ is a constant, and α_s is the strong coupling. The $G\tilde{G}$ term can be written as a total derivative of products of gauge fields. However, the gauge fields do not necessarily go to zero at infinity, and thus the term cannot be integrated away. The quark mass matrix $\mathcal{M}_q$ will in general be non-diagonal and complex. The rotation needed to form a real, diagonal mass matrix will include a chiral rotation of the

quark fields. This rotation gives a contribution of $\arg\det \mathcal{M}_q$ to θ through the chiral anomaly. This leads to the new CP-violating parameter $\overline{\theta} = \theta + \arg\det \mathcal{M}_q$, which is the sum of two contributions with very different physical origins. The term leads to a neutron electric dipole moment (EDM) proportional to $\overline{\theta}$, which has not been observed. The current upper limit on the neutron EDM is $|d_n| < 2.9 \times 10^{-26}\, e\,\mathrm{cm}$ [24], which leads to the constraint $\overline{\theta} \lesssim 10^{-10}$. Thus it seems there is an unnatural cancellation between two parameters that are not related. This is the strong CP problem, for a review see [25].

The most popular solution is the axion. A new U(1) 'PQ' (named after Peccei and Quinn) global chiral symmetry is introduced to the Lagrangian [26, 27], which is spontaneously broken by a PQ-charged scalar field getting a vacuum expectation value, $\Phi = (f_{PQ} + \psi)e^{i\phi}$. The Goldstone mode ϕ is the axion [28, 29], it is a pseudo-scalar and inherits a continuous shift symmetry from the U(1) PQ symmetry, $\phi \to \phi + \alpha$. The axion has an anomalous coupling to $G_{\mu\nu}\tilde{G}^{\mu\nu}$, and thus through the field redefinition

$$\phi \to \phi' = \phi - \overline{\theta}, \tag{3.2}$$

the $\overline{\theta}$ parameter is promoted to a dynamical field. QCD instantons induce a potential for the axion which has a minimum at $\phi' = 0$, the strong CP problem is thus solved dynamically.

The canonically normalised axion field is $a = f_{PQ}\phi$, and thus the coupling of the axion to other fields is suppressed by powers of $1/f_{PQ}$. In order to evade constraints, which we discuss in section 3.3, the scale of PQ-symmetry breaking f_{PQ} should be very high $f_{PQ} \gg m_{weak}$, and thus the axion is very weakly interacting. The shift symmetry also means the axion can only couple to other fields through derivative $\partial_\mu a$ couplings. The coupling to gluon fields is induced through the colour anomaly by either SM quarks carrying PQ charge, or new, heavy quarks, which carry colour and $U(1)_{PQ}$ charge.

The potential that sets $a = 0$ induces a mass for the axion. The neutral pion gets its mass from the same source and thus the axion and pion mix. The axion mass is

then related to the pion mass by $f_a m_a \sim f_\pi m_\pi$, which leads to the famous relationship $m_a \propto f_a^{-1}$, where $f_a = f_{PQ}/N$, and N is the colour anomaly. A full treatment leads to the relationship [30],

$$m_a \approx 6 \text{ meV} \left(\frac{10^9 \text{ GeV}}{f_a} \right). \tag{3.3}$$

The mixing with the pion leads to a model-independent coupling of the QCD axion to nucleons. There is also a model-dependent coupling if the SM quarks are PQ charged. The only other model-independent coupling of the axion is to photons. Again the mixing with the pion leads to the model-independent term $a\, F_{\mu\nu}\tilde{F}^{\mu\nu}$. If the heavy quarks that carry PQ charge also carry electric charge, or the axion couples to the SM fermions, then there will be an additional model-dependent contribution from the electromagnetic anomaly, leading to

$$\mathcal{L}_{a\gamma} = -\frac{1}{4} g_{a\gamma} a\, F_{\mu\nu}\tilde{F}^{\mu\nu}, \tag{3.4}$$

with $g_{a\gamma} = \frac{\alpha_{em}}{2\pi f_a}\xi$, $\xi = \frac{E}{N} - \xi_\pi$, with E the model-dependent electromagnetic anomaly, and ξ_π is the contribution from pion mixing.

In addition, the axion can have model dependent derivative couplings to SM fermion fields, provided the SM fields carry PQ charge,

$$\mathcal{L}_{af} = \frac{C_j}{2f_a} \overline{\Psi}_j \gamma^\mu \gamma_5 \Psi_j \partial_\mu a, \tag{3.5}$$

where C_j is dependent on the PQ charge of the fermion field.

Finally, let us note that the QCD axion can be dark matter [31]. In the early universe when QCD instanton effects are unimportant, the axion has no potential so can take any value between 0 and $2\pi f_a$. Thus, just like we saw for moduli, the axion will start out far away from the minimum of its potential when QCD effects become important. The axion field will then oscillate around its minimum, and as we showed in Chapter 2 this oscillating field will behave like dark matter. The oscillating axion field has energy density, see e.g. [30],

$$\Omega_a h^2 \approx 0.11 \left(\frac{f_a}{5 \times 10^{11} \text{ GeV}} \right)^{1.19}, \tag{3.6}$$

where we have assumed the axion is initially $\mathcal{O}(f_a)$ away from the minimum of the QCD potential. We have also assumed that the PQ-symmetry is broken before inflation, such that the axion takes a constant value throughout the observable universe. In the opposite case, the axion will take different initial values in different patches of the universe, and there will be an additional production mechanism of axions from cosmic strings. Thus for axion decay constants $f_a \sim 10^{12}$ GeV the axion can account for the observed dark matter.

3.2 String Axions and Axion-like Particles

We saw in section 2.2 that the four-dimensional effective theories arising from compactifications of string theory contain many scalar fields called moduli, and also many pseudo-scalar fields called string axions, coming from the compactification of massless fields. I now review the origin of these in slightly more detail. I again point the interested reader to the reviews [15, 16] for more information.

In the different string theories, there are massless p-form fields C_p in the ten-dimensional spectrum. These fields are higher-dimensional analogues of gauge fields. Decomposing these p-forms into four and six-dimensional mdoes and compactifying gives many massless scalar fields, just like in the case of a simple scalar field in five dimensions studied in section 2.2. We consider an example of the appearance of axions here using type IIB string theory. Consider, for example, the massless four-form, part of its decomposition into modes in the compact space includes the term

$$C_{4,10d} \supset \rho_a(x)\omega^a. \tag{3.7}$$

Here $\rho_a(x)$ is a function of the four space-tme coordinates, ω^a is a four-form for which $\Delta_{6d}\omega^a = 0$, Δ_{6d} is the six-dimensional Laplacian, and a is an index which runs over the number such four-forms. The fields ρ_a then obey the four-dimensional equation of motion for a massless field, they are called string axions.

The forms ω^a are dual to the four-cycles Σ_4^b of the compact geometry, $\int_{\Sigma_4^b} \omega^a \sim \delta_b^a$. Thus we can also think of axions as arising from the integral of the C_4 field over an

internal four-cycle

$$\rho_a(x) \sim \int_{\Sigma_a} C_4. \tag{3.8}$$

The size of these four-cycles is dynamical and corresponds to the moduli discussed earlier which control the size of the compact space, they are called the Kähler moduli. Thus it is easy to see that we get the same number of Kahler moduli as axions in the effective theory. In fact the moduli and axions combine into supersymmetry multiplets, and we can write them as a complex scalar $T_a = \tau_a + i\rho_a$, where τ_a is a modulus.

The p-form fields are invariant under gauge transformations in the higher dimensional theory, and the pseudo-scalar fields inherit this invariance as a shift symmetry. Thus the string axions are massless pseudo-scalars with a shift symmetry, similar to the QCD axion, hence their name. The typical decay constant of these axions is $f_a \sim M_{string}$, see e.g. [32].

Moduli stabilisation can also give high scale masses to some of the axions, removing them from the low energy spectrum. Generally though some of these axions can remain massless through the stabilisation process. However non-perturbative effects can give very small masses to these fields. It has been shown that generally many very light string axions can remain in the low energy spectrum, e.g. [23, 33, 34], this has been called the 'axiverse'.

Many different theories of BSM physics lead to similar fields in the low energy effective theory. Theories with higher-dimensional gauge symmetries, or theories with broken U(1) symmetries like the PQ symmetry lead to these axion-like fields. If such a theory contains many of these fields a_i, each with decay constant f_i, then the Lagrangian for these fields, including potential couplings to gluons and photons will be

$$\mathcal{L} = \frac{1}{2}\partial_\mu a_i \partial^\mu a_i + \frac{\alpha_s}{8\pi}\frac{C_{g,i}}{f_i} a_i G\tilde{G} + \frac{\alpha}{8\pi}\frac{C_{\gamma,i}}{f_i} a_i F\tilde{F} + ..., \tag{3.9}$$

where the ... includes any other couplings to hidden sector fields or SM fermions, and

the $C_{\alpha,i}$ are constants. The combination of fields

$$\frac{a_{QCD}}{f_a} = \sum_i \frac{C_{g,i}}{f_i} a_i, \tag{3.10}$$

solves the strong CP problem, and is the QCD axion discussed in section 3.1.

The orthogonal combinations of fields then do not couple to QCD, but will still couple to photons. They may pick up small masses from non-QCD non-perturbative physics. These fields are called axion-like particles (ALPs), they have the Lagrangian

$$\mathcal{L} = \frac{1}{2}\left(\partial_\mu a \partial^\mu a - m_a^2 a^2\right) + \frac{1}{4M} a F_{\mu\nu}\tilde{F}^{\mu\nu}, \tag{3.11}$$

where we have dropped indices labelling the ALPs. The inverse coupling M is defined in this way, how this parameter is related to the decay constant f_a is dependent on the model-dependent parameter C_γ. Importantly, there is no relation between M and m_a as there is for the QCD axion.

When we refer to ALPs in the rest of this thesis, we refer to a generic field which has the Lagrangian 3.11 above.

We note finally, that just like the QCD axion, ALPs can make up some or all of the dark matter by the same mechanism. The energy density of the oscillating ALP field is a function of the ALP mass as well as its decay constant, and also on the temperature dependence of the ALP mass.

3.3 Searching for ALPs

The coupling of an ALP to electromagnetism in equation 3.11 can be written

$$\mathcal{L}_{a\gamma\gamma} = \frac{1}{4M} a F_{\mu\nu}\tilde{F}^{\mu\nu} = \frac{1}{M} a \vec{E}\cdot\vec{B}. \tag{3.12}$$

This makes it clear that in the presence of either an external electric or magnetic field an ALP can convert into a photon and vice versa. It was realised early on that this coupling would allow searches for the QCD axion and ALPs [35, 36].

In the next section, 3.3.1, I review the theory of ALP–photon oscillations in an external magnetic field. I then review searches for ALPs, and constraints on the ALP–

photon coupling, both from astrophysical observations (section 3.3.2), and laboratory experiments (section 3.3.3). The main focus will be on the ALP–photon coupling, but where relevant I will also mention constraints on couplings to other particles, or properties of the QCD axion.

3.3.1 ALP–photon Mixing

The conversion of a photon into an ALP and vice versa in an external magnetic field will be of central importance to later chapters. Here we show how this conversion is possible, first worked out in [37]. Note that since the photon has spin 1 and the ALP has spin 0, it is only possible for this conversion to happen in a magnetic field which is transverse to the propagation of the ALP/photon.

Let us start with the photon and ALP Lagrangian,

$$\mathcal{L} = -\frac{1}{4}F_{\mu\nu}F^{\mu\nu} + \frac{1}{2}\left(\partial_\mu a \partial^\mu a - m_a^2 a^2\right) + \frac{1}{4M}\, a\, F_{\mu\nu}\tilde{F}^{\mu\nu} - J_\mu A^\mu, \tag{3.13}$$

where J_μ is the electromagnetic current. This Lagrangian gives the equations of motion

$$\begin{aligned} \partial_\mu F^{\mu\nu} &= J^\nu + \frac{1}{M}\partial_\mu a \tilde{F}^{\mu\nu}, \\ \partial_\mu \partial^\mu a + m_a^2 a &= -\frac{1}{4M} F_{\mu\nu}\tilde{F}^{\mu\nu}, \end{aligned} \tag{3.14}$$

along with the usual condition $\partial_\mu \tilde{F}^{\mu\nu} = 0$. Now we write the field strength tensor as a sum of the photon A^μ and the external field,

$$F^{\mu\nu} = F^{\mu\nu}_{\text{ext}} + \partial^\mu A^\nu - \partial^\nu A^\mu. \tag{3.15}$$

With the gauge-fixing condition $\partial_\mu A^\mu = 0$, and setting $A^0 = 0$, assuming the photon field is weaker than the external field we get the equations of motion

$$\begin{aligned} \partial_\mu \partial^\mu \vec{A} &= -\frac{\vec{B}_{\text{ext}}}{M}\frac{\partial a}{\partial t}\,, \\ \partial_\mu \partial^\mu a + m_a^2 a &= \frac{1}{M}\frac{\partial \vec{A}}{\partial t}\cdot \vec{B}_{\text{ext}}\,. \end{aligned} \tag{3.16}$$

We can decompose A into components perpendicular to and parallel to the external magnetic field, assuming plane wave forms travelling in the z-direction we get the

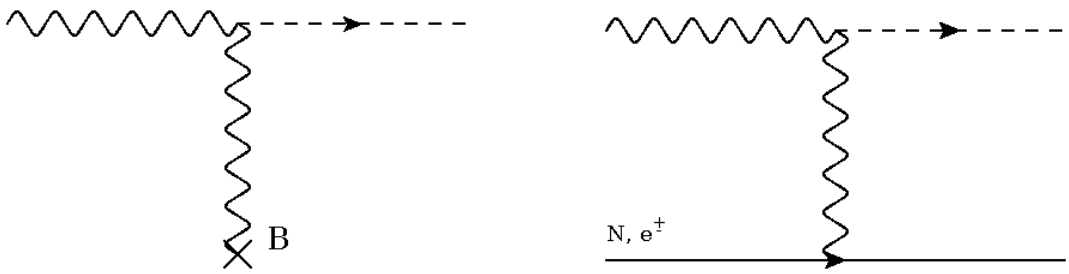

Figure 3.1: *Left*—Feynman diagram for the conversion of a photon into an ALP in an external magnetic field. *Right*—Feynman diagram for a photon scattering off the electric field of a charged nucleon or electron, producing an ALP.

coupled equations of motion

$$\left(\omega^2 + \partial_z^2 + \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & B_{\rm ext}\omega/M \\ 0 & B_{\rm ext}\omega/M & -m_a^2 \end{pmatrix}\right)\begin{pmatrix} A_\perp \\ A_\parallel \\ a \end{pmatrix} = 0. \tag{3.17}$$

We can simplify this by taking $\omega^2 + \partial_z^2 = (\omega + i\partial_z)(\omega - i\partial_z) \approx (\omega + k)(\omega - i\partial_z) \approx 2\omega(\omega - i\partial_z)$. Then we get the linearised equation of motion

$$\left(\omega - i\partial_z + \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & B_{\rm ext}/2M \\ 0 & B_{\rm ext}/2M & -m_a^2/2\omega \end{pmatrix}\right)\begin{pmatrix} A_\perp \\ A_\parallel \\ a \end{pmatrix} = 0. \tag{3.18}$$

If one starts with a pure ALP (or photon) state, then along the propagation, due to the off-diagonal mixing terms, the ALP (photon) wavefunctions develop non-zero photon (ALP) components. The conversion probability is then just the probability that the state will be measured as a photon (ALP). This process is shown in figure 3.1 (left).

There is a similar process, depicted in figure 3.1 (right), known as the Primakoff process. In this case, a photon can scatter off the electric field of an electron or nucleus producing an ALP in the final state.

In both cases the mis-match between ALP and photon masses leads to a momentum transfer, either to the magnetic field, or to the charged particle in the scattering. When propagating in a plasma, the photon dispersion relation is modified to $\omega^2 = k^2 + \omega_{\rm pl}^2$, and the photon develops an effective mass of $\omega_{\rm pl}$. The conversion is

suppressed for large momentum transfer. This means that conversion in a plasma is most efficient for ALPs with a mass similar to or less than the plasma frequency, $m_a^2 \lesssim \omega_{\rm pl}^2$, and highly suppressed for larger ALP masses. The result is that its much easier to search for lighter ALPs than heavier ALPs, but as discussed in section 3.2 this is a very interesting region for BSM physics and string compactification models.

3.3.2 Astrophysics

In this section I will discuss the very stringent astrophysical bounds on various parts of the axion and ALP parameter space.

Stars have been used to derive axion and ALP bounds. Photons in the stellar interior would convert into axions or ALPs by the Primakoff process. This would be a very efficient energy loss mechanism for the star. The energy carried away by ALPs or axions should not exceed the luminosity of the star. The constraint from ALP production in the solar interior is $M > 4 \times 10^8$ GeV [38].

It was noted early on that the helium burning horizontal branch stars are a more promising place to look [39]. ALP production would dramatically shorten the time stars exist in a stable helium burning phase, reducing the number of these stars relative to other populations. The study [40] obtained the lower bound $M > 1.5 \times 10^{10}$ GeV at 95% confidence. In addition they found a slight $1 - 2\sigma$ preference for an ALP with $M = 1.5 \times 10^{10}$ GeV. This constraint holds for ALPs with masses $m_a \lesssim$ keV.

In addition to this hint, there are also hints for ALPs from the cooling of white-dwarfs. It appears that the cooling of white dwarves exhibit small anomalies which can be explained if the ALP couples to electrons acting as an additional cooling mechanism. It is even possible to explain both of these hints by the same coupling to electrons. See [41] for a recent review of cooling hints from stars and their explanation in terms of ALPs.

The supernova SN1987a leads to a very strong bound on the QCD axion and ALPs. SN1987a is a core collapse supernova, located in the Large Magellanic Cloud. At the end of its life, the iron core collapsed until it reached nucleon densities forming a

neutron star. This caused a 'bounce', creating a shock wave which produced the visible supernova explosion. The resulting proto-neutron star initially had a temperature in the tens of MeV. The core of the supernova is opaque to photons, but not to the weakly interacting neutrinos who efficiently carry energy away from the core. The observations of a burst ($\sim 5-10$ seconds) of neutrinos by Kamiokande [42] and the Irvine-Michigan-Brookhaven water Cerenkov detector [43] provide valuable insight into core collapse supernovae.

The scattering of photons off the electric fields of charged particles in the proto-neutron star would produce a large amount of ALPs, which, like neutrinos, can freely leave the core. On their way to earth, these gamma ray energy ALPs can re-convert to photons in the Milky Way's magnetic field, leading to a short burst of gamma rays at earth coincident with the neutrino burst. Such a burst was not observed [44], allowing a constraint to be put on the ALP-photon coupling. Early studies produced bounds $M \gtrsim 1-3\times 10^{11}$ GeV, depending on the model of the supernova, and the Milky Way's magnetic field [45,46]. This was recently updated to $M > 2\times 10^{11}$ GeV, for $m_a \lesssim 10^{-10}$ eV, by [47], using updated models of core collapse supernova, and modern models for the Milky Way's magnetic field. This is the most stringent bound on low mass ALPs, and will be relevant for us later.

These ALPs, and also the QCD axion, will lead to an additional source of energy loss for the supernova. The required coupling to photons is too small to produce large effects, however if ALPs or the axion couples to nucleons, then they efficiently carry energy away from the core. An ALP coupling to nucleons is model-dependent, however the QCD axion will get a model-independent coupling due to mixing with the neutral pion. The original bound on the axion was determined to be $f_a \gtrsim$ few $\times 10^{10}$ GeV [48–52]. This was later relaxed somewhat to $f_a \gtrsim 7\times 10^9$ GeV [53,54]. This remains the most stringent lower bound on the axion decay constant. Note the precise value of the bound is uncertain due to model-dependent $\mathcal{O}(1)$ parameters.

Photon-ALP mixing in magnetic fields leads to energy-dependent distortions to astrophysical spectra. The H.E.S.S. satellite study [55] looked at the spectrum of the

active galaxy PKS 2155-304 to derive bounds on the ALP–photon coupling. They derive the bound $M > 5 \times 10^{10}$ GeV, at 95% confidence, for the small ALP mass window 1×10^{-8} eV $\lesssim m_a \lesssim 6 \times 10^{-8}$ eV. The bound is relevant for only these masses because to get these spectral distortions in the right energy window, 'resonance' is required between the ALP mass and photon mass in the converting medium. This bound is more constraining than the He-burning stars bound, and also in a mass region not covered by the bounds from SN1987a γ-rays. We will study a similar phenomenon for a larger mass range in Chapter 6.

It has been claimed that the universe is more transparent to very high energy gamma rays than expected [56–58]. The situation is not settled however since another analysis claims to find no such anomaly [59]. A possible explanation of this effect is if the gamma rays can oscillate into an ALP and then back into a photon. Very high energy gamma rays should be attenuated due to absorption on the extragalactic-background light (EBL)—the light emitted by stars and galaxies—which would pair produce electrons. Allowing photons to oscillate into ALPs reduces the probability that they will scatter off the EBL, increasing the universe's transparency. The oscillation can either happen as the photon travels through the intergalactic medium [60–62], and thus takes place in the unknown inter-galactic magnetic field (IGMF), or in the source AGN or host galaxy cluster (with the back-conversion happening in the galactic magnetic field) [63, 64], or a mixture of both [65, 66]. If the conversion takes place in the IGMF then the changes to AGN spectra will be redshift-dependent, otherwise, the spectra of all AGNs, or just AGNs in galaxy clusters, would be altered. The detailed analysis of [62] found an ALP with coupling to photons $M \sim 1 \times 10^{11}$ GeV fit observations. Couplings smaller than this were excluded since they should produce even larger effects. These bounds are valid for ALP masses $m_a \lesssim 10^{-7}$ eV, and thus partly overlaps with the SN1987a gamma ray burst bound.

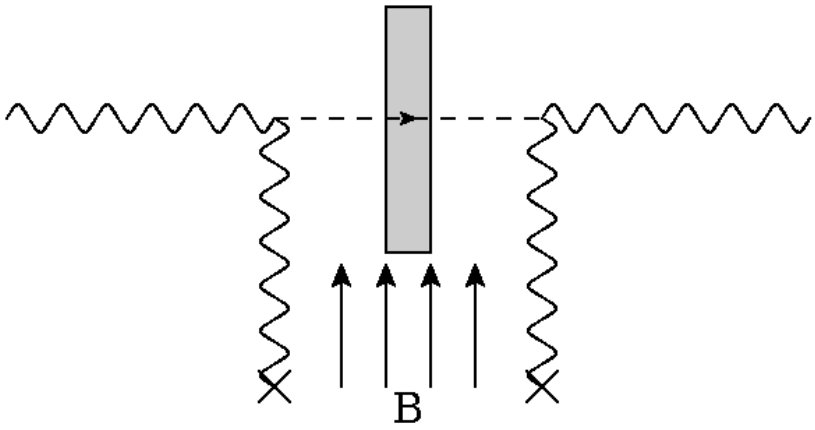

Figure 3.2: A light-shining-through-a-wall experiment. A laser produces a photon beam which travels through a magnetic field, some of these photons will transform into ALPs which can travel freely through a 'wall', unlike the photons. The ALP can then convert back into a photon using the magnetic field on the other side of the wall.

3.3.3 Terrestrial Experiments

There are two main experimental techniques used to search for ALPs on earth. These are light-shining-through-a-wall (LSW) and helioscope experiments.

In the first of these, an ALP gives a non-zero probability that light can pass through a wall in the presence of a magnetic field. The photon may oscillate into an ALP one side of the wall, and instead of getting reflected, travels through the wall to the other side where it may oscillate back into a photon to be detected. This process is illustrated schematically in figure 3.2.

The experiment OSQAR has the current best bound for an LSW experiment of $M \gtrsim 2.9 \times 10^7$ GeV at 95% confidence, for masses $m_a \lesssim 2 \times 10^{-4}$ eV [67]. A significant advance on this bound is expected to come in the next few years from the ALPS-II (Any Light Particle Search) LSW experiment. ALPS-II is an upgrade on the ALPS-I experiment [68], which yielded competitive bounds to OSQAR, by several orders of magnitude. It is hoped that ALPS-II will probe values of the coupling up to $M \sim 5 \times 10^{10}$ GeV [69], and will thus push into so far uncovered regions in the mass range $m_a \sim 10^{-9} - 10^{-4}$ eV. Data taking is expected to start in 2018/19.

The second type of search is a helioscope experiment. This uses the potential large flux of ALPs produced in the solar interior. The scattering of photons in the solar

interior can produce ALPs by the Primakoff process, which freely stream away from the interior and out of the sun. The ALPs produced are all of keV energies. The luminosity of ALPs produced, and the flux at earth is [70]

$$\begin{aligned} L_a &= \left(\frac{10^{10}\ \text{GeV}}{M}\right)^2 1.85 \times 10^{-3} L_\odot, \\ \Phi_a &= \left(\frac{10^{10}\ \text{GeV}}{M}\right)^2 3.75 \times 10^{11}\ \text{cm}^{-2}\text{s}^{-1}. \end{aligned} \tag{3.19}$$

Helioscope experiments use a large magnet to re-convert these ALPs into X-ray photons. The best bounds from this type of experiment come from CAST (CERN Solar Axion Telescope). The lack of any observation during the running of CAST leads to the lower limit on M of [71]

$$M > 1.1 \times 10^{10}\ \text{GeV}, \tag{3.20}$$

for ALP masses $m_a \lesssim 10^{-2}$ eV. This bound is comparable in magnitude to the bound for He-burning horizontal branch stars, and is the most stringent laboratory bound.

There are proposals for the next generation helioscope experiment IAXO (International Axion Observatory), which is still at the technical design stage. The proposed IAXO experiment will improve on the CAST bound by using a stronger and bigger magnetic field. The sensitivity of the instrument should probe up to [72]

$$M \sim 1 - 5 \times 10^{11}\ \text{GeV}, \tag{3.21}$$

for ALP masses up to $m_a \lesssim 0.25$ eV, far surpassing the previous bound from CAST by an order of magnitude. IAXO should also be able to probe ALP–photon couplings beyond the bound from the lack of a SN1987a gamma ray burst, and will also probe the region needed to explain the anomalous gamma ray transparency of the universe.

Other types of terrestrial searches focus on ALPs or axions as a dark matter candidate. Examples of these are ADMX and ADMX-HF [73, 74] which look for conversion of dark matter axions or ALPs in a microwave cavity, and also CASPEr [75] which aims to detect the oscillating axion or ALP dark matter field which should induce an oscillating neutron electric dipole moment. Finally, independent of the

axion or ALP being dark matter, there is a proposed experiment to search for the spin-dependent force induced by the QCD axion, which will probe the QCD axion couplings to fermions [76].

4 ALP Dark Radiation Production in Large Volume Scenarios

In this chapter I use early universe cosmology to study a class of string theory compactifications called the Large Volume Scenario (LVS). A generic prediction of these models is a hierarchical moduli spectrum, and at least one (nearly) massless ALP, which is the partner of the lightest modulus. This makes it a very tractable model to deal with, the lightest modulus is responsible for reheating, and computing its decay rates to various products gives a predictive model for dark radiation. The modulus will decay with a sizeable branching fraction to its ALP partner. I firstly briefly review the tree level dark radiation prediction in these models, which was first done in [77, 78]. I then look at loop corrections to this prediction. There are potentially large effects of running the coupling of the modulus to the visible sector from the string scale to the mass of modulus. These may suppress or enhance dark radiation production in these models, allowing them to evade the tight constraints. This work was published in [1], and was done in collaboration with Stephen Angus, Joseph P. Conlon and Ulrich Haisch.

4.1 Dark Radiation

The possible existence of dark radiation is interesting from both a theoretical and observational perspective. Observationally, dark radiation refers to additional radiation density beyond that predicted in the ΛCDM model of standard Big Bang cosmology. As discussed in Chapter 2, at the time of big bang nucleosynthesis (BBN) the Standard Model (SM) predicts radiation density only in photons and three neutrino species. The energy density of possible dark radiation is conventionally parametrised by the effective excess number of neutrino species, $\Delta N_{\rm eff} \equiv N_{\rm eff} - 3.046$, then

$$\rho_{rad} = \rho_\gamma \left(1 + \frac{7}{8} N_{\rm eff} \left(\frac{4}{11}\right)^{4/3}\right). \tag{4.1}$$

Cosmic microwave background (CMB) experiments have developed increasing sensitivity to $\Delta N_{\rm eff}$. For a long time CMB experiments saw hints of a non-zero $\Delta N_{\rm eff}$, although the situation was confused due to the slight tension between local Hubble Space Telescope (HST) measurement of the Hubble constant H_0 [79] and the CMB derived value. The degeneracy between $N_{\rm eff}$ and H_0 meant the tension could be resolved by allowing larger $N_{\rm eff}$. When combined with the HST measurement, the results from the Wilkinson Microwave Anisotropy Probe (WMAP), the Atacama Cosmology Telescope (ACT), and the South Pole Telescope (SPT) are $N_{\rm eff} = 3.84 \pm 0.40$ [80], $N_{\rm eff} = 3.50 \pm 0.42$ [81], $N_{\rm eff} = 3.71 \pm 0.35$ [82], respectively.[1] However, the Planck satellite [83] has now measured the CMB with the greatest ever precision, resulting in a measurement of $N_{\rm eff} = 3.13 \pm 0.32$ at 68% confidence [8]. The constraint is even more stringent if polarisation data is used, and becomes $N_{\rm eff} = 2.99 \pm 0.20$ [8]. This leads to the conclusion that at the 3σ level, $\Delta N_{\rm eff} > 1$ is excluded.[2]

These values are quite consistent with the independent measurements of $N_{\rm eff}$ from BBN measurements alone [30]. Combining the latest Planck results, with the latest

[1] Without including direct measurements of H_0, the determinations using only CMB and baryon acoustic oscillations data are $N_{\rm eff} = 3.55 \pm 0.60$ (WMAP), $N_{\rm eff} = 2.87 \pm 0.60$ (ACT) and $N_{\rm eff} = 3.50 \pm 0.47$ (SPT).

[2] However, there is still some $2 - 3\sigma$ tension between the measurement of the Hubble constant by Planck using the CMB, and 'local' astrophysical measurements using standard candles [84]. This tension can be alleviated if $N_{\rm eff}$ is higher than in the standard model $N_{\rm eff} \approx 3.5 - 4$.

deuterium [85] and helium [86] abundances, gives the very stringent 95% C.L. constraints $N_{\rm eff} = 2.95^{+0.52}_{-0.52}$ and $N_{\rm eff} = 3.11^{+0.59}_{-0.57}$ respectively [8]. Thus we take values of $\Delta N_{\rm eff} \gtrsim 1$ to be ruled out, and values $\Delta N_{\rm eff} \sim 0.5$ to be disfavoured.

Dark radiation is also interesting theoretically, as it is a simple and natural extension of ΛCDM. It is believed that after inflation the Universe was reheated from the decays of a scalar field. Dark radiation is produced whenever this field has a non-zero branching ratio to light hidden-sector particles. Examples of such particles include axion-like particles. From this perspective, it is not a presence, but an absence of dark radiation that would be a surprise—dark radiation is only absent if the reheating field has no decay modes to light hidden-sector particles.

Dark radiation also provides an arena to make contact between observations and models of Planck-scale physics. In particular, it allows us to make connections to string theory models. As discussed in Chapter 2, in theories with moduli it is expected that the lightest modulus field is responsible for reheating. Thus given a string theory model with lightest modulus Φ, it is possible to compute the prediction for dark radiation production given Φs decay rates to light, hidden sectors, compared to the visible sector. In this Chapter we do exactly this. We look at a promising compactification scenario of type IIB string theory, a sequestered form of the Large Volume Scenario (LVS) [87,88], and compute the amount of dark radiation produced in a minimal version of this model, including relevant radiative corrections.

4.2 The Large Volume Scenario

In the Large Volume Scenario (LVS) [87,88] the volume $\mathcal{V}$ of the compact dimensions is stabilised at exponentially large values. This stabilisation mechanism creates a naturally small expansion parameter, i.e. the inverse volume, and leads to a distinctive hierarchy of scales, given in the sequestered LVS by [89]

$$M_{\rm string} \sim \frac{M_{pl}}{\mathcal{V}^{1/2}}\,, \qquad m_\Phi \sim \frac{M_{pl}}{\mathcal{V}^{3/2}}\,, \qquad M_{\rm soft} \sim \frac{M_{pl}}{\mathcal{V}^{2}}\,, \qquad m_a \lesssim M_{pl}\, e^{-2\pi\mathcal{V}^{2/3}} \sim 0\,. \tag{4.2}$$

Here $M_{pl} = 2.4 \times 10^{18}\,\text{GeV}$ is the reduced Planck mass, while M_{string}, m_Φ, M_{soft} and m_a denote the string scale, the mass of the volume modulus Φ (the modulus that controls the overall volume of the extra dimensions), the scale of the SUSY breaking soft masses and the mass of the volume ALP a, respectively. In what follows, we will assume that the level of volume sequestering is the same for both scalar and gaugino masses, so that $M_{\text{soft}} \sim m_0 \sim m_{1/2}$. To solve the gauge hierarchy problem, i.e. $M_{\text{soft}} \sim 1\,\text{TeV}$, one needs $\mathcal{V} \sim 5 \times 10^7$, resulting in $M_{\text{string}} \sim 3 \times 10^{14}\,\text{GeV}$ and $m_\Phi \sim 7 \times 10^6\,\text{GeV}$.

The LVS is tractable to analyse as it has a unique lightest modulus, the volume modulus Φ, which is parametrically lighter than any other modulus. As argued in Chapter 2, the presence of a single lightest Planck-coupled modulus implies that within these models reheating should be driven by decays of the volume modulus, independently of the details of the high-scale inflationary model. Here we work with an MSSM matter content with low energy supersymmetry so as to solve the electroweak hierarchy problem.

This is an attractive model to study, as the majority of the volume modulus' couplings are calculable in a model-independent fashion. In fact, there are two important couplings. The first is to the volume ALP a, which is the ALP partner of the volume modulus, it is a hidden-sector state and thus the corresponding decay channel $\Phi \to a\,a$ gives rise to dark radiation. The second coupling is to the bilinear $H_u H_d$ of Higgs fields. This interaction leads to the only competitive visible-sector decay mode, $\Phi \to H_u H_d$, and induces the reheating of the Universe. The corresponding coupling Z is an undetermined constant with a natural value of $\mathcal{O}(1)$ at the string scale M_{string}. However, if the Higgs sector has an exact shift symmetry (see [90, 91] for explicit string theory constructions of such a symmetry), then Z is fixed to 1 at M_{string}. The case of a shift-symmetric Higgs sector with pure MSSM matter content is then completely defined and predictive. We will refer to this specific LVS as the minimal LVS (MLVS).

In section 4.3, I review the tree level computation of dark radiation in this model that was performed in [77, 78] (see also [92]). These authors found that in the

MLVS, dark radiation was over-produced, with $\Delta N_{\rm eff} \approx 1.7$. However the coupling Z, defined at the string scale, receives potentially large corrections of the form $1/(4\pi)\ln(M_{\rm string}/m_\Phi)$, which given the hierarchy $M_{\rm string} \gg m_\Phi$ could be $\mathcal{O}(1)$ and thus need to be resummed using renormalisation group (RG) techniques. The rest of the chapter presents this calculation. We find that loop corrections to the MLVS do not improve the picture. We scan over MSUGRA parameter space, and find in general that the coupling to the visible sector is actually reduced. Even for the most optimistic set of parameters, it is not possible to reduce $\Delta N_{\rm eff}$ sufficiently to evade current bounds within the MLVS.

4.3 Tree Level Dark Radiation Prediction

4.3.1 Decay Rates

Let us compute the branching fractions to the various sectors. To leading order, the volume of the compactified space can be given by the volume modulus $\tau_b = T_b + \overline{T}_b$ as $\mathcal{V} \approx \tau_b^{3/2}$. The volume modulus has a partner volume ALP, given by $a_b = -i(T_b - \overline{T}_b)$. The Kähler potential of the moduli fields is

$$K = -2\ln\mathcal{V} = -3\ln\left(T_b + \overline{T}_b\right), \tag{4.3}$$

where we have set $M_{pl} = 1$ for now.

The volume modulus coupling to its ALP partner comes about through the kinetic terms. The kinetic terms are $\mathcal{L} = -K_{i\bar{j}}\partial_\mu\Phi^i\partial^\mu\Phi^{\bar{j}}$ where K_i is the derivative of the Kähler potential with respect to the field Φ_i. Given the leading order Kähler potential above, we get

$$\mathcal{L} = \frac{3}{4\tau_b}\left(\partial_\mu\tau_b\partial^\mu\tau_b + \partial_\mu a_b\partial^\mu a_b\right). \tag{4.4}$$

After canonically normalising both fields, $\tau_b = \sqrt{\frac{3}{2}}\ln\Phi$ and $a_b = \sqrt{\frac{2}{3}}a$, we get the kinetic terms

$$\mathcal{L} = \frac{1}{2}\partial_\mu\Phi\partial^\mu\Phi + \frac{1}{2}\exp\left(-\sqrt{\frac{8}{3}}\Phi\right)\partial_\mu a\,\partial^\mu a, \tag{4.5}$$

giving the interaction

$$\mathcal{L} = -\sqrt{\frac{2}{3}}\frac{\Phi}{M_{pl}}\partial_\mu a \partial^\mu a, \tag{4.6}$$

where we have put M_{pl} back in. This leads to the decay of the volume modulus into the volume ALP, with width

$$\Gamma = \frac{1}{48\pi}\frac{m_\Phi^3}{M_{pl}^2}, \tag{4.7}$$

which will correspond to the production of dark radiation. In principle there can be many more light ALPs in the theory, if so the volume modulus would in general decay into all of these. In the MLVS we only consider there to be one light ALP, the model-independent volume ALP, but we note with the inclusion of these model-dependent ALPs, the amount of dark radiation produced by the volume modulus' decay would likely increase.

We must now compute the rate of decay of Φ into visible sector fields, this must be at a rate comparable to or larger than equation 4.7 in order to reheat the visible sector, and not overproduce dark radiation at the same time. The volume modulus couples to gauge bosons, through its contribution to the gauge kinetic function $f_a(\Phi_i)$ where Φ_i are the moduli fields. In the MLVS however such a coupling only comes about at one loop in the SM couplings, and the decay rate is thus α_{SM}^2 suppressed with respect to equation 4.7.

Next consider the couplings to MSSM scalar fields. The Kähler potential, when extended to include these fields, takes the form

$$K = -3\ln(T_b + \overline{T}_b) + \frac{C\overline{C}}{T_b + \overline{T_b}}, \tag{4.8}$$

which after canonically normalising the kinetic terms, gives the interactions

$$\mathcal{L} = \frac{1}{\sqrt{6}}\Phi(\overline{C}\Box C + C\Box\overline{C}). \tag{4.9}$$

The decay rate induced by this coupling is $\Gamma \propto m_C^2 m_\Phi \ll m_\Phi^3$ and is thus suppressed compared to the decay rate to ALPs. The decay rate to fermions has the same dependence, and so is again suppressed.

However, the quantum numbers of the two Higgs doublets allow us to write an

extra term in the matter part of the Kähler potential, the so-called Giudice-Masiero term [93],

$$K = -3\ln(T_b + \overline{T}_b) + \frac{H_u\overline{H}_u + H_d\overline{H}_d + (ZH_uH_d + h.c.)}{T_b + \overline{T}_b}, \tag{4.10}$$

which gives the extra interaction term

$$\mathcal{L} = \frac{1}{\sqrt{6}M_{pl}}\left(ZH_uH_d\Box\Phi + h.c.\right), \tag{4.11}$$

where again we have put M_{pl} back in. This leads to the unsuppressed decay of Φ into the Higgs fields, with decay rate

$$\Gamma = \frac{2Z^2}{48\pi}\frac{m_\Phi^3}{M_{pl}^2}. \tag{4.12}$$

This decay into Higgs fields is the only competitive decay mode to that of dark radiation.

4.3.2 Dark Radiation Prediction

Let us now compute ΔN_{eff} for this model. First, lets define

$$B_a = \frac{\Gamma_{\Phi\to aa}}{\Gamma_{tot}} = \frac{1}{1+2Z^2}, \tag{4.13}$$

as the branching fraction to dark radiation (ALPs). The ratio of dark radiation to SM radiation does not stay constant due to small boosts in the SM temperature due to particle decoupling. The value of N_{eff} then does not change from neutrino decoupling onwards, since the neutrino sector now no longer gets these boosts. During the expansion of the universe the comoving entropy, $s = g(T)a^3T^3$ remains constant, so the temperature of the SM obeys

$$g(T_{\nu,dec})^{1/3}\, a\, T_{\nu,dec} = g(T_{rh})^{1/3}\, a\, T_{rh} \tag{4.14}$$

with T_{rh} the reheat temperature, defined in Chapter 2. This tells us then that, remembering $\rho \sim gT^4$,

$$\frac{\rho_{dark,\nu}}{\rho_{SM,\nu}} = \left(\frac{g(T_\nu)}{g(T_{rh})}\right)^{1/3} \frac{\rho_{dark,init}}{\rho_{SM,init}} = \left(\frac{g(T_{\nu,dec})}{g(T_{rh})}\right)^{1/3} \frac{B_a}{1-B_a}, \tag{4.15}$$

where the subscript ν means the quantity is evaluated at neutrino decoupling. Then using $N_{\text{eff}} = N_\nu(1 + \rho_{dark}/\rho_\nu)$ and that $\rho_\nu/\rho_{SM} = g_\nu/g_{SM} = 21/43$ at neutrino decoupling. Then (setting $N_\nu = 3$ for convenience),

$$\Delta N_{eff} = 3\frac{\rho_{dark}}{\rho_\nu} = \frac{43}{7}\frac{\rho_{dark}}{\rho_{SM}} = \frac{43}{7}\left(\frac{g(T_{\nu,dec})}{g(T_{rh})}\right)^{1/3} \frac{B_a}{1-B_a}. \tag{4.16}$$

At neutrino decoupling, $g(T_{\nu,dec}) = 10.75$, and $g(T_{rh}) = 247/4 - 345/4$ depending on the exact reheat temperature. This gives the prediction for ΔN_{eff} in this model to be

$$\frac{1.56}{Z^2} \leq \Delta N_{\text{eff}} \leq \frac{1.74}{Z^2}. \tag{4.17}$$

At tree level in the MLVS one has $Z(m_\Phi) = Z(M_{\text{string}}) = 1$, which implies $\Delta N_{\text{eff}} \simeq 1.7$. However the measured values of N_{eff} require $\Delta N_{\text{eff}} \lesssim 0.5$, which translates into $Z^2 \gtrsim 4$. The MLVS tree-level prediction for ΔN_{eff} is hence in conflict with observation.

However, even if the Higgs sector is exactly shift symmetric at the compactification scale such that $Z = 1$, this symmetry is broken by the gauge and Yukawa couplings. In consequence, the coupling Z will receive logarithmically-enhanced corrections from MSSM loop diagrams. An immediate question then arises as to whether the radiative corrections are large enough to make the MLVS compatible with the measurements of N_{eff}. The next subsections will address this question.

4.4 Analytic Results

We have shown that the MLVS at tree level produces dark radiation a factor of a few above the current upper bounds. In this section we will derive some analytic results for the running of the coupling Z, to determine whether radiative corrections can enhance the modulus' branching fraction to the visible sector.

4.4.1 Running of Volume Modulus Higgs Coupling

In contrast to the Φaa coupling, which receives only Planck-suppressed corrections, the $\Phi H_u H_d$ coupling is modified by the virtual exchange of MSSM particles. As a result the coupling Z will evolve logarithmically from M_{string} to m_Φ, where the volume modulus decays. The scale dependence of Z is determined by the following RG equation

$$\frac{d}{dt}Z = \gamma_Z Z\,, \tag{4.18}$$

where $t \equiv \ln(Q/Q_0)$ with Q denoting the renormalisation scale and Q_0 a reference scale, and γ_Z is the corresponding anomalous dimension.

As a consequence of the supersymmetric non-renormalisation theorem [94,95], and since the volume modulus field itself does not renormalise, the anomalous dimension γ_Z can be written in terms of Higgs wave-function renormalisations. Thus

$$\gamma_Z = \gamma_{H_u} + \gamma_{H_d}\,, \tag{4.19}$$

where γ_{H_u} and γ_{H_d} are the anomalous dimensions of the Higgs superfields. To verify the correctness of (4.19), we have calculated the one-loop correction $\gamma_Z^{(1)}$ to the anomalous dimension γ_Z explicitly. The corresponding Feynman diagrams are depicted in figure 4.1. We performed the calculation of the self-energy and vertex diagrams using dimensional regularisation with modified minimal subtraction (i.e. the $\overline{\text{DR}}$ scheme). Our results for the individual diagrams agree with those given in [96–98]. Keeping only the third-family Yukawa couplings $y_{t,b,\tau}$, we obtain

$$\gamma_Z^{(1)} = \frac{1}{(4\pi)^2}\left[-\frac{3g_1^2}{5} - 3g_2^2 + 3|y_t|^2 + 3|y_b|^2 + |y_\tau|^2\right], \tag{4.20}$$

which equals the sum $\gamma_{H_u}^{(1)} + \gamma_{H_d}^{(1)}$ of one-loop superfield anomalous dimensions as given e.g. in the review [99]. Here the couplings g_1 and g_2 are given in terms of the conventional $U(1)_Y$ and $SU(2)_L$ SM gauge couplings by $g_1 = \sqrt{5/3}\,g'$ and $g_2 = g$, respectively.

Employing the one-loop anomalous dimension (4.20) to solve the RG equation

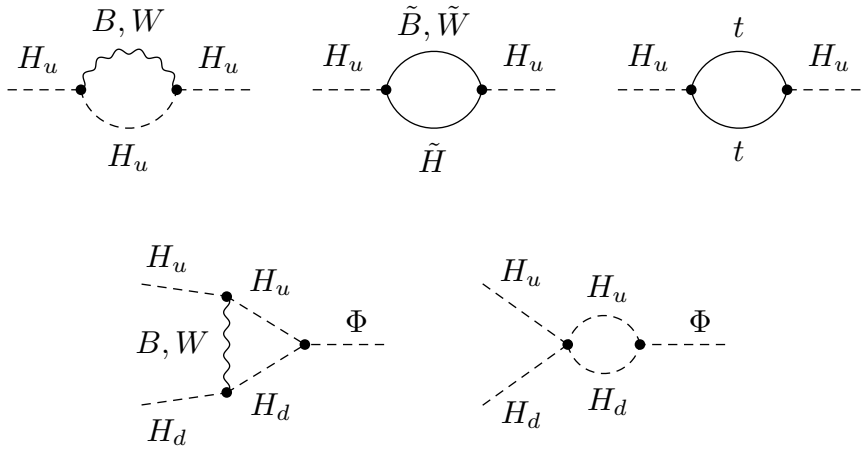

Figure 4.1: The $H_u H_u$ self-energy diagrams (upper row) and $\Phi H_u H_d$ vertex diagrams (lower row) that contribute to the one-loop anomalous dimension $\gamma_Z^{(1)}$. The $H_d H_d$ self-energy diagrams needed to determine the wave-function renormalisation factor of H_d are not shown.

(4.18), we find to leading logarithmic accuracy

$$K \equiv \frac{Z(m_\Phi)}{Z(M_{\rm string})} \simeq 1 - \gamma_Z^{(1)} \ln\left(\frac{M_{\rm string}}{m_\Phi}\right) \simeq 1 - \frac{18}{(4\pi)^2}\left(-1.7 + \frac{1.5}{\sin^2\beta} + \frac{1.6\cdot 10^{-4}}{\cos^2\beta}\right), \tag{4.21}$$

where $\gamma_Z^{(1)}$ is calculated at the string scale. To arrive at the numerical expression we have employed $g_1(M_{\rm string}) \simeq 0.65$, $g_2(M_{\rm string}) \simeq 0.69$, $y_t(M_{\rm string}) \simeq 0.70/\sin\beta$, $y_b(M_{\rm string}) \simeq 6.0\times 10^{-3}/\cos\beta$ and $y_\tau(M_{\rm string}) \simeq 7.2\times 10^{-3}/\cos\beta$, corresponding to $M_{\rm string} = 3\times 10^{14}\,{\rm GeV}$ and $m_\Phi = 7\cdot 10^6\,{\rm GeV}$.[3] In the final result in (4.21), we have shown the contributions arising from the terms $g_{1,2}^2$, $|y_t|^2$ and $|y_{b,\tau}|^2$ separately. The different overall signs multiplying the contributions from the gauge and the Yukawa couplings imply that the individual terms in (4.20) tend to cancel. In fact, the numerical expression for $\gamma_Z^{(1)}$ used in (4.21) is less than 0 for $3 \lesssim \tan\beta \lesssim 35$ and vice versa. We hence expect to find that loop corrections suppress (enhance) the partial decay rate $\Gamma(\Phi \to H_u H_d)$ for small and large (moderate) ratios of the Higgs vacuum expectation values, $\tan\beta$. In order to obtain a reliable prediction for K, however, the large logarithm appearing in (4.21) has to be resummed by solving (4.18) together with the RG equations describing the scale dependence of the gauge and Yukawa

[3]These values for the couplings have been obtained using the numerical procedure outlined in section 4.5.1.

couplings.

4.5 Numerical Results

After presenting the analytic result for the anomalous dimension of the Guidice-Masiero coupling, we now turn to the numerical RG analysis of the $\Phi \to H_u H_d$ decay mode. Our methodology is detailed in the following.

4.5.1 Solution of RG Equations

The system of differential equations describing the renormalisation scale dependence of the coupling strength Z as well as those of the gauge and Yukawa couplings is solved iteratively with the help of SoftSusy 3.3.7 [100]. The calculation is performed including all relevant one-loop and two-loop effects.[4] The fine structure constant $\alpha(m_Z) = 1/127.973$, the Fermi constant $G_F = 1.16637 \cdot 10^{-5}\,\mathrm{GeV}^{-2}$, the strong coupling $\alpha_s(m_Z)$, the pole mass m_t of the top quark, the bottom mass $m_b = 4.2\,\mathrm{GeV}$ and the tau mass $m_\tau = 1.777\,\mathrm{GeV}$ serve as SM inputs and constraints in the RG evolution. The low-energy boundary conditions are applied at the Z-boson mass $m_Z = 91.1875\,\mathrm{GeV}$. At the string scale M_{string} we impose minimal supergravity (MSUGRA) boundary conditions, which just leaves five free SUSY parameters: common scalar and gaugino masses, m_0 and $m_{1/2}$, universal trilinear terms A_0, the bilinear soft SUSY breaking term B and the SUSY μ parameter. Following common practice, we use the one-loop corrected electroweak symmetry breaking (EWSB) conditions (see e.g. [101]) to trade B and the magnitude $|\mu|$ in favour of $\tan\beta$ and the sign of μ. Notice that the assumed scaling of $m_0 \sim m_{1/2} \sim M_{pl}/\mathcal{V}^2$ naturally requires $B \sim M_{pl}^2/\mathcal{V}^4$ and $\mu \sim M_{pl}/\mathcal{V}^2$ to achieve EWSB. We assume that these scalings are realised by an appropriate volume sequestering, and furthermore take $A_0 \sim M_{pl}/\mathcal{V}^2$. The SUSY scale is determined by the geometric mean $M_{\mathrm{soft}} \equiv \sqrt{m_{\tilde{t}_1} m_{\tilde{t}_2}}$ of the masses $m_{\tilde{t}_{1,2}}$ of the stop mass eigenstates. Finally, the mass of the volume modulus is obtained

[4] Two-loop effects have however a very minor impact on the analysis, and therefore we only reported the result of the one-loop anomalous dimension $\gamma_Z^{(1)}$ in (4.20).

from $m_\Phi = M_{\text{string}}^3/M_P^2$.

4.5.2 SM and MSUGRA Parameter Dependencies

Before studying the effects of varying the MSUGRA parameters on K, in equation 4.21, we consider the impact of the parametric SM errors. The dominant sources of SM uncertainties arise from the top mass and the strong coupling constant. This is to be expected because the anomalous dimension of Z (equation 4.20) is quadratic in the top Yukawa coupling and the RG evolution of y_t depends sensitively on the low-energy initial conditions for m_t and α_s. The more critical ingredient is the top mass for which the latest Tevatron measurements find $m_t = (173.2 \pm 0.9)\,\text{GeV}$ [102]. The latest world average, including new LHC ATLAS and CMS results [30], is within this uncertainty, thus will not affect our conclusions. However, the exact meaning of the mass parameter measured by experiments such as CDF, D0, ATLAS and CMS via a kinematical reconstruction of the top decay products and comparison to Monte Carlo simulations is unclear and so is its connection to y_t. A theoretically well-defined determination of m_t can, on the other hand, be obtained from the total cross section for top-quark pair production. While such extractions (see e.g. [103]) give values for m_t that are compatible with the mass determinations from direct reconstruction, the achieved accuracy is notably worse, with an uncertainty of around $\pm 5\,\text{GeV}$. We use the following value of the strong coupling evaluated at the Z-boson mass $\alpha_s(m_Z) = 0.1184 \pm 0.0007$ [104]. This value of α_s is obtained from a large set of measurements with significant spreads between them. To account for this fact we will also give results employing the 3σ error ± 0.0021 of the α_s world average. Note the latest world average for $\alpha_s(M_Z)$ [30] is again well within the 1σ quoted errors we use in our calculation, and will thus not affect our conclusions at all.

Our predictions for K as a function of M_{soft} are shown in figure 4.2. The results displayed in the left panel correspond to $m_0 = m_{1/2} = A_0$, $\tan\beta = 10$ and $\text{sign}\,\mu = +1$. Almost identical predictions are obtained for different choices of A_0 and $\text{sign}\,\mu = -1$. The solid black curve corresponds to $m_t = 173.2\,\text{GeV}$ and $\alpha_s(m_Z) = 0.1184$, while the

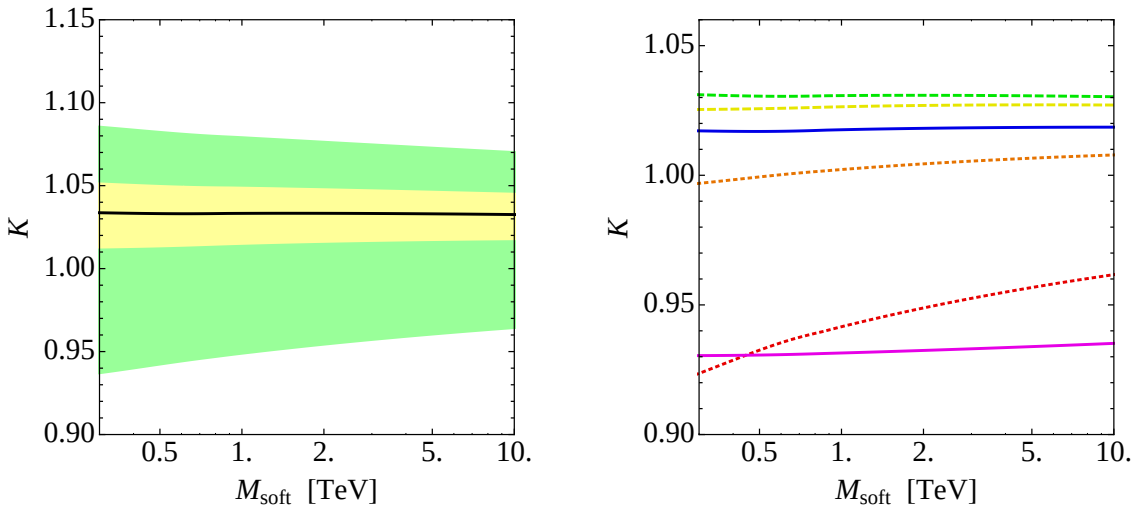

Figure 4.2: Left: Predictions for K for fixed MSUGRA input. The solid black line indicates the result obtained for the central choice of SM inputs while the coloured bands reflect the uncertainties associated with the errors in the top mass and the strong coupling constant. Right: Predictions for K for different values of $\tan\beta$. See text for further explanations.

yellow (green) band has been obtained by varying m_t and $\alpha_s(m_Z)$ by $\pm 0.9\,\text{GeV}$ and ± 0.0007 ($\pm 5\,\text{GeV}$ and ± 0.0021) around their central values. We see that the ratio (4.21) is largely independent of the SUSY scale $M_{\rm soft}$, but that the exact value of K depends to some extent on the low-energy input m_t and $\alpha_s(m_Z)$. Numerically, we find that the variations of $\pm 0.9\,\text{GeV}$ and ± 0.0007 ($\pm 5\,\text{GeV}$ and ± 0.0021) lead to shifts in K of less than $\pm 2\%$ $(^{+5\%}_{-10\%})$ relative to the central values. The largest value of (4.21) is thereby attained for the smallest value of m_t and the largest value of $\alpha_s(m_Z)$, and vice versa.

We now analyse the dependence of K on the choice of $\tan\beta$. Our numerical results are shown in the right panel of figure 4.2. All curves have been obtained for $m_0 = m_{1/2} = A_0$, $\text{sign}\,\mu = +1$, $m_t = 173.2\,\text{GeV}$ and $\alpha_s(m_Z) = 0.1184$. The dotted red, dotted orange, dashed yellow, dashed green, solid blue and solid magenta lines correspond to $\tan\beta = 2, 3, 5, 15, 25$ and 50, respectively. As anticipated, we find that for $\tan\beta \lesssim 3$ the predictions for the ratio (4.21) are below 1, while for moderate values of $\tan\beta$ one obtains ratios above 1. In fact, the values of K saturate for $\tan\beta \simeq 10$ and increasing $\tan\beta$ further leads to a suppression of the ratio $\big($a feature that is also reproduced by the simple formula (4.21)$\big)$. For large $\tan\beta$ values, the ratio K then ends up below 1. We see furthermore that varying $\tan\beta$ in the range $[2, 50]$ shifts K

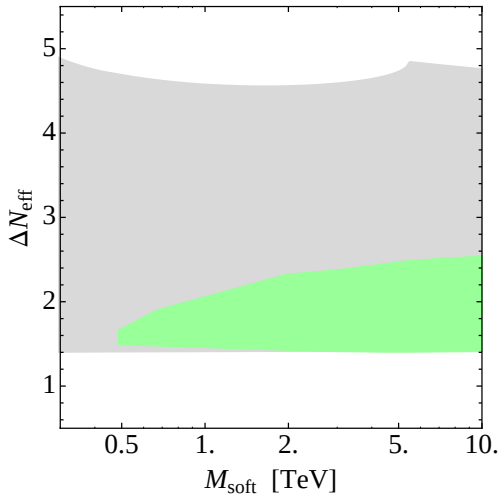

Figure 4.3: Predictions for the effective excess number of neutrino species in the MLVS framework. The coloured wedge-shaped region indicates the possible values of $\Delta N_{\rm eff}$ consistent with the LHC measurements of a Higgs-like state near 126 GeV. For comparison the accessible parameter space without imposing the Higgs constraint is underlaid in grey. For further details see text.

by only $^{+3\%}_{-7\%}$ away from 1. The dependence on the other MSUGRA parameters are even less pronounced than that on $\tan\beta$.

4.5.3 Predictions for the Effective Excess Number of Neutrinos

It is well-known that the mass m_h of the Higgs boson puts stringent constraints on the MSUGRA parameter space. This is particularly true after the discovery of a relatively heavy Higgs-like state with a mass of around 126 GeV by ATLAS [105] and CMS [106].

We assess the impact of the LHC measurements of the Higgs mass on the predictions for $\Delta N_{\rm eff}$ by performing a global scan in the MSUGRA parameter space. Only points that lead to $m_h \in [123, 129]$ GeV are retained, which is the range allowed by the ATLAS and CMS data (see [107] for the latest ATLAS and CMS combination) if one accounts for the theoretical uncertainties in the MSSM calculation of the Higgs mass (see e.g. [108]). We generate a large sample of points allowing the MSUGRA parameters to take random values within $m_{0,1/2} \in [0.1, 10]$ TeV,

$A_0 \in [-30, 30]\,\text{TeV}$ and $\tan\beta \in [1, 60]$, permitting μ to be of either sign. In order to incorporate SM uncertainties we let the top mass and the strong coupling constant float within $m_t = (173.2 \pm 5)\,\text{GeV}$ and $\alpha_s(m_Z) = 0.1184 \pm 0.0021$, respectively.

The predicted value of $Z(m_\Phi)$ at each parameter point is then converted into a predicted range of ΔN_{eff} using

$$\frac{1.56}{Z^2} \leq \Delta N_{\text{eff}} \leq \frac{1.74}{Z^2}\,. \tag{4.22}$$

Note that the range takes into account the uncertainty associated with the value of the reheating temperature. Since we effectively scan over all individual sources of uncertainties, the derived limits on ΔN_{eff} should be considered very conservative.

Our results of the MSUGRA scan are shown in figure 4.3. The accessible parameter before (after) imposing the Higgs-mass constraint is indicated by the grey (coloured) region. We see that in the MLVS the values for ΔN_{eff} compatible with the m_h constraint lie in the narrow range of about $[1.4, 2.6]$, and that the width of the allowed region is essentially constant for $M_{\text{soft}} \gtrsim 5\,\text{TeV}$. The constraint due to the Higgs mass influences the predictions for ΔN_{eff} only indirectly by narrowing down the possible values of M_{soft} and $\tan\beta$. This effect is most visible for $M_{\text{soft}} \lesssim 2\,\text{TeV}$, since such relatively low values of M_{soft} require large values of $\tan\beta$ to push the Higgs mass up to around 126 GeV. Notice also that the constraint from m_h cuts away the parts of the parameter space with $\Delta N_{\text{eff}} \gtrsim 2.6$ and $M_{\text{soft}} \lesssim 0.5\,\text{TeV}$. Both regions are inaccessible because they correspond to either $\tan\beta \lesssim 2$ or to a too light stop spectrum. We expect that other low-energy constraints (such as e.g. flavour physics) have an even smaller impact on the limits obtained for ΔN_{eff} than m_h. The latest Planck measurement of N_{eff} [8] gives $\Delta N_{\text{eff}} = 0.08 \pm 0.32$ ($\Delta N_{\text{eff}} = -0.06 \pm 0.20$ with polarization data). The minimal value of $\Delta N_{\text{eff}} \simeq 1.4$ that is attainable in the MLVS framework thus corresponds to a discrepancy of about $4\,\sigma$ ($7\,\sigma$) between theory and experiment. These findings basically rule out the MLVS as a model of dark radiation.

4.6 Conclusions

The latest Planck results have ushered in a new era of precision cosmology. These measurements support the standard ΛCDM cosmological model, and the bounds on the presence of dark radiation are very stringent. This is especially interesting since many beyond the standard model theories predict light hidden particles which could be produced abundantly during reheating. In light of this, in this chapter we have analysed the predictions for dark radiation in a class of string theory compactifications called the Large Volume Scenario. We pick a minimal version of these models with one light ALP, an MSSM visible sector, and shift-symmetric Higgs sector called the Minimal LVS (MLVS). After reviewing the tree level prediction, we have analysed loop corrections to $\Delta N_{\rm eff}$ in this MLVS.

In this class of models, additional contributions to the effective excess number of neutrinos are an unavoidable consequence of the presence and the interactions of a light volume modulus Φ: the decays of this field to the visible sector drive the reheating of the Universe after inflation, while dark radiation arises from its decays to an ultralight ALP partner a. The only visible-sector decay mode that can compete with the ALP channel is the decay into Higgs pairs induced by a Giudice-Masiero term. The interplay between the two channels, $\Phi \to a\,a$ and $\Phi \to H_u H_d$, fixes the relative fraction of dark radiation uniquely in terms of the coupling strength Z between Φ and the bilinear $H_u H_d$. Under the assumption that the coupling Z is set to 1 at the string scale by means of a shift-symmetric Higgs sector, the ratio of branching ratios of visible-sector and hidden-sector decays can then be predicted accurately. At the tree level such a calculation leads to $\Delta N_{\rm eff} \simeq 1.7$, at variance with observation.

Unlike the coupling of the volume modulus to its ALP partner, which receives only Planck-suppressed contributions, the $\Phi\, H_u H_d$ coupling is modified by MSSM loops. These radiative corrections induce large logarithms that are formally of $\mathcal{O}(1)$, and hence have to be resummed to all orders. In our work, we have calculated the anomalous dimension γ_Z of the composite operator $H_u H_d \Box\Phi$ needed to perform such

a resummation. We found that the size of the leading-logarithmic corrections to the coupling strength Z depends sensitively on the ratio of the Higgs vacuum expectation values, $\tan\beta$, through the top Yukawa coupling. As a result, loop corrections suppress $\Gamma(\Phi \to H_u H_d)$ for $\tan\beta \lesssim 3$ and $\tan\beta \gtrsim 35$, while the partial decay rate to Higgs pairs is enhanced for all other $\tan\beta$ values. The maximal enhancements occur for $\tan\beta \simeq 10$, but amount to below 10% only.

This simple pattern of suppressions and enhancements is also reproduced by our high-statistics MSUGRA scan, which includes all relevant two-loop effects. Specifically, we find that in the MLVS the values of ΔN_{eff} that are compatible with a Higgs-boson mass close to 126 GeV all lie in the range $[1.4, 2.6]$. The spread of the predictions is rather insensitive to the exact values of the MSUGRA parameters m_0, $m_{1/2}$ and $\operatorname{sign}\mu$, and is influenced by the Higgs mass requirement only indirectly because this constraint needs tuning of A_0 and $\tan\beta$. In consequence, it turns out that for moderate values of $\tan\beta$, radiative corrections tend to suppress the tree-level prediction $\Delta N_{\text{eff}} \simeq 1.7$. The loop-induced effects are however always small, leading to a robust lower bound of $\Delta N_{\text{eff}} \gtrsim 1.4$. This limit corresponds to a $\sim 4\sigma$ tension between theory and experiment, which essentially excludes the MLVS.

The production of ALP dark radiation by moduli decay at reheating appears to be unavoidable. For now we put aside the question of whether a certain string model evades the bounds on dark radiation, and instead ask whether these ALPs could be observed. Since these ALPs do not interact, they should still exist as a homogeneous background akin to the CMB. Such an observation would be a direct test of early universe physics and in particular string theory models. We look at the possibility of using galaxy clusters to observe these ALPs in the next chapter.

Since the publication of the paper this chapter is based on, several studies have looked at dark radiation production in other string theory compactification models. Dark radiation production has been studied in the context of anisotropic versions of the Large Volume Scenario [109], it was found that such models over-produce dark radiation by several orders of magnitude and thus are ruled out. The authors of [110]

considered three different models again within the Large Volume Scenario, where the volume modulus' decay to SM gauge bosons was not loop suppressed and thus found regions of parameter space where dark radiation bounds could be evaded. Dark radiation bounds were also found to be evaded if one considered split supersymmetry-like models (see [111, 112] for a discussion of split-supersymmetry), instead of the MSSM [113]. Here decays to scalars of a similar mass to the modulus enhances the branching fraction to visible-sector fields. Finally, [114] looked at dark radiation in the context of the axiverse, and showed that when specific conditions are held, dark radiation could again be reduced. They showed these conditions could be realised in both the Large Volume Scenario, and for M-theory compactified on a G_2-holonomy manifold. These studies confirm the ability of dark radiation to probe string models.

5 A Cosmic ALP Background and the Galaxy Cluster Soft X-ray Excess

In this chapter I will look at the question of whether a population of relativistic ALPs produced in the early universe could be detected using X-ray astronomy. I have shown in the previous chapters the production of light (massless) ALPs during reheating is generic in string theory compactifications. Here we focus on a potential signal of this ALP dark radiation. I start by reviewing the work of [115] which showed that today these ALPs would form a homogeneous background. For an appropriate choice of modulus mass, this background has X-ray energies. In the rest of the chapter I show using detailed simulations that this background of ALP dark radiation could explain a long-standing puzzle of galaxy cluster X-ray astronomy—the soft X-ray excess. The idea that the soft X-ray excess could be explained in this way was originally proposed in [116]. This chapter is based on the publications [2], work done and written in collaboration with Stephen Angus, Joseph P. Conlon, David Marsh and Lukas Witkowski, and [3].

5.1 Introduction

We have seen in Chapter 4 that non-thermal production of ALPs in the early universe is well-motivated in models arising from compactifications of string theory to four dimensions. These ALPs arise from moduli decays at the time of reheating, and, due to their very weak interactions, would linger today as a homogeneous and isotropic background. We call this background a cosmic ALP background (CAB). The CAB has a non-thermal spectrum determined by the expansion of the universe during the time of moduli decay [115]. For moduli masses $m_\Phi \approx 10^6\,\text{GeV}$ the present energy of these ALPs is $E_a \sim 0.1 - 1\,\text{keV}$. In this chapter we will use this background of ALPs to explain an excess of soft X-ray photons seen in many galaxy clusters.

Galaxy clusters are the largest gravitationally bound objects in the universe and have historically served as powerful indicators of novel fundamental physics [117]. In addition to the dark matter component comprising around 80% of the cluster mass, around 15% of the mass is in a hot ionised intra-cluster medium (ICM) with typical temperatures of $T \approx 10^8\,\text{K}$ (corresponding to $\omega \approx 7\,\text{keV}$) and number densities $n \sim 10^{-1} - 10^{-3}\ \text{cm}^{-3}$. The ICM represents the large majority of a cluster's baryonic mass and generates diffuse X-ray emission through thermal bremsstrahlung.

Observations of a large number of galaxy clusters have found evidence at low energies around $E \lesssim 0.4\,\text{keV}$, for excess emission above that from the hot ICM. This soft excess was initially observed in the Virgo and Coma clusters in 1996 [118–120] and has since been found in many other clusters [121, 122]. There are two candidate astrophysical explanations: emission from a warm $T \approx 0.1$ keV gas; and inverse Compton scattering of non-thermal electrons on the cosmic microwave background (CMB). The former explanation has difficulty with rapid cooling times of a warm gas and the lack of associated line emission; the latter has difficulty remaining consistent with the observed level of synchrotron radio emission and the failure to detect clusters in gamma rays.

In [116], the cluster soft excess was proposed to arise from conversion of a CAB

into photons in the magnetic field of galaxy clusters.

The existence of such a CAB is indirectly probed through its contribution to dark radiation, see Chapter 2. The latest Planck satellite [83] measurements place very stringent constraints on the amount of dark radiation in the early universe. They measure $N_{\rm eff} = 3.13 \pm 0.32$ at 68% confidence [8], dark radiation is even more constrained if one includes Planck polarisation data $N_{\rm eff} = 2.99 \pm 0.20$ [8]. This leads to the conclusion that at the 3σ level, $\Delta N_{\rm eff} > 1$ is excluded. In combination with the latest deuterium [85] and helium [86] abundances, the Planck 95% C.L. constraints $N_{\rm eff} = 2.95^{+0.52}_{-0.52}$ and $N_{\rm eff} = 3.11^{+0.59}_{-0.57}$ respectively [8]. Thus the energy density of the CAB is very tightly constrained.

In this chapter we will take the ALP mass to be zero, although in practice the physics is unaffected for any mass $m_a \lesssim 10^{-12}$ eV. We will also set the energy density of the CAB such that the current limits on dark radiation are saturated, i.e. we set $\Delta N_{\rm eff} = 0.5$, but where applicable we will discuss how our results scale with $\Delta N_{\rm eff}$, should the bounds become more restrictive (or a discovery made). It should also be noted that if more than one ALP is produced as dark radiation, not all of these ALPs may couple to photons, and thus the $\Delta N_{\rm eff}$ we use for our results would need to be smaller than the current bound.

In the presence of a magnetic field, axion-like particles can directly convert into photons of the same energy via the coupling,

$$\mathcal{L}_{a\gamma\gamma} = \frac{1}{4M} a F_{\mu\nu} \tilde{F}^{\mu\nu} = \frac{1}{M} a \mathbf{E} \cdot \mathbf{B} \,, \tag{5.1}$$

see Chapter 3. In the enlightening case of sufficiently high ALP energies or small ambient electron densities, the conversion probability for a fixed domain is given by

$$P_{a\to\gamma} = \frac{1}{4} \left(\frac{B_\perp L}{M} \right)^2 \,, \tag{5.2}$$

where $B_\perp$ denotes the magnetic field component transverse to the ALP velocity and L denotes the corresponding coherence length [35]. This conversion allows the potential detection of a CAB through ALP-photon conversion.

Galaxy clusters support magnetic fields that are modest in magnitude ($B \approx \mu$G) but are extended over megaparsec distances and have kiloparsec coherence scales, allowing observationally significant ALP-photon conversion probabilities. In [116], a crude model with a constant magnitude and fixed coherence length for the magnetic field was used to estimate the ALP-photon coupling M that would be required to reproduce the soft excess in Coma from a CAB, finding $M \approx 10^{13}$ GeV.

In this chapter we continue the study of ALP-photon conversion in galaxy clusters using a far more detailed magnetic field model, first constructed in [123]. In this model of the cluster magnetic field, the magnetic field has a simple power-law power spectrum, with cut-offs such that the magnetic field only has power on scales between $\Lambda_{\rm min}$ and $\Lambda_{\rm max}$. Using this stochastic model, we construct a numerical simulation of the magnetic field in four clusters: Coma, A665, A2199, and A2255. We then numerically propagate ALPs through it and quantitatively study the resulting predictions for the soft excess magnitude and morphology.

This chapter is organised as follows. I start in section 5.2 by introducing what we call a cosmic ALP background (CAB), produced by moduli decay in the early universe. Section 5.3 reviews the cluster soft excess phenomenon, the proposed astrophysical explanations and the constraints on these explanations. Section 5.4 reviews knowledge of galaxy cluster magnetic fields, introduces the model used in this work and describes the method used to produce a simulation of the magnetic field. Then in section 5.5 I describe ALP–photon mixing, and the numerical methodology involved in calculating ALP-photon conversion probabilities. The results for the detailed study of the Coma cluster, first published in [2], are presented in section 5.7. In section 5.8, I present a follow-up study of three more clusters: A665, A2199, and A2255, published in [3]. In section 5.9 I present the conclusions from these studies.

5.2 A Cosmic ALP Background

As discussed in Chapter 2 if moduli fields arising from the compactification of string theory exist, general arguments imply they should be responsible for the reheating of

the Standard Model degrees of freedom. Almost independently of the detailed model of inflation, moduli become displaced from their final metastable minimum during inflation and begin to oscillate at the end of inflation. The energy density stored in this oscillating field will come to dominate the energy density of the universe. This modulus-dominated stage of the universe's history lasts until the moduli decay into visible and hidden sector matter and radiation, thus inducing reheating.

The energy density of the universe at the time the modulus decays, $\tau_{\rm decay}^{-1} \sim H \sim \Gamma \sim \frac{1}{8\pi}\frac{m_\Phi^3}{M_{pl}^2}$, is

$$V_{\rm decay} \sim 3H^2 M_{pl}^2 \sim \frac{m_\Phi^6}{M_{pl}^2}\,. \tag{5.3}$$

The visible sector decay products of the modulus rapidly thermalise and initiate the Hot Big Bang at a temperature

$$T_{\rm reheat} \sim \frac{m_\Phi^{3/2}}{M_{pl}^{\frac{1}{2}}} \sim 1\,{\rm GeV}\left(\frac{m_\Phi}{10^6\,{\rm GeV}}\right)^{3/2}. \tag{5.4}$$

However, the gravitational origin of the moduli implies that moduli can also decay to any hidden sector. Indeed we have already shown in Chapter 4 that one expects large branching fractions to hidden sectors, which include decays into hidden sector massless particles with extremely weak interactions (such as ALPs) [1,77,78,92] (also see [124]).

The decay of a modulus field into ALPs is induced by the Lagrangian coupling $\frac{\Phi}{M_{pl}}\partial_\mu a \partial^\mu a$, resulting in ALPs with an initial energy $E_a = m_\Phi/2$. Since they are weakly interacting, the ALPs do not thermalise and the vast majority of ALPs propagate freely to the present day, where they form a homogeneous and isotropic background—a cosmic ALP background. Being relativistic, they contribute to the dark radiation energy density of the universe, and thus the energy density of the CAB is probed indirectly through dark radiation measurements.

The characteristic ALP energy today is set by the initial ALP energy, redshifted to the present. Since the current CMB temperature is found simply by redshifting the primordial thermal plasma (up to small $\left(\frac{g_{*,\rm now}}{g_{*,\rm init}}\right)^{-1/3}$ boosts as species decouple),

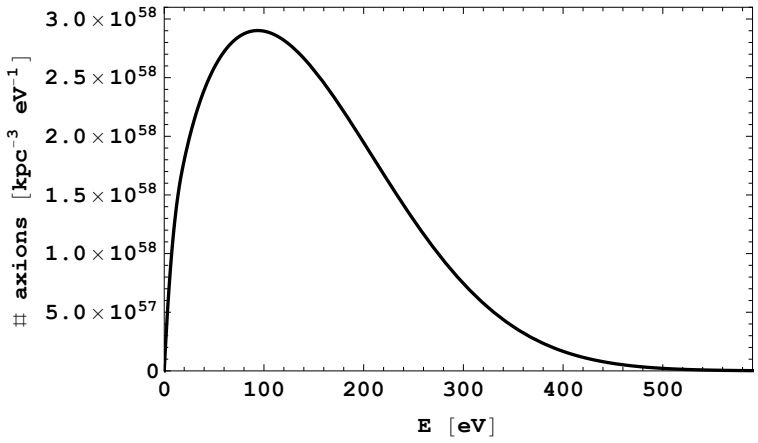

Figure 5.1: A typical ALP number density per $(\mathrm{kpc})^3$ for a CAB with $\langle E_{CAB}\rangle = 150$ eV, which contributes to dark radiation with $\Delta N_{\mathrm{eff}} = 0.5$. The precise location of the energy peak depends on the value of m_Φ.

we have

$$\frac{E_{a,\mathrm{now}}}{T_{\gamma,\mathrm{now}}} \simeq \frac{E_{a,\mathrm{init}}}{T_{\gamma,\mathrm{init}}} \sim \left(\frac{M_{pl}}{m_\Phi}\right)^{1/2} .$$

For moduli masses $m_\Phi \approx 10^6\,\mathrm{GeV}$, this gives $E_a \sim 10^6\, T_{\mathrm{CMB}} \approx 200\,\mathrm{eV}$.

To find the exact spectral shape of the CAB, we must account for the fact that moduli do not decay instantaneously, and meanwhile the expansion rate of the universe changes as it transitions from matter (modulus) domination prior to reheating into radiation domination after all moduli have decayed. Moduli that decay early give rise to present-day lower-energy ALPs as they have more time to redshift, whereas more energetic ALPs arise from late-decaying moduli. The spectral shape was computed numerically in [115] and may be described as 'quasi-thermal', with an exponential fall-off at high energies (c.f. figure 5.1). The overall magnitude is normalised to the ALP contribution to ΔN_{eff}, and the peak location is determined by the mass of the modulus and its lifetime.

While a CAB can be indirectly probed through studies of dark radiation, it can be directly observed only through its couplings to visible-sector matter and gauge bosons, as mediated, for example, by the operator

$$\frac{a}{M}\mathbf{E}\cdot\mathbf{B}\,.$$

In the presence of a magnetic field this induces 'oscillations' of ALPs into photons in

a process analogous to neutrino oscillations [35,36], see Chapter 3. The observational consequences of this conversion of the CAB have been considered in [116, 125, 126]. ALPs may also play a role by scattering off ambient particles in the thermal plasma, which was considered in [115]. Giving the value of M, the total ALP energy density, and the central CAB energy specifies an entirely predictive model. In this model, the spectrum and number of photons arising from ALP–photon conversion in any astrophysical magnetic field can be computed.

5.3 The Cluster Soft X-ray Excess

Galaxy clusters are the largest virialised structures in the universe, with typical masses $\sim 10^{14}-10^{15} M_{\odot}$ and spatial extents of $\mathcal{O}(1\,\mathrm{Mpc})$. Clusters emit light of all frequencies ranging from radio waves to X-rays. The space between galaxies is suffused with an energetic ionised plasma, termed the intra-cluster medium (ICM). The energy of the ICM arose from the release of gravitational potential energy as the cluster formed through accretion and merger of sub-clusters. The hot ICM has a temperature $T \sim 2 - 8\,\mathrm{keV}$ and emits X-rays via thermal bremsstrahlung, with typical X-ray luminosities of $\mathcal{L} \sim 10^{42-45}\,\mathrm{erg\ s^{-1}}$ (where $1.6 \times 10^{-3}\,\mathrm{erg} = 1\,\mathrm{GeV}$).

The aim of this section is to summarise the observational history of the soft excess from galaxy clusters and to discuss the astrophysical models proposed to account for it. We also review the status of these models in light of more recent measurements of the Coma magnetic field as well as the (null) observations of galaxy clusters in gamma rays. A comprehensive review of the soft excess phenomenon from 2008 is given in [122].

5.3.1 History of Soft Excess Observations

Excess soft X-ray emission from galaxy clusters has been consistently found by a number of satellites and in a significant number of clusters. In this section we review how the different satellite missions have contributed to observational evidence for the excess.

Accurate measurements of soft X-ray fluxes from distant sources require accurate measurements of the local soft X-ray background. The soft X-ray sky has gradients on the scale of a few degrees and is also subject to temporal variation due to solar flares and charge exchange scattering between the solar wind and the Earth's exosphere. Galaxy clusters are large objects, with a typical diameter $d \approx 1\,\mathrm{Mpc}$, and the intra-cluster gas generates diffuse emission with relatively low surface brightness. The ideal background is therefore one which is both spatially and temporally contiguous. This is most easily accomplished if the observing telescope has a large field of view, which can then accommodate the entire cluster and allow the background to be taken from the edge region of the detector.

The original discovery of the cluster soft excess phenomenon was made in 1996 with observations of the Virgo and Coma clusters [118–120] using the Extreme Ultraviolet Explorer (EUVE). The EUVE Lex/B detector [127] had peak response at 138 eV and very limited spectral resolution. The properties of the ICM gas were determined by measurements on complementary instruments with X-ray sensitivity. The EUVE observations revealed a large EUV excess over the level of emission expected from thermal bremsstrahlung from the hot gas. By default, these observations were interpreted as evidence for a warm gas ($T \approx 5 \times 10^5\,\mathrm{K} \simeq 50\,\mathrm{eV}$) within the cluster.

An excess was then found in the clusters A1795 and A2199 [128, 129], where very large soft excesses were reported. However there was disagreement about the details of background subtraction for the EUVE satellite and the variation in telescope sensitivity over the field of view. Several re-analyses have disagreed over the presence and detailed properties of the excesses in specific clusters [130–135]. The analysis [136] found EUV soft excesses present in Coma, Virgo, A1795, A2199, and A4059. The magnitude of the soft excess became more pronounced at large radii. No consensus was reached over the existence of an EUV soft excess in the other clusters observed with EUVE.

The other main satellite in which the soft excess has been observed was ROSAT [137] which carried out an all-sky survey. ROSAT was the most suitable satellite for

studying the excess. It had an energy range of $0.1 - 2.4\,\mathrm{keV}$ with a large field of view of 2° diameter allowing the full cluster to be seen, and a low internal background. The large field of view allows the entire cluster to be seen in one pointing, allowing a temporally contiguous background to be estimated from the edge of the image. The spectral resolution was however limited, with $\Delta E/E = 0.43\sqrt{0.93\,\mathrm{keV}/E}$.

ROSAT has been used to study the soft excess both through individual papers on particular clusters and through a large statistical survey [121]. A soft X-ray excess was seen in many clusters individually: A3558, A3560, A3562 and A3571 [138]; clusters in the Hercules supercluster [139]; and Sersic-159 (a.k.a. AS1101) [140, 141].

The largest study to date of the soft excess was carried out in [121], in which 38 clusters were studied. This study looked for excess emission in the ROSAT $0.1 - 0.28$ keV band compared to expectations from the hot gas. In most cases (except for the nearest clusters), the large ROSAT field of view enabled an *in situ* background measurement from the peripheral detector regions.

The study found that soft excess emission was a general feature of galaxy clusters. A statistically significant detection was observed in 30% of the sample. In some cases, such as Coma, at very high statistical significance. A soft excess was only ruled out in a small number of clusters. Soft excess emission was found for both nearby and distant clusters, for both relaxed and disturbed clusters, and at a wide range of Milky Way neutral hydrogen absorption column densities. In all cases the relative magnitude of the soft excess compared to the thermal ICM emission did not exceed 30 percent. The soft excess appeared to follow an identifiable morphological trend in that excess emission was preferentially found at radii $r \gtrsim 175\,\mathrm{kpc}$, outside the very centre of clusters. In later sections we will comment on how the properties of ALP–photon conversion in cluster magnetic fields may explain this trend.

Soft excess emission has also been studied with XMM-Newton, Suzaku and Chandra. These telescopes have better spectral resolution than ROSAT but have much smaller fields of view and energy ranges not covering all of soft excess range.

Statistical studies of large samples of clusters with XMM-Newton were carried out

in [142–145]. Searches for emission lines with XMM proved null [146–149] (although see [150] for a marginal detection), with previous detections [143, 144, 151] explained by the local or galactic background [152–154].

The current state of soft excess detection in XMM-Newton remains unclear however. The XMM-Newton detections of soft excess were challenged in [153], where it was argued that they arose from an incorrect background subtraction. This issue was therefore revisited in [154]. With a new calibration, these authors found that soft excesses were consistently detected by one of the two instruments (MOS) while being absent in the other (PN).

Let us also mention the (limited) studies of the soft excess that have been carried out with Suzaku and Chandra. The limited field of view again makes these suboptimal instruments for studies of diffuse soft emission. The soft excess emission for the cluster A3112 was studied in [155] (with Chandra) and in [148] (with Suzaku), both finding an excess. Sersic 159 (AS1101) was studied with Suzaku in [147]. This study confirmed the existence of the soft excess but did not find the tentative O_{VII} emission line reported by earlier XMM-Newton studies. Chandra data on the same cluster was considered in [141], again an excess was found.

5.3.2 Specific Clusters

Throughout this work we will compare our simulations to the observations of [121], which is the largest complete study of the soft X-ray excess in 38 galaxy clusters, and currently appears unchallenged. ROSAT is also the best instrument for the study of the soft excess, and unlike the case for EUVE there is no controversy about the precise mechanism of background subtraction. The excess is usually phrased as the fractional excess, ξ, defined as

$$\xi \equiv \frac{p-q}{q}, \tag{5.5}$$

where p and q are the observed and expected counts in the $0.1 - 0.28$ keV band respectively. We show the fractional excess in several concentric annuli for the clusters involved in this study in figure 5.2. Here we summarise the observational standing of

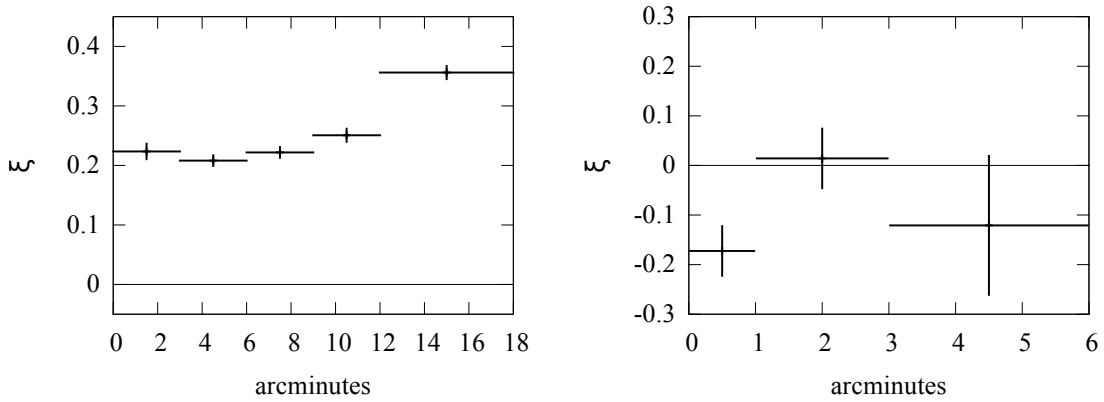

(a) Coma

(b) A665

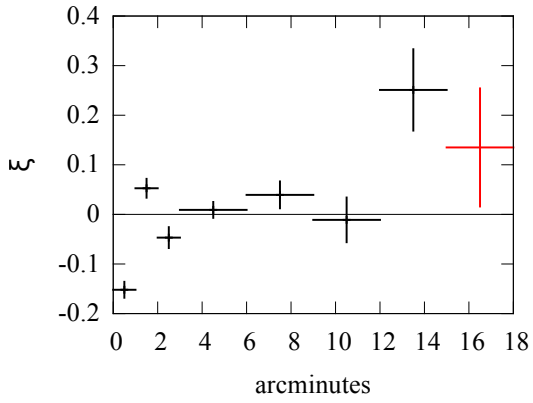

(c) A2199

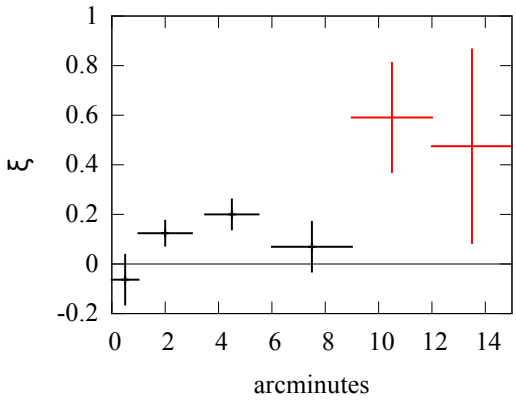

(d) A2255

Figure 5.2: The fractional excess, ξ, over the expected thermal bremsstrahlung emission in the $0.1 - 0.28$ keV band, in several annuli for the four clusters a) Coma, b) A665, c) A2199, and d) A2255. Data taken from [121]. Points coloured red indicate that the total (thermal and soft excess) flux in the $0.1 - 0.28$ keV band does not exceed the background by 25%.

the four clusters involved in this study, focusing on ROSAT data.

Coma

Coma was one of the first clusters in which a soft excess was found. It is also one of the few clusters for which EUVE conclusively found a soft excess. The centre of the Coma cluster was studied with ROSAT in [121], and two studies have also been performed for the outskirts of the cluster [156, 157]. In the centre of the cluster, a large $20 - 30\%$ soft excess was found at very high statistical significance, in addition

the excess was seen to be increasing with radius.

A study of cluster outskirts, where there is less intrinsic signal, requires a careful subtraction of the background. The first study (2002) [156] used offset pointings to measure the background and found a soft excess extending out to 2.6 Mpc. A later study (2009) was [157], which used the observation of Coma during the ROSAT all-sky survey. Although this observation was of limited temporal duration, it involves a background that is both spatially and temporally contiguous, being measured as the satellite slews across the sky. This found a soft excess in the R2 channel ($0.14 - 0.28$ keV) extending up to 5 Mpc from the cluster centre.

A665

ROSAT observations of the galaxy cluster Abell 665 in [121] found no excess. Indeed the inner part of the cluster seems to show a deficit of soft X-rays which is taken as a sign of some cooler gas resulting in absorption of the soft X-ray component of the ICM. Observations at larger radii show results consistent with no excess. No other analyses of observations of A665 have been performed, that we are aware of.

A2199

The second galaxy cluster we consider is Abell 2199. Many analyses of the soft X-ray excess in A2199 have been performed. An excess in A2199 has been observed with three satellites: the EUVE satellite, and the X-ray satellites BeppoSAX and ROSAT [129, 130, 134, 136, 158–161]. However due to controversies about the background subtraction in EUVE and BeppoSAX, the status of observations of the soft excess in A2199 with these satellites is unclear.

Averaging over the whole of the cluster, the ROSAT analysis of [121] found no compelling evidence for a soft X-ray excess. Excluding a large deficit at < 1 arcminute, there is small overall excess, however due to the large scatter in the data it has very low statistical significance.

A2255

Observations of Abell 2255 in [121] revealed an excess, though averaged over the cluster it is still at low significance. Although the excess is not as statistically significant as the Coma cluster, an excess of $\sim 20\%$ was seen between 1 and 6 arcminutes from the cluster centre. An even larger excess of $\sim 50\%$ was seen between 9 and 15 arcminutes, though the total signal (thermal emission from the hot ICM plus any excess emission) over background in this region was too poor to draw conclusive results and were thus not analysed further. This is the only analyses of the soft X-ray excess that has been performed on A2255 that we know of.

5.3.2.1 Summary

In summary, soft excess emission has consistently been found with many different satellites across many different clusters. While calibration and background subtraction is difficult, the consistency between different instruments with very different sources of systematic error strongly suggests that the effect is real and not an instrumental artefact. The precise energy range of the excess is unclear. Some XMM-Newton studies suggest that the excess reaches up to 1 keV, but this is also a satellite where background and calibration issues are more pronounced.

5.3.3 Astrophysical Models of the Soft Excess

Two main astrophysical models have been proposed to explain the soft excess. The first is a 'warm' gas, which either coexists with the hot gas of the ICM, or is located at the outskirts of the clusters as a filamentary inter-galactic medium. The soft excess then arises from thermal bremsstrahlung from this warm gas. The second scenario involves the non-thermal generation of the soft excess from inverse Compton scattering of relativistic electrons on the CMB. Here I review both proposed explanations.

5.3.3.1 A Warm Gas

Thermal emission from a 'warm', $T \sim 0.5\text{–}1 \times 10^6\,\mathrm{K}$ [118–120], gas can come in two forms. Either as an extra cooler component of the ICM [118–120], or from a warm hot intergalactic medium (WHIM) located at the outskirts of the cluster [162–164]. Such a WHIM is thought to be a filamentary structure linking galaxy clusters, and explains the missing baryons problem [165]. The main problem in the first case is cooling of the gas. To maintain pressure balance between the warm and hot components of the ICM, the lower temperature gas must have higher electron densities, increasing the rate of collisional cooling and resulting in a gas that cools in a shorter time than the typical age of the cluster [166], thus without cooling flows one would not expect to see this warm component of the gas. In the second case, it seems possible to explain the soft X-ray excess at the outskirts of clusters, where the filaments typically reside, but the central excesses are harder to explain. Explaining the soft excess here requires very long filaments aligned along the line of sight to each cluster where a central excess is observed [164].

A second problem is that a warm gas should also have thermal emission lines associated to it. However searches for lines have proved null [146–149, 167, 168], and claimed detections [143, 144, 151] have subsequently been challenged [152–154].

While these two problems render the warm gas proposal at the cluster centre very problematic, there remains a possibility it could be associated to a filamentary WHIM at the outskirts of the cluster.

5.3.3.2 Inverse Compton Scattering

In the second scenario, proposed in [169–172], the soft excess arises from inverse Compton up-scattering of CMB photons with a relativistic electron population, hence abbreviated as IC-CMB. Up-scattering CMB photons to the soft X-ray regime requires a population of non-thermal electrons with

$$\gamma \sim 500 \left(\frac{E_{\mathrm{excess}}}{200\,\mathrm{eV}} \right)^{\frac{1}{2}} .$$

Such electrons have presumably been produced by supernovae, radio galaxies or by particle creation in intra-cluster shocks, and the IC-CMB explanation therefore ties the soft excess to the non-thermal cosmic ray content of clusters.

One attractive feature of this scenario is that significant populations of non-thermal relativistic electrons are known to be present in some galaxy clusters. In Coma, for example, this is evidenced by the presence of a large radio halo which (as I will review in section 5.4) indicates the presence of non-thermal relativistic electrons with $\gamma \approx 2000$. In the first versions of the IC-CMB scenario, the electrons at $\gamma \approx 500$ were assumed to connect with the higher energy electrons at $\gamma \approx 2000$ by a simple spectral power law.

Observations of the radio emission fix the electron population's spectral index, and a combination of the magnetic field magnitude and electron number density. The magnitude of the soft excess produced by IC-CMB emission is dependent on just the number density of electrons. Thus given a simple power law model of the relativistic electrons between the soft excess and synchrotron emitting electrons, matching the observed fluxes predicts the galaxy cluster magnetic field.

Indeed, early studies of the Coma cluster [169–171,173–178] found that both emissions could be explained in this way, provided the cluster magnetic field is $B \lesssim 1\,\mu\mathrm{G}$. Alternatively, extrapolation of the radio population to the soft X-ray regime in [169] found that a fraction $\sim \left(\frac{B}{0.4\,\mu\mathrm{G}}\right)^{-2.34}$ of the soft excess in the Coma cluster could be explained by the radio population.

As we will discuss in section 5.4, observations of Faraday rotation through the Coma magnetic field imply that the magnetic field has a strength $B \sim 3 - 5\ \mu\mathrm{G}$. Thus it is not possible to explain both radio emission and soft excess emission using the same population of electrons. Instead there must be two populations of relativistic electrons, and the soft X-ray emitting electron population should have a sharp cut-off above $\gamma \sim 500$ [179, 180].

Given a single injection event, such a cut-off naturally occurs. Due to radiative losses electrons lose energy as $\frac{dE}{dt} \sim -E^2$, thus high energy electrons are naturally

degraded. However this a time-dependent cut off, as such, in order that the cut-off falls exactly in the region between the two emission regions, $\gamma \sim 500 - 1000$, the soft X-ray emitting electrons would have to have been injected at a rather specific time in the past. For Coma, this was estimated to be 1 to 1.4 billion years ago [180]. The radio emitting population must then be generated by other (weaker) injection events later than this original injection event.

We expect similar constraints on the cluster's injection histories to hold for the other clusters in which an excess has been observed. Thus it seems that the observed high magnetic fields in galaxy clusters require rather special initial conditions to generate the electron populations necessary to produce the soft X-ray and radio emission.

A further constraint on the IC-CMB explanation of the soft excess comes from gamma ray emission. Gamma ray emission arises from non-thermal bremsstrahlung from scattering of the IC-CMB electrons off the thermal proton population, and is independent of the magnetic field. Since it is physically implausible that whichever initial event accelerated the electrons did not simultaneously accelerate protons, there is in addition expected to be gamma ray emission from π_0 secondaries produced by collision of cosmic-ray protons with the thermal protons (note that at the top of the atmosphere cosmic-ray ions outnumber electrons by a factor $\approx 10^2$). So far, however, there have been no positive detections of diffuse gamma ray emission from the intra-cluster medium [181–186], despite a predicted gamma ray flux from the IC-CMB scenario above current limits [186]. While such tuned initial conditions may potentially be possible for any single cluster, it is difficult to see how the IC-CMB scenario can reconcile both the generic presence of the soft excess phenomenon and the generic absence of gamma rays from clusters at the level accessible to Fermi-LAT.

A further difficulty for the IC-CMB explanation of the soft excess is the observation of soft excess emission at large radii from the cluster centre [156, 157]. In the case of Coma, soft excess halo emission has been detected up to radii of $\approx 5\,\mathrm{Mpc}$. This emission is well beyond the radius at which the hot gas can be detected and at which

the cluster meaningfully exists. For an IC-CMB explanation of this halo, it is unclear where this necessary relativistic electron population would come from.

Overall, while it would be premature to conclude that astrophysical explanations cannot work, the above difficulties and observations motivate alternative scenarios.

5.4 Cluster Magnetic Fields

In this section I discuss magnetic fields in galaxy clusters. In section 5.4.1 I briefly review the observational methods used to infer the existence of cluster magnetic fields, and in section 5.4.2 I review the magnetic field model of [187]. Finally in section 5.4.3 I describe the numerical details of how the discretised magnetic field is produced for the simulation.

5.4.1 Magnetic Field Observations

Galaxy clusters are known to be magnetised, see [188,189] for reviews of observations and methods to measure cluster magnetic fields. It is well established that clusters feature magnetic fields of strengths $\mathcal{O}(1-10\ \mu\text{G})$, which are coherent over scales ~ 10 kpc, with a physical extent stretching over Mpc scales. Observations in radio waves have shown that a large number of clusters contain large areas of diffuse radio emission (radio halos/relics) which cannot be attributed to known sources. Such radio structures are produced by synchrotron emission of a relativistic electron population in the magnetic field of the cluster. Observations of polarised sources in and behind galaxy clusters also exhibit Faraday rotation consistent with treating the cluster as a 'Faraday screen', whose magnetic field rotates the plane of polarisation of the radio emission by making the intra-cluster medium birefringent.

Little is known about the origin of the magnetic field in galaxy clusters, consensus is that gas motions during cluster formation amplify some seed field by many orders of magnitude. There is thus two open questions, what is the origin of this seed field, and what mechanism amplifies it to the $\sim \mu$G strengths we see in clusters today. It is thought that amplification happens through turbulent gas motions during

cluster formation which acts as a dynamo explosively amplifying the seed field, see e.g. [190, 191]. This seed field is expected to still be visible as the cosmological (or inter-galactic) magnetic field (IGMF), for a review of the cosmological magnetic field see e.g. [192]. The cosmic magnetic field strength is bounded both from above and from below, $B \sim 10^{-16} - 10^{-9}$ G [193, 194]. There are two main mechanisms which are thought to be able to produce such a field: early universe magnetogenesis, i.e. magnetic field production during inflation or a phase transition; or from outflows from astrophysical bodies, such as supernovae, active galaxies, starburst regions, and galactic winds prior to cluster formation [195].

These two mechanisms lead to different magnetic field structures. In the first case, the primordial field on cluster and galactic scales will go through turbulent, cascade decay, and will thus develop a turbulent, Kolmogorov power spectrum, the evolution of these fields have been studied in e.g. [196–198]. In the outflow case this will not be the case and the field is unlikely to have a Kolmogorov power spectrum, thus upon measurement of the IGMF, it should be possible to tell these two situations apart. In the first case the cosmological magnetic field will be pervasive throughout the whole inter-galactic medium (IGM). In the second case, some fraction of IGM will contain such a magnetic field, depending on how efficiently such magnetic fields are transported through the inter-stellar voids. Indeed, coming with the lower bound on the magnetic field strength, there is also a lower bound on the filling factor of the IGMF [199]. The cosmological magnetic field thus should fill $\sim 50\%$ of voids, which is hard to explain in the astrophysical outflow origin for the field, and favours a primordial origin.

The magnetic field in galaxy clusters has however gone through amplification and the nature of the seed field is less important compared to the amplification process. Magneto-hydrodynamical simulations hint that the cluster field is independent of the initial conditions of the seed field, except for the field magnitude [200, 201]. In the very outskirts of clusters and in particular in inter-stellar voids where there has been no processing during cluster formation, the field will remain in its original seed

form. However in the central parts of the clusters (which we consider) where some amplification has occurred, the resulting cluster magnetic field structure is likely to be very different from its cosmological form.

Detailed measurement of the magnetic field in clusters is challenging. Synchrotron radio emission is dependent on both the strength of the magnetic field in the cluster, and the size of the electron population. The emission is degenerate in these parameters, and, with no other information about the electron population, the magnetic field strength cannot be measured. An estimate can be made through 'equipartition' arguments. This is a minimum energy criterion that the relativistic electron and magnetic field energy densities should be the same, however this assumption is not easily verifiable. This assumption was used in [202] who found the Coma magnetic field, averaged over the central $1\,\mathrm{Mpc}^3$, to be $B \sim 0.7 - 1.9\,\mu\mathrm{G}$.

The relativistic electrons in the clusters will inverse Compton scatter off CMB photons to produce hard X-ray or gamma ray photons. The lack of observation of these hard X-ray photons in Coma bounds the electron population, and in doing so has led to lower bounds on the magnetic field strength. This bound is $B > 0.2\,\mu\mathrm{G}$ for the Coma radio halo, and $B > 1\,\mu\mathrm{G}$ for the Coma radio relic [203].

The effect of Faraday rotation of polarised light in a magnetic medium is described by a 'rotation measure'

$$\Psi_{\mathrm{obs}}(\lambda) = \Psi_0 + \lambda^2\,\mathrm{RM}\,. \tag{5.6}$$

The rotation measure is given by the line-of-sight integral of the parallel component of the magnetic field multiplied by the electron density,

$$\mathrm{RM} = \frac{e^3}{2\pi m_e^2}\int_{\mathrm{l.o.s.}} n_e(l)B_{\parallel}(l)dl\,, \tag{5.7}$$

where, by convention, a magnetic field pointing towards the observer gives rise to a positive RM.

Since the electron density distribution may be inferred from X-ray observations, studies of rotation measures provide a sensitive probe of the magnitude of the cluster magnetic field. By measuring Ψ_{obs} at several frequencies for a given radio source, the

value of RM may be inferred.

The Faraday rotation observations are limited by the number of polarised radio sources behind an individual cluster and only depend on the integral of the parallel component of magnetic field along the line of sight. If one makes an assumption that the magnitude of the magnetic field is constant along many single-sized domains, the variance of the Faraday rotation measures gives the size of these domains, which is a hint to the typical coherence length of the field along the line of sight. Maps of Faraday rotation exhibit a patchy structure which similarly give hints to the size of the coherence lengths.

More in-depth analyses involve simulating a magnetic field, given a specific field model, and using the observed radio halo and Faraday rotation data to constrain the parameters of the model. Using the power spectrum of the Faraday rotation is another method used to constrain the magnetic field power spectrum. However, it has been argued that the use of Faraday rotation measurements does not take into account the intrinsic rotation in the source of the polarised signal and thus is ill-suited to constraining the cluster-wide magnetic field [204].

In this study we assume the cluster magnetic fields are well represented by the model first proposed in [123]. The simulated magnetic field is a stochastically-generated, Gaussianly-distributed field that is tangled over a range of scales, and has a power-law power spectrum. The magnitude of the magnetic field is assumed to decrease with radius as some power of the thermal ICM electron density. Such fields have been shown to be good fits to Faraday rotation measurements and synchrotron observations in many clusters [123, 187, 205–208]. However, we caution that it is an idealised model and thus may not capture the full details of the magnetic fields in these clusters.

5.4.2 Magnetic Field Model

Here we review the magnetic field model used for the clusters we study.

We start by randomly drawing the Fourier-space vector potential components.

The magnitude of each component of the vector potential is drawn randomly from a Rayleigh distribution with power law scaling $|A_k|^2 \sim k^{-n}$, where we chose the normalisation factor to be one,

$$p(\tilde{A}_k) = \frac{\tilde{A}_k}{|k|^{-n}} \exp\left(-\frac{\tilde{A}_k^2}{2|k|^{-n}}\right), \tag{5.8}$$

and the complex phase of each component is drawn randomly from a uniform distribution. We then choose the correct normalisation of the real space magnetic field later. We compute the Fourier-space magnetic field as $\tilde{B}(k) = ik \times \tilde{A}(k)$. Then we Fourier transform this to position space to generate a Gaussianly-distributed, isotropic, tangled, divergence-free magnetic field, with power spectrum

$$4\pi k^2 |B_k|^2 \sim k^{-\zeta}, \tag{5.9}$$

where $\zeta = n - 4$. The Fourier vector potential is set to zero outside of the momentum range $k_{\rm min}$ to $k_{\rm max}$, corresponding to physical scales $\Lambda_{\rm max} = 2\pi/k_{\rm min}$ and $\Lambda_{\rm min} = 2\pi/k_{\rm max}$. There is observational evidence that the magnetic field magnitude falls as a function of radius from the cluster centre. To emulate this, the magnetic field is modulated by multiplying by some function $f(n_e) \propto (n_e)^\eta$, and normalised such that the core of the cluster has a magnetic field of magnitude B_0, such that

$$B(r) = B_0 \left(\frac{n_e(r)}{n_0}\right)^\eta. \tag{5.10}$$

We take time here to point out there are two interesting values for η: the first is the case where the magnetic field energy density ($\varepsilon_B \propto B^2$) falls with the electron energy density, i.e. $B^2 \propto n_e$ and $\eta = 0.5$; the second is the case where the magnetic field lines are 'frozen-in' to the plasma,[1] which corresponds to the case $\eta = 2/3$.

The electron/gas density of clusters follows an isothermal β-model profile

$$n_e(r) = n_0 \left(1 + \frac{r^2}{r_c^2}\right)^{-\frac{3}{2}\beta}, \tag{5.11}$$

[1]Briefly, the magneto-hydrodynamical ICM is in a regime where the magnetic Reynolds number is large (corresponding to turbulence). In this regime the magnetic field lines 'move with the plasma', specifically this regime corresponds to $B \cdot A = $ const (with A the cross-sectional area) and hence $B \propto n_e^{2/3}$.

with n_0 the central electron density and r_c the core radius of the cluster. In many clusters the core has cooled (such as A2199), leading to a spike in the electron density in the centre [209]. The electron density profile in these cool-core clusters then usually follows a double β-model, which is the sum of two β-models,

$$n_e(r) = n_0 \left(1 + \frac{r^2}{r_c^2}\right)^{-\frac{3}{2}\beta} + n_{0,\mathrm{cool}} \left(1 + \frac{r^2}{r_{c,\mathrm{cool}}^2}\right)^{-\frac{3}{2}\beta_{\mathrm{cool}}}, \tag{5.12}$$

with $r_{c,\mathrm{cool}} \ll r_c$ and typically $n_{0,\mathrm{cool}} \geq n_0$.

The magnetic field model then has 5 parameters: the power spectrum index ζ, the maximum and minimum lengthscales the magnetic field is 'tangled' over Λ_{max} and Λ_{min},[2] the radial scaling of the magnetic field η, and the magnetic field magnitude in the centre of the cluster B_0. These parameters are constrained using Faraday rotation maps or synchrotron radio halos.

The natural expectation value for the power law index parameter is $\zeta = 5/3$, corresponding to a Kolmogorov turbulent spectrum. Observations have hinted that the ICM is itself turbulent, thus a Kolmogorov index is the theoretically favoured value for the power spectrum index. However the amplification of the seed field, the motions of the gas, and the pollution from expulsion of magnetic fields from galactic environments could all change this spectrum and lead to different values for the power law index. Hence these power law indices are fit using observational data in all cases as we will see later, except for A665, where the theoretically favoured value of $\zeta = 5/3$ is chosen.

Coma

The magnetic field in the centre of the Coma cluster was studied using this model in [187], using the rotation maps of several polarised radio sources, located in or behind the cluster. For Coma, the beta-model parameters have been experimentally determined to be $n_0 = 3.44 \times 10^{-3}\mathrm{cm}^{-3}$, $r_c = 291\,\mathrm{kpc}$, and $\beta = 0.75$ [210]. Mock Faraday rotation maps were produced, scanning over magnetic field parameters.

[2]Note that the scale Λ corresponds to the wavelength of a full period of magnetic field 'oscillation', the field will thus be coherent over a scale $\lesssim \Lambda/2$.

It is important to note that while the parameters ζ, $\Lambda_{\rm min}$, $\Lambda_{\rm max}$, η and B_0 are free parameters of the magnetic field model which can be constrained by fitting to rotation measures [187], there is an effective degeneracy between ζ and $\Lambda_{\rm max}$. Larger ζ can be compensated for by lowering $\Lambda_{\rm max}$ and vice versa, giving an equally good fit when comparing to RMs from Faraday rotation. The value $\zeta = 5/3$ corresponds to a Kolmogorov-like turbulent power-law slope for the one-dimensional power spectrum of the magnetic field and was the headline value adopted in [187]. This power-law slope corresponds to a best fit value of $\Lambda_{\rm max} = 34\,{\rm kpc}$, with $\Lambda_{\rm min}$ found to be $2\,{\rm kpc}$. However, as discussed in [187], the Faraday rotation measurements are degenerate along a curve in $(\zeta, \Lambda_{\rm max})$ space, with a flatter spectrum as $\Lambda_{\rm max}$ is increased. Equally good fits to Faraday rotation measures are provided by a flat one-dimensional power spectrum, i.e. $\zeta = 0$, with $\Lambda_{\rm max}$ increased to $100\,{\rm kpc}$. These spectra have more power on small scales compared to the Kolmogorov spectrum.

Due to these degeneracies, we will pick two parameter sets that fit the Faraday rotation data, to simulate the Coma magnetic field. These are called Models 1 and 3 in table 5.1. We will also find it illustrative to include a third, less physically-motivated magnetic field, Model 2, when discussing ALP–photon conversion probabilities, thus we also put these model parameters in table 5.1.

Let us note that these values are quite consistent with other methods used to determine the magnetic field in the Coma cluster. Equipartition arguments have been used to estimate the magnetic field averaged over the central $1\,{\rm Mpc}^3$, to be $B \sim 0.7 - 1.9\,\mu{\rm G}$ [202]. Faraday rotation, assuming the magnetic field to be comprised of cells of constant magnetic field, but pointing in different directions, gives an estimate $B \approx 2\,\mu{\rm G}$, with a tangling scale of $13 - 40\,{\rm kpc}$ [211].

A665

Here I summarise the method and the results of [206] in constraining the above model parameters in the A665 galaxy cluster, which we show in table 5.2. The parameters for A665 are constrained by simulating mock radio halo images, upon assuming the

	Model 1	Model 2	Model 3
$\Lambda_{\rm min}$	2 kpc	2 kpc	2 kpc
$\Lambda_{\rm max}$	34 kpc	5 kpc	100 kpc
ζ	5/3	5/3	0
B_0	$3.9 - 5.4\,\mu$G	$5.4\,\mu$G	$5.4\,\mu$G
η	$0.4 - 0.7$	0.7	0.7

Table 5.1: The parameter values for three parameter choices for the Coma magnetic field. The first (Model 1) is a Kolmogorov spectrum that fits the Faraday rotation data. The second (Model 2) is designed to show the effect of concentrating all power on small scales but is not a fit to Faraday rotation data. The third (Model 3) is a flat spectrum (in k-space) that fits Faraday rotation data and has most power on small scales.

form of the relativistic electron population, and comparing to the observed radio halo. The analysis proceeds in two steps: first, equipartition arguments are used to determine the radial scaling of the magnetic field, the parameter η. Secondly mock radio halos are produced and compared to data to constrain the field strength and power spectrum.

In the first step the electron population is taken as a power law where the spectral index of the electron population is related to the spectral index of the observed radio halo. The number of relativistic electrons in the cluster is set point-by-point such that at that point there is equipartition between the relativistic electrons and magnetic field. We note that the spatial form for the relativistic electron population is an assumption, since little is known about it or how it is produced. Equipartition is the assumption that the energy stored in the magnetic field and in the cosmic ray electrons is equal. The radial profile of this spherically-symmetric radio halo is compared to the radial profile of the observed radio halo, giving a best fit value of $\eta = 0.47 \pm 0.03$, which we note is very close to the interesting value mentioned earlier, where the energy density of the magnetic field falls proportional to the thermal electron energy density.

A full 3D magnetic field is produced using the method described earlier. The same form for the electron population as used in the first step is used to simulate mock radio halos for this field. The predicted radio emission from the idealised, spherically-symmetric radio model used in the first step is subtracted from both the observed radio halo and the simulated radio halo. The RMS of the residuals for the simulated

halos are then plotted as a function of $\Lambda_{\rm max}$, and show a clear trend that they increase with $\Lambda_{\rm max}$. The simulated RMS matches the observed radio halo's for $\Lambda_{\rm max} \sim 450$ kpc.

On top of this there is a more qualitative result that there is a clear change in shape of the radio halo when going from small to large $\Lambda_{\rm max}$, at larger $\Lambda_{\rm max}$ the halo becomes more anisotropic and also there is a clear separation between the centre of the ICM plasma and the centre inferred from radio observations. One can plot this offset distance for the different values of $\Lambda_{\rm max}$, there is a clear trend for larger offsets for larger $\Lambda_{\rm max}$ values. The value where the simulations match the offset of the observations is slightly higher, at roughly 500 kpc.

We note that there are a lot of assumptions made in this analysis, especially about the electron population used to produce the mock radio halos. It is not clear whether these assumptions hold, thus there is likely to be a large systematic uncertainty on these magnetic field parameters.

A2199

The parameters of the model for the magnetic field of A2199 have been constrained by producing mock Faraday rotation images and comparing to the observed Faraday rotation in a radio source located at the centre of the cluster [207]. All Faraday rotation data is contained within 20 kpc of the centre of the cluster, thus in our analysis we are extending the magnetic field beyond the range which it is known to be valid. The X-ray brightness of the cluster is not completely spherically-symmetric, it contains two elliptical regions to the east and the west of the cluster where the electron density drops significantly. The simulations of [207] showed that there was very little effect on results when they were included as when they were left out. Since we are averaging the probabilities over concentric annuli, and since the features are located within 20 kpc and we simulate out to > 500 kpc, we do not expect these regions to change the analysis in any appreciable way, we thus do not include these elliptical areas in our simulation.

A2199 is a cool-core cluster and thus the gas density follows a double β-model, see equation 5.12. The electron density in the cool-core is ~ 30 times that in Coma and A665 for instance. The cool-core electron density spike drops rapidly beyond 9 kpc, whereas the normal ICM electron population has a core radius of 26 kpc. Again since we simulate out to ~ 400 kpc we do not expect the cool-core to have a big effect on the results.

The magnetic field was simulated scanning over all five parameters of the model, using the power spectrum of the Faraday rotation to constrain the magnetic field power spectrum parameters, and maps of Faraday rotation to constrain the radial scaling. The parameters of the magnetic field of A2199 are not well constrained from this analysis. The parameters B_0 and η are degenerate with respect to Faraday rotation measurements, i.e. the same average magnetic field can be obtained by reducing (increasing) B_0 whilst simultaneously reducing (increasing) η such that the field falls off less (more) rapidly with radius. The resulting values of the central magnetic field and radial parameters are $B_0 = 11.9 \pm 9$ and $\eta = 0.9 \pm 0.5$. There is a similar degeneracy between the maximum lengthscale the field is tangled over and the power spectrum of the field, large spectral indices put more power at large scales, increasing $\Lambda_{\rm max}$ has the same effect. Thus again these parameters are poorly constrained, although we keep these parameters fixed at their central values in these simulations. We list these parameters in table 5.2.

A2255

The magnetic field parameters in A2255 are constrained in [205] (listed in table 5.2), using polarised emission from three galaxies. In this cluster a synchrotron radio halo is simulated, as well as mock Faraday rotation maps. The radial scaling of the magnetic field is constrained in two ways, firstly the root mean square of the Faraday rotation measures in each galaxy should be proportional to the X-ray brightness of the ICM, to some power. This poorly constrains η to anywhere between -0.5 and 1. Secondly there is the observation that the X-ray brightness profile follows the radio

	A665 [206]	A2199 [207]		A2255 [205]	
		Cool-core	Hot ICM	Inner	Outer
r_c (kpc)	340	9	25	432	
β	0.76	1.5	0.39	0.74	
n_0 (10^{-3} cm^{-3})	3.44	74	27	2.20	
ζ	5/3	0.6		0	2
$\Lambda_{\rm min}$ (kpc)	4	0.7		4	64
$\Lambda_{\rm max}$ (kpc)	450	35		1000	1000
B_0 (μG)	1.3	$6.0 - 11.9$		2.5	
η	0.5	$0.5 - 0.9$		$0.5 - 0.7$	

Table 5.2: Magnetic field and β-model parameters for the three clusters considered. The A2255 Inner and Outer columns refer to the magnetic field within and without the core radius respectively. The A2199 cool-core and hot ICM columns refer to the parameters for the two β-models which make up the cool-core and the normal hot ICM respectively.

halo brightness profile thus it is assumed that the populations of thermal and non-thermal electrons, and hence also the magnetic field (by equipartition), follow the same profile. Thus it is assumed in their simulations that $n_e \propto B^2$. In these simulations we will consider two values: $\eta = 0.5$ and $\eta = 0.7$, which (roughly) correspond to the two interesting cases mentioned earlier.

It was shown that the Faraday rotation is best fit for a magnetic field that has increasing power spectrum index as a function of radius. The best fit global magnetic field is one in which the magnetic field has $\zeta = 0$ for $r < r_c$ and then an exponential fall off beyond r_c, and then a non-Gaussian $\zeta = 2$ magnetic field with larger $\Lambda_{\rm min}$ for radii beyond the core radius. In our simulations we use this model, but assume both parts of the magnetic field follow a Gaussian distribution.

This choice of a rising spectral index with radius results in the magnetic field coherence lengths increasing with radius. This sort of behaviour is expected on general theoretical grounds: as the electron density decreases, the typical lengthscale of the problem $\propto 1/n_e^{1/3}$ gets larger. Such behaviour of the magnetic field was considered for ALP$-$photon conversion in the outskirts of Coma in [212]. Though the model of the field as a 'stitching' together of two fields with differing spectral index with a smoothing function between them is unrealistic, it is nevertheless a reasonable test model to study the CAB in A2255.

5.4.3 Magnetic Field Generation

Following the detailed prescription reviewed in section 5.4.2, we generate numerical models of the clusters' magnetic fields. We do this in C++. For Coma we generate a 2000^3, for the other three clusters we generate a 2000×1000^2 grid, with differing unit cell sizes. This grid size is chosen to save computer time, as spherical symmetry means that it is sufficient to consider one quarter of the plane perpendicular to the ALP–photon direction only.

The cell size should be chosen so as to evade the Nyquist criterion that the sampling rate of a dataset must be greater (ideally much greater) than twice the frequency of the dataset is satisfied for fields with structure only on scales larger than $\Lambda_{min} > 2s$. We note that such a small unit cell size places a limit on the size of the field we can generate. For Coma the cell size is $s = 0.5$ kpc, resulting in a simulation of size 1 Mpc.

As outlined in section 5.4.2, the values of the Fourier coefficients of the vector potential are generated randomly for all modes in the range $\Lambda_{\rm min}$ to $\Lambda_{\rm max}$. After computing the momentum space magnetic field, the real space representation is obtained by performing a discrete Fourier transform using FFTW 3.3.3 routines [213].

The real-space magnetic field is modulated so as to exhibit attenuation over cluster scales. For Coma, as in [187] the normalisation of the magnetic field is chosen so that the average magnitude of the magnetic field within the core radius, r_c, of the cluster is equal to the parameter B_0. For the other clusters, the magnetic field normalisation is chosen slightly differently. In those cases the magnetic field is normalised such that the average magnetic field across the simulation region, after it has been modulated, is equal to some value B_{ave}.

5.5 ALP–photon Propagation

In this section we solve the ALP–photon equation of motion. I start in section 5.5.1 by solving for a homogeneous magnetic field, or a single magnetic field domain, and then

in section 5.5.2 solve for an inhomogeneous magnetic field applicable to the magnetic field we simulate.

5.5.1 Homogeneous Solution

As we saw in 3.3.1, the presence of the term $aF_{\mu\nu}\tilde{F}^{\mu\nu} \sim a\vec{E}\cdot\vec{B}$ leads to the coupled ALP–photon equation of motion [37]

$$\left(\omega + \begin{pmatrix} \Delta_\gamma & 0 & \Delta_{\gamma ax} \\ 0 & \Delta_\gamma & \Delta_{\gamma ay} \\ \Delta_{\gamma ax} & \Delta_{\gamma ay} & \Delta_a \end{pmatrix} - i\partial_z\right)\begin{pmatrix} |\gamma_x\rangle \\ |\gamma_y\rangle \\ |a\rangle \end{pmatrix} = 0\,. \tag{5.13}$$

Here ω denotes the energy of the photon and ALP modes. We have also set to zero the mixing between photon polarisation states caused the magnetic field, which represents Faraday rotation, since it is unimportant for X-ray energies.

The refractive index for photons in the plasma is given by $\Delta_\gamma = -\omega_{\rm pl}^2/2\omega$, where $\omega_{\rm pl} = \sqrt{\frac{4\pi\alpha n_e}{m_e}}$ denotes the plasma frequency of the ICM, and is an effective mass for the photon. The ALP–photon mixing is induced by the matrix element $\Delta_{\gamma ai} = B_i/2M$. The mass of the ALP determines the final diagonal matrix element: $\Delta_a = -m_a^2/2\omega$.

We can write the solution to equation (5.13) as

$$\begin{pmatrix} |\gamma_x\rangle \\ |\gamma_y\rangle \\ |a\rangle \end{pmatrix}(L) = \exp\left(-i\int_0^L \mathcal{M}(z)\mathrm{d}z\right)\begin{pmatrix} |\gamma_x\rangle \\ |\gamma_y\rangle \\ |a\rangle \end{pmatrix}(0)\,, \tag{5.14}$$

with

$$\mathcal{M}(z) = \begin{pmatrix} \Delta_\gamma(z) & 0 & \Delta_{\gamma ax}(z) \\ 0 & \Delta_\gamma(z) & \Delta_{\gamma ay}(z) \\ \Delta_{\gamma ax}(z) & \Delta_{\gamma ay}(z) & \Delta_a(z) \end{pmatrix}, \tag{5.15}$$

and we have left off the overall phase factor $\exp(-i\omega L)$.

In section 5.8.2 we will describe the results of numerically integrating equation (5.13) for the inhomogeneous magnetic field discussed in section 5.4.3. However, it is illuminating to first consider the simpler case of a homogeneous electron density

and magnetic field in some domain of size L. Furthermore, since only photons with polarisation parallel to the magnetic field couple to ALPs, a simple rotation in the x-y plane reduces the non-trivial part of the problem to that of a 2-body system of $\left|\gamma_{\parallel}\right\rangle$ and $|a\rangle$. The non-trivial part of the z-evolution generator $\mathcal{M}$ can then be diagonalised by a rotation of angle θ satisfying

$$\tan(2\theta) = \frac{2\Delta_{a\gamma}}{\Delta_a - \Delta_\gamma}\,, \tag{5.16}$$

where now $\Delta_{a\gamma} = B/2M$. Thus, in a single domain with a homogeneous magnetic field, the probability that an ALP converts into a photon is given by

$$P(a \to \gamma) = \sin^2(2\theta)\sin^2\left(\frac{\Delta}{\cos 2\theta}\right), \tag{5.17}$$

where $\tan 2\theta = \frac{2B_\perp \omega}{M m_{\rm eff}^2}$, $\Delta = \frac{m_{\rm eff}^2 L}{4\omega}$ and $m_{\rm eff}^2 = m_a^2 - \omega_{\rm pl}^2$.

For a single domain with a coherent magnetic field, the ALP-photon conversion probability is completely determined by the angles θ and Δ,

$$\theta \approx \frac{B_\perp \omega}{M m_{eff}^2} = 5.6 \cdot 10^{-4} \left(\frac{10^{-3}\ \mathrm{cm}^{-3}}{n_e}\right)\left(\frac{B_\perp}{2\,\mu\mathrm{G}}\right)\left(\frac{\omega}{200\ \mathrm{eV}}\right)\left(\frac{10^{13}\ \mathrm{GeV}}{M}\right), \tag{5.18}$$

$$\Delta = 2.7 \left(\frac{n_e}{10^{-3}\ \mathrm{cm}^{-3}}\right)\left(\frac{200\ \mathrm{eV}}{\omega}\right)\left(\frac{L}{10\,\mathrm{kpc}}\right), \tag{5.19}$$

where we have taken $m_a = 0$. To get an estimate of the conversion probabilities, we take an approximation known as the small-angle regime $\theta \ll 1$, $\Delta \ll 1$, which is always true for θ, but not true for Δ for large values of n_e (and hence small radii), large L, or low ALP energies. Then the probability an ALP will convert to a photon over D/L domains, with D the full cluster size is simply

$$P(a \to \gamma) = 0.9 \cdot 10^{-3} \left(\frac{D}{1\ \mathrm{Mpc}}\frac{L}{10\ \mathrm{kpc}}\right)\left(\frac{B_\perp}{2\,\mu\mathrm{G}}\frac{10^{13}\ \mathrm{GeV}}{M}\right)^2. \tag{5.20}$$

We thus see that clusters can be extremely efficient at ALP$-$photon conversion.

While the above discussion is illuminating, it can only give us an order of magnitude estimate for the typical conversion probability in clusters. A constantly varying, turbulent, multi-scale magnetic field requires a full numerical calculation to calculate the conversion probabilities. However, we will see that the dependence on radius of

the conversion probabilities can be well understood by looking at the single-domain formula. Typically there are two regimes, one where the angle Δ is large (low ALP energies, or high electron densities), and one where it is small and the full small angle regime formula can be used. In the large-Δ regime

$$P(a \to \gamma) \propto \frac{\omega^2}{M^2} \left(\frac{B(r)}{n_e(r)} \right)^2 \left(\frac{D}{L} \right). \tag{5.21}$$

Thus in this regime the conversion probabilities are energy dependent, and behave with radius as $P \propto {n_e}^{2(\eta-1)}$. This will (for $0 < \eta < 1$) produce conversion probabilities which increase with radius. For small-Δ, we see in equation 5.20 that the probability will fall with radius, $P \propto {n_e}^{2\eta}$, and that the conversion is no longer dependent on the energy.

5.5.2 Inhomogeneous Magnetic Fields

The ALP–photon conversion probabilities are computed by numerically simulating the propagation of an ALP through the discretised magnetic field model discussed in section 5.4.3. Since the lattice spacing is always much smaller than the cluster radius r_c, the electron density is slowly varying over each zone of constant magnetic field and may consistently be approximated as constant within each lattice zone. Thus within each zone the unitary 'z-evolution' matrix is constant, and equation (5.14) can be solved recursively from the nth lattice point to the next as

$$\begin{pmatrix} |\gamma_x\rangle \\ |\gamma_y\rangle \\ |a\rangle \end{pmatrix}_{n+1} = \exp\left(-i \begin{pmatrix} \Delta_{\gamma,\, n} & 0 & \Delta_{\gamma ax,\, n} \\ 0 & \Delta_{\gamma,\, n} & \Delta_{\gamma ay,\, n} \\ \Delta_{\gamma ax,\, n} & \Delta_{\gamma ay,\, n} & \Delta_{a,\, n} \end{pmatrix} s \right) \begin{pmatrix} |\gamma_x\rangle \\ |\gamma_y\rangle \\ |a\rangle \end{pmatrix}_n , \tag{5.22}$$

where we have again denoted the lattice spacing by s. This way, an initial pure ALP state will develop non-vanishing photon components as the particle propagates through the cluster.

The solution to equation (5.22) is obtained, just as in the single-domain case, by first rotating to a basis in which the magnetic field is aligned with one of the coordinate

axes and then diagonalising the non-trivial part of the z-propagation generator. This way, the ALP–photon propagation can be solved exactly for each lattice point.

The propagation of the full 3-body system through the lattice is then achieved recursively by diagonalising, propagating the new fields to the next grid point, and finally rotating back to obtain the state with respect to the original reference basis. Thus the state at the $(n+1)$th step is given by

$$|n+1\rangle = U_{1,n}^{\mathrm{T}} U_{2,n}^{\mathrm{T}} \mathcal{M}_n U_{2,n} U_{1,n} |n\rangle \ , \tag{5.23}$$

where $U_{1,n}$ denotes the rotation required to align a coordinate axis with the local magnetic field direction, and $U_{2,n}$ denotes the diagonalisation of the unitary z-evolution operator at the nth step.

The probability of the ALP converting into a photon is then computed as the sum of the squares of the $|\gamma_x\rangle$ and $|\gamma_y\rangle$ components of the final state. This procedure is done for each of the grid points in the x-y plane as a function of ALP energy, and for a vanishing ALP mass. By setting the ALP mass to zero we make the approximation that m_{eff}^2 is dominated by the plasma frequency. The plasma frequency is given as

$$\omega_{\mathrm{pl}} = 1.2 \times 10^{-12} \sqrt{\frac{n_e}{10^{-3}\mathrm{cm}^{-3}}}\,\mathrm{eV} \tag{5.24}$$

and is never much less than $\approx 10^{-12}\,\mathrm{eV}$ in the systems we are considering. ALP masses $m_a \leq 10^{-13}\,\mathrm{eV}$ therefore behave equivalently to a vanishing ALP mass. For ALP masses $m \gg 10^{-12}\,\mathrm{eV}$ (including the case of a QCD axion), $\theta \sim 1/m_a^2$ and $\Delta \sim m_a^2$, so that it is reasonable to expect that the conversion probabilities become suppressed relative to those we have obtained by a factor of $(10^{-13}\,\mathrm{eV}/m_a)^4$. For an ALP mass $m \approx 10^{-12}\,\mathrm{eV}$, detailed simulation would be required to study the resulting morphology and how it differs from the zero-mass case.

Note that the very low mass region applicable to this study is exactly the region well-motivated by string theory, see section 3.2.

5.6 Consistency Checks

Before going on to discussing the results in detail, we will quickly discuss some of the consistency checks of our numerical procedure.

Firstly, let us start with the computation of the conversion probabilities. Two methods were initially created to calculate conversion, both developed in Mathematica: the method described above of rotating the wavevector, and a method which used the matrix exponential function in Mathematica. These two methods agreed well with each other, with only a small error of $\mathcal{O}(0.1\%)$, which was likely due to a numerical error. Both of these conversion probabilities agreed with the single domain formula when computed along a constant magnetic field. It was then checked that the C++ implementation agreed with these. Thus we are confident in the accuracy of the conversion probabilities we calculate.

Similarly, two versions of the magnetic field code were developed independently and checked against each other to eliminate bugs in the final version. We checked that the radial scaling and normalisation of the magnetic field gave an average magnetic field ~ 2 μG, which was comparable to [187]. There are several other checks which we will come across later in the results section, but we will list now. Firstly, the best fit values of the inverse coupling from our analysis agree very well with that from the simplified analysis of [116]. That study used a constant, fixed domain model for the magnetic field, and calculated conversion probabilities in the small angle approximation, finding a best fit inverse coupling of $M \sim 10^{13}$ GeV. Second, as we have discussed, there are two distinct radial profiles, determined by the value of θ, for low energies this results in a transition at intermediate radii between increasing and decreasing conversion probabilities as a function of radius. The radius at which this happens is essentially fixed (in a constant coherence length model) by the coherence length of the magnetic field. We can calculate an effective coherence length of our magnetic field using this turning point, which we found to be ~ 10 kpc for the canonical magnetic field model in Coma. This agrees very well with both observations and

the expected coherences lengths given the distribution between $\Lambda_{\min}$ and $\Lambda_{\max}$.

Finally we checked that the random generation of the magnetic field was stable by generating multiple versions of the same magnetic field model. We calculated conversion probabilities and found that they changed little between the realisations. The largest differences were around 10%, which occurred for small inverse couplings $M \sim 10^{11}$ GeV, where the conversion probabilities were $\mathcal{O}(1)$. We note this is not a rigorous check of the numerical stability but nevertheless all the above points make us confident in our model. We will use this observed 10% fluctuation as a guide to the numerical error on our results.

5.7 Coma

This section describes the results of the detailed simulation of the Coma cluster, which was published in [2]. We compute the expected luminosity from a CAB passing through the Coma magnetic field and compare to the results of [121].

However, let us start by stating why we focus on the Coma cluster in detail. There are several reasons for this. First, the soft excess in the Coma cluster is particularly well established and well studied. Second, the cluster is at a redshift such that a significant fraction of the cluster fits into the ROSAT field of view. This gives confidence that the properties of the soft excess are associated to the cluster as a whole (a contrasting example is the Virgo cluster, which is so close that the soft excess is only studied within the central 50 kpc, and so one might worry that properties of the cluster may be conflated with properties of the large dominant central galaxy M87). However the Coma cluster is still close enough that the count rate is large enough to give high statistical confidence on the soft excess. Finally, the magnetic field structure of the Coma cluster has been studied in detail. This is crucial for our purposes, as without the magnetic field structure it is impossible to calculate the $a \to \gamma$ conversion rate.

5.7.1 Luminosity Calculation

The simulation generates a 2000^2 grid of $a \to \gamma$ conversion probabilities, each representing the probability of a single ALP at energy ω, traversing Coma through a unit of cross-sectional area $(0.5\,\mathrm{kpc})^2$, converting into a photon of the same energy. We need to convert these probabilities into intrinsic source luminosities. A redshift-distance converter is in [214]. Coma is at a well-determined redshift of $z = 0.023$. We use the parameters from the magnetic field model of [187], with $H_0 = 71\,\mathrm{km\,s^{-1}}$ corresponding to 0.460 kpc per arcsecond and a luminosity distance of 98.9 Mpc. We note the soft excess analysis of [121] assumed a Hubble constant of $H_0 = 75\,\mathrm{km\,s^{-1}}$, corresponding to an angular scale of 0.434 kpc per arcsecond and a luminosity distance to Coma of 93.6 Mpc. These differences are small enough to be neglected compared to the other statistical and systematic uncertainties in the extraction of the soft excess.

The overall CAB energy density is $\rho_{\mathrm{CAB}} = \Delta N_{\mathrm{eff}} \frac{7}{8} \left(\frac{4}{11}\right)^{4/3} \rho_{\mathrm{CMB}}$. Associated with this is a CAB number density $\mathrm{d}N_a/\mathrm{d}E$ set by the spectral shape, such that

$$\int \mathrm{d}E E \frac{\mathrm{d}N_a}{dE} = \rho_{\mathrm{CAB}}\,.$$

In terms of these the intrinsic excess luminosity associated with ALP–photon conversion is given by

$$\mathcal{L}_{\mathrm{excess}} = D_{\mathrm{Coma}}^2 \int \mathrm{d}\Omega \mathrm{d}E\, E \frac{\mathrm{d}N_a}{\mathrm{d}E} c P_{a\to\gamma}(\Omega, E)\,, \tag{5.25}$$

where D_{Coma} is the physical distance to Coma and $\mathrm{d}\Omega$ is a solid angle element (measured in arcmin^2). We restrict the energy integral to the $0.2 - 0.4\,\mathrm{keV}$ range, in accordance with the luminosities calculated in [121].

5.7.2 Results

In this section I discuss the results of the numerical simulation of the conversion probabilities, and present for the first time a detailed description of the predictions of the CAB conversion scenario for the cluster soft excess.

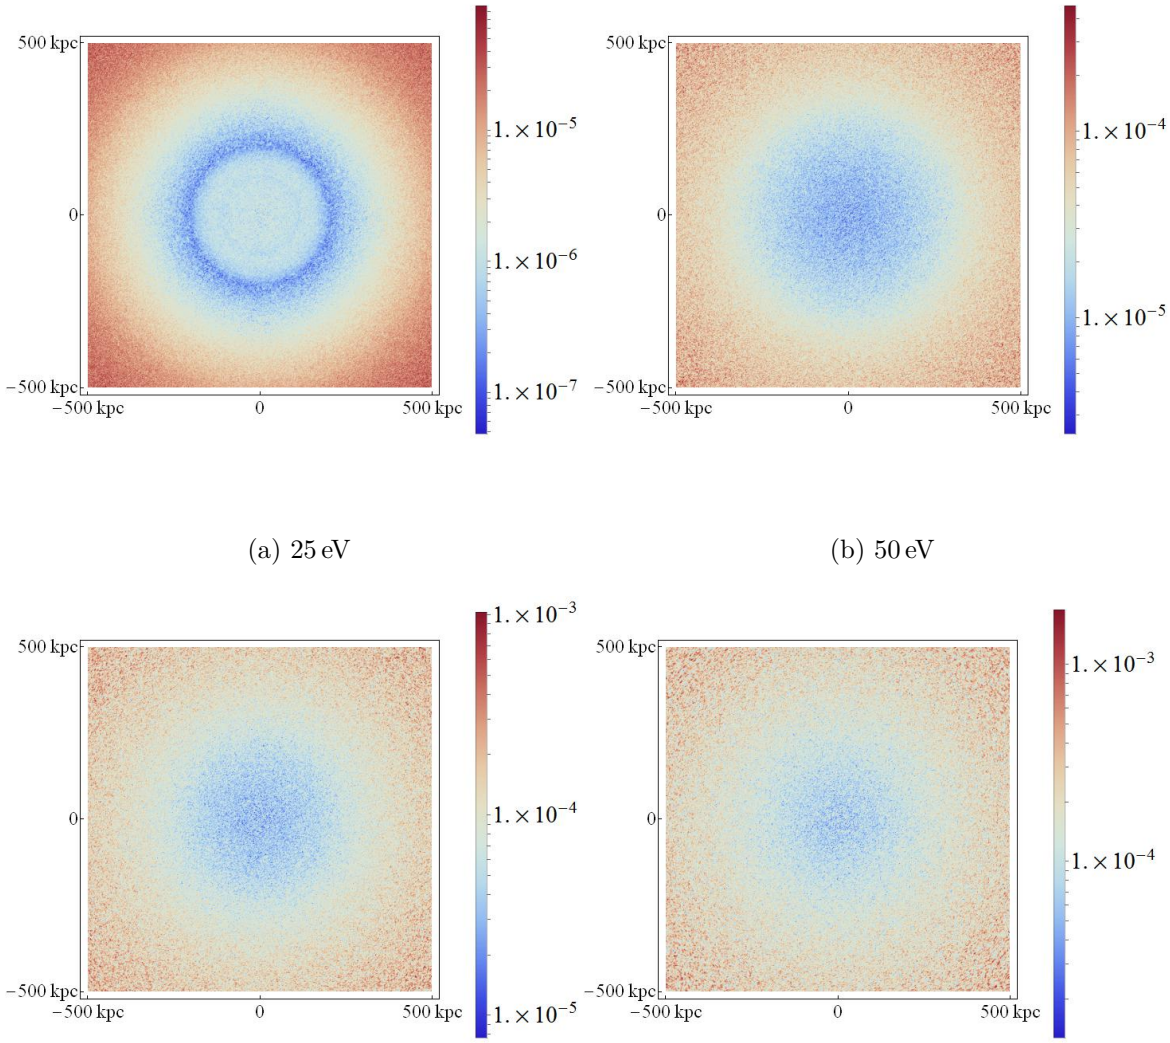

(a) 25 eV (b) 50 eV

(c) 75 eV (d) 100 eV

Figure 5.3: Conversion probabilities for energies 25 eV to 100 eV for Model 1, with $\eta = 0.5$, $B_0 = 4.7\,\mu$G and $M = 7 \times 10^{12}$ GeV.

5.7.2.1 General features of ALP–photon conversion

While several properties of the simulated conversion probabilities and soft X-ray luminosities are sensitive to the detailed magnetic field model, there are also general features that are shared by all models we have considered. In this section we highlight these properties by using Model 1 of table 5.1 as our main example.

In this model, the stochastic magnetic field is generated with a Kolmogorov power spectrum ($\zeta = 5/3$), with coherence lengths in the range $2-34$ kpc. The best-fit values

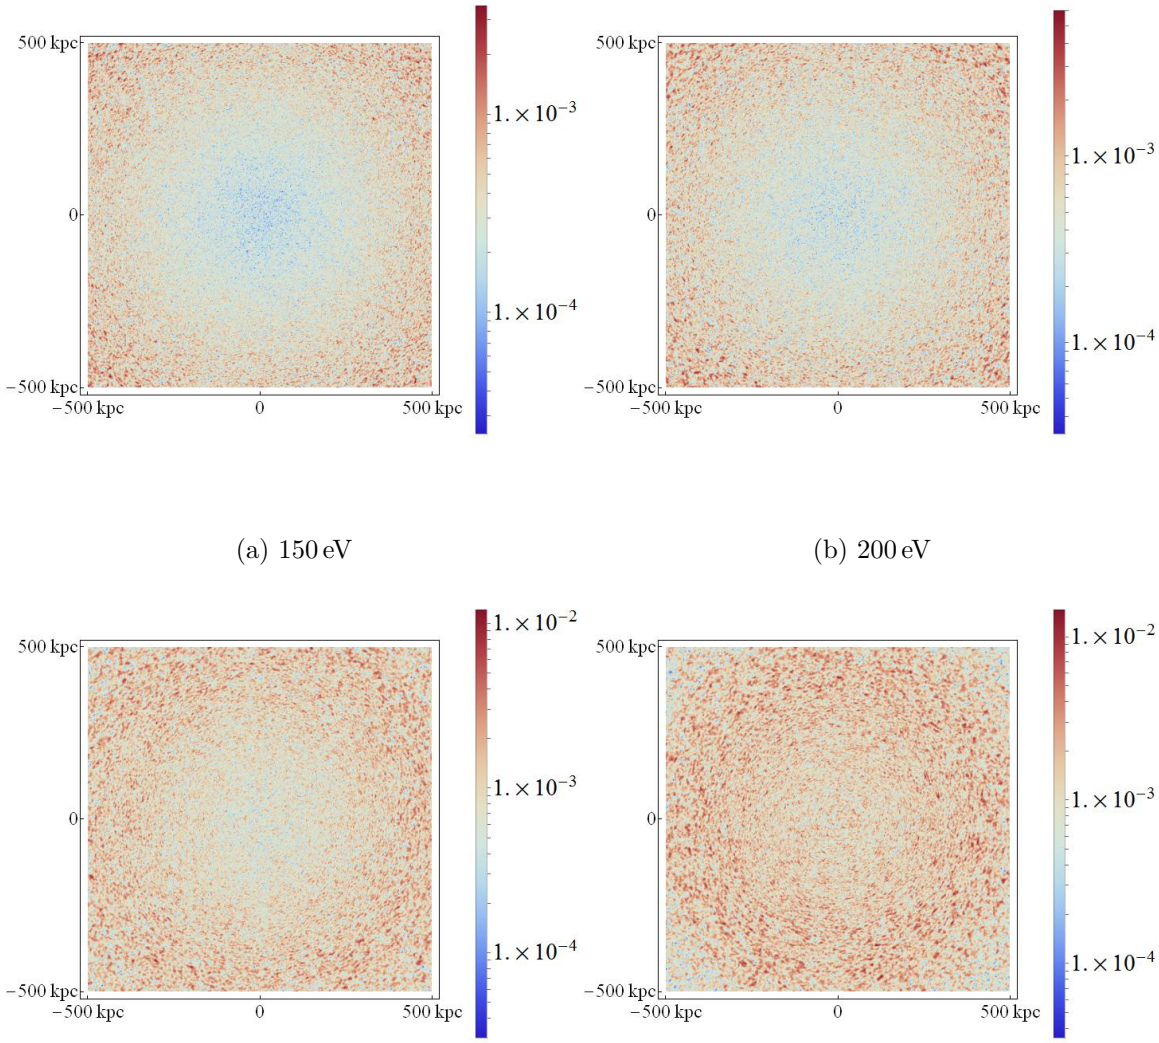

(a) 150 eV

(b) 200 eV

(c) 300 eV

(d) 400 eV

Figure 5.4: Conversion probabilities for energies 150 eV to 400 eV for Model 1, with $\eta = 0.5$, $B_0 = 4.7\,\mu$G and $M = 7 \times 10^{12}$ GeV.

of the scaling of the total magnetic field with electron density, η, and the central value of the magnetic field, B_0, are then $\eta = 0.5$ and $B_0 = 4.7\,\mu$G, respectively. In section 5.7.2.2 we will also consider the effects of 1σ variations of the parameters of the magnetic field model on the resulting conversion probabilities.

The simulated conversion probability maps (best viewed in colour) for this model are shown in figures 5.3, 5.4 and 5.5 with a pixel size of $(2\,\text{kpc})^2$. Figure 5.6 shows the conversion probabilities as a function of the impact parameter, for 8 energies from 25 keV to 2 keV.

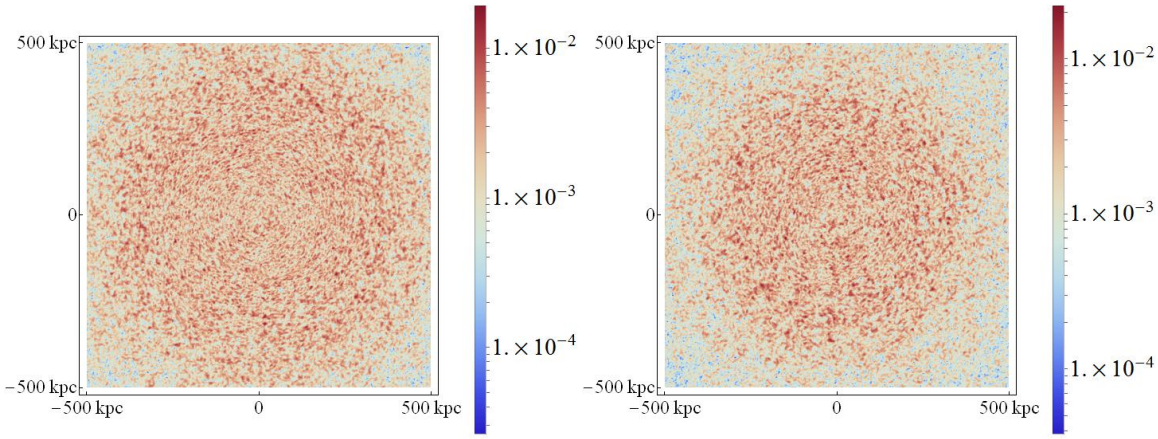

(a) 600 eV

(b) 800 eV

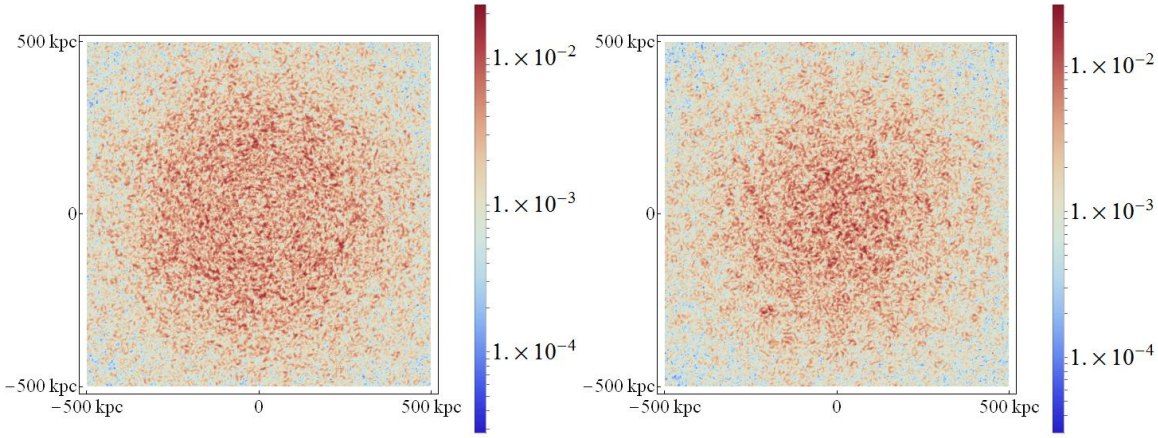

(c) 1000 eV

(d) 2000 eV

Figure 5.5: Conversion probabilities for energies 600 eV to 2 keV for Model 1, with $\eta = 0.5$, $B_0 = 4.7\,\mu$G and $M = 7 \times 10^{12}$ GeV.

These figures illustrate two key features of the results that are ubiquitous in all the magnetic field models we have studied. First, the overall conversion probabilities increase with energy, up to a maximal energy at which they saturate.

Second, the morphology of the conversion probabilities is quite distinct at low and high energies. At low energies, the conversion probabilities are lowest for ALPs which pass through the very centre of the cluster. As a function of increasing impact parameter, the conversion probabilities increase, reach a maximum at some intermediate radius, before again decreasing towards the edge of the cluster. This behaviour

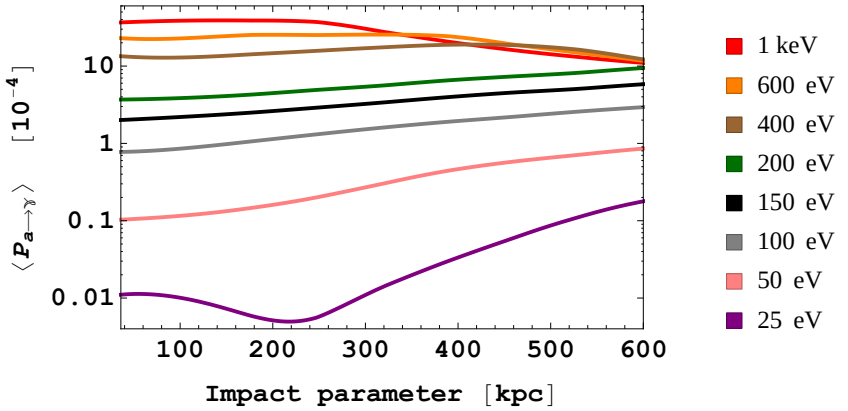

Figure 5.6: Conversion probabilities as a function of impact parameter for Model 1 with $\eta = 0.5$ and $B_0 = 4.7\ \mu$G.

is clearly visible in the 400 eV plot (figure 5.4d) — for lower energies than this the point of maximal conversion probability lies beyond the range of the simulation. At ultra-low energies (c.f. the 25 eV plot, figure 5.3a), a curious, unanticipated ring-like structure is visible. Here the conversion probabilities increase, decrease and then increase again. We discuss the origin of this at greater length below.

In contrast, higher energy ALPs have a maximal conversion probability at the centre of the cluster, with a monotonic decrease in conversion probability on going to larger radii. A crossover between the high-energy regime of 'central dominance' and the low-energy regime of 'central deficit' can be observed for modes with energies of $400\,\text{eV} < \omega < 1\,\text{keV}$ for Model 1.

In fact, both these generic features of the conversion probabilities can be understood from the single-domain solution of equation 5.17, even though it is not fully applicable to the multi-scale fields considered here. The single domain conversion probability is (to leading order in θ) given by

$$P(a \to \gamma) \propto B_0^2 \frac{\tilde{\omega}^2}{M^2} \tilde{n}(r)^{2(\eta-1)} \sin^2\left(0.93 \frac{\tilde{L}\tilde{n}(r)}{\tilde{\omega}}\right), \tag{5.26}$$

where $\tilde{n}(r) = \frac{n_e(r)}{n_0}$, $\tilde{L} = \frac{L}{(1\,\text{kpc})}$ and $\tilde{\omega} = \frac{\omega}{(200\,\text{eV})}$. Here we have factored out the dependence of the total magnetic field on the electron density. The fractional electron density $\tilde{n}(r)$ is completely determined by the β-model, and decreases from unity at the cluster centre to ≈ 0.15 at $r = 600\,\text{kpc}$.

We now note that for either sufficiently large ω or sufficiently small $\tilde{n}(r)$, the argument of the sin function becomes small. In the small-Δ approximation the conversion probabilities are given by

$$P(a \to \gamma)_{\text{single domain}} \propto \frac{B_0^2 L^2}{M^2} \tilde{n}_e^{\eta}(r)\,, \tag{5.27}$$

$$P(a \to \gamma)_{\text{per unit length}} \propto \frac{B_0^2 L}{M^2} \tilde{n}_e^{\eta}(r)\,. \tag{5.28}$$

Thus, according to the single-domain formula, sufficiently far away from the cluster centre the small-θ *and* small-Δ approximations should be valid for all energies above a certain cut-off. Evidently, at large radii the small-angle approximation appears as an 'attractor' with a radial dependence completely determined by the modulation of the magnetic field with $\tilde{n}(r)^{\eta}$.

As the impact parameter is decreased and $\tilde{n}(r)$ increases, some modes will cease to be well described by the small-Δ approximation and will rather require the full equation (5.26). Such modes will leave the small-angle 'attractor' solution. According to the single-domain formula, modes with sufficiently low energies may undergo several 2π rotations of Δ as the impact parameter is decreased towards the centre of the cluster, and these modes will in particular exhibit rings of decreased probabilities as Δ comes close to an integer multiple of π.

However, for an ALP traversing multiple magnetic field domains with slightly varying electron densities and coherence lengths, some of the detailed features of the single-domain probabilities can be expected to be 'washed out'. In particular, in the large Δ regime it is reasonable to approximate $\langle \sin^2 \Delta \rangle = \frac{1}{2}$ for ALPs traversing multiple domains of slightly varying size. Then for $B \propto B_0 n_e^{\eta}$, we have

$$P(a \to \gamma)_{\text{single domain}} \propto \frac{B_0^2 \tilde{\omega}^2}{M^2} \tilde{n}_e^{2(\eta-1)}\,, \tag{5.29}$$

$$P(a \to \gamma)_{\text{per unit length}} \propto \frac{B_0^2 \tilde{\omega}^2}{LM^2} \tilde{n}_e^{2(\eta-1)}\,. \tag{5.30}$$

As in all cases we consider $\eta < 1$, this gives increasing conversion probabilities as the electron density decreases with radii. This increase will continue until the small-Δ regime is reached, when the conversion probability is again described by equation

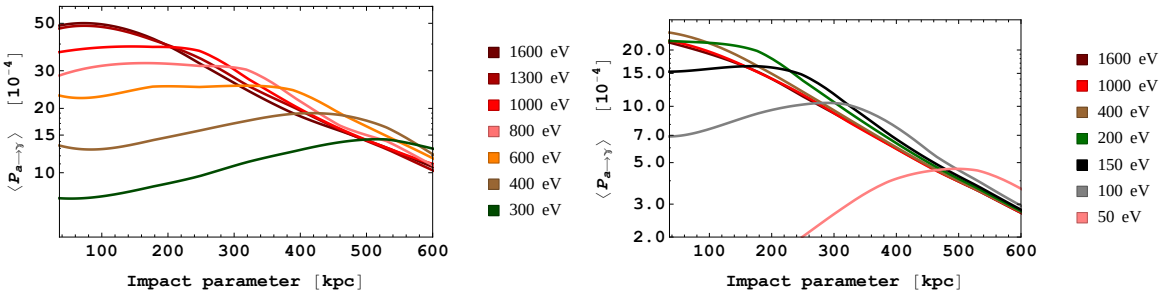

(a) Model 1

(b) Model 2

Figure 5.7: At large impact parameter the conversion probabilities tend to the small angle approximation, as here illustrated for Model 1 and Model 2.

(5.27).

We now note that several of the features predicted from the single-domain formula also appear in the radial probabilities emerging from the full numerical simulation with multi-scale magnetic fields, c.f. figures 5.6 and 5.7a. From these figures, we note that modes with $\omega \gtrsim 400\,\mathrm{eV}$ have equal conversion probabilities at large radii (small $\tilde{n}(r)$), consistent with these modes entering the small-Δ approximation. At smaller impact parameter (corresponding to larger maximal $\tilde{n}(r)$), the lower energy modes begin to decouple from the small-Δ approximation with ultimately decreasing probabilities as a result, perfectly consistent with the single-domain result. Only the highest energy modes stay in the attractor curve for any impact parameter, as illustrated in figure 5.7a by the modes with $\omega = 1300\,\mathrm{eV}$ and $\omega = 1600\,\mathrm{eV}$.

The single domain analysis also predicts the presence of regions with highly suppressed conversion probabilities as Δ approaches an integer multiple of π, and we note that this feature most likely explains the ring structure of the conversion probability in the 25 eV plot.

We can use the qualitative consistency of the conversion probabilities with the single-domain result to generate a heuristic estimate of an 'effective coherence length' of the magnetic field. For the high-energy modes in figure 5.7a we have argued that $\Delta \ll 1$ at large impact parameter, and we may heuristically associate the radius of maximum conversion probability, $r_{\max}(\omega)$, for each mode with the phase $\Delta = \pi/2$ in

the single-domain formula. By furthermore noting that the largest contribution to the conversion probability for modes close to the small-Δ approximation will come from the region closest to the cluster centre, we have

$$\frac{\pi}{2} = 0.93 \frac{\tilde{L}\tilde{n}(r_c(\omega))}{\tilde{\omega}} \,. \tag{5.31}$$

From this formula, we can extract an 'effective coherence length' $\tilde{L}$ for each mode that has a peaking conversion probability within the range studied. For Model 1 this range corresponds to modes with energies $300\,\mathrm{eV} \leq \omega \leq 1\,\mathrm{keV}$, as indicated in figure 5.7a. The decoupling of all modes is consistent with the single-domain estimate for effective coherence lengths in the $13 - 15\,\mathrm{kpc}$ range. In comparison, the full numerical simulation involves magnetic fields coherent over all scales from $2\,\mathrm{kpc}$ to $34\,\mathrm{kpc}$, with a mean coherence length of $\approx 10\,\mathrm{kpc}$.

The single-domain intuition also holds for other magnetic field models. From figure 5.7b, we note that Model 2 may be associated with an 'effective coherence length' of the magnetic field in the $2.0 - 2.2\,\mathrm{kpc}$ range, based on the peak positions of modes with $50\,\mathrm{eV} < \omega < 200\,\mathrm{eV}$. The full multi-scale model has coherence lengths in the $2 - 5\,\mathrm{kpc}$ range, with a mean value of $3.2\,\mathrm{kpc}$. These estimates indicate that the physical picture of ALP–photon conversion motivated by the single-domain analysis is also quite accurate for more complicated magnetic field configurations such as the multi-scale configurations considered in this chapter.

Model 3 is distinguished by having the largest range of scales in the magnetic field, from $2 - 100\,\mathrm{kpc}$. As we see in figure 5.8, for this model at large radius ($500 - 600\,\mathrm{kpc}$) the conversion probabilities have not converged to a small-angle approximation. We can again understand this behaviour using the single-domain formula.

The greater range of coherence lengths imply that even ALPs traversing the cluster at large impact parameter are likely to encounter domains in which the small-Δ approximation is not valid. Increasing ω always has the effect of decreasing Δ and thereby bringing a larger fraction of the traversed distance into the small-angle approximation, which results in an increased overall conversion probability according

to equations (5.28) and (5.30). This explains the increase in conversion probabilities with ω in Model 3, even at the largest radii.

A further difference in figure 5.8 compared to figures 5.7a or 5.7b can also be understood. For model 3, at large radii the conversion probabilities are roughly similar for the higher energy modes, with consistent small increases as the energy of the ALP increases. This behaviour is absent for models 1 and 2. We can understand this through the different power spectra of the models. For the Kolmogorov spectra of models 1 and 2, power dominantly lies in the largest coherence lengths. The conversion probabilities are then highly suppressed for ALP modes that have insufficient energy to reach the small-Δ regime, given this range of coherence lengths. For the truncated spectrum of model 2 and the flat spectrum of model 3, more power of the magnetic field is allocated to shorter scales, and even lower energy modes can reach the small-Δ approximation. Owing to the wide range of coherence lengths in the Model 3 magnetic field, ALPs traversing the cluster will still pass through regions which are not well-described by a small-Δ approximation — even for relatively high energy modes. As the ALP energy is increased, only a small additional fraction of the large coherence lengths are brought into the small-Δ approximation, which explains the gradual approach to the small-Δ approximation in this case.

Let us finally for the sake of clarity remark on features of the conversion probabilities that are not well-captured by the single-domain formula. As Δ approaches $\pi/2$, the single-domain formula predicts a relative decrease in the conversion probability with respect to the small-Δ approximation, yet the full numerical simulations exhibit conversion probabilities with a clear trend of 'overshooting' the small-angle attractor at this point.

A final general comment about the simulations: it is useful to know how much variation one can expect purely from repeating a simulation with identical magnetic field parameters. By repeating simulations with the Model 1 parameters, we found that the averaged conversion probabilities within each annulus varied by at most 9%, where in most cases the difference was less than 5%. We note that the most significant

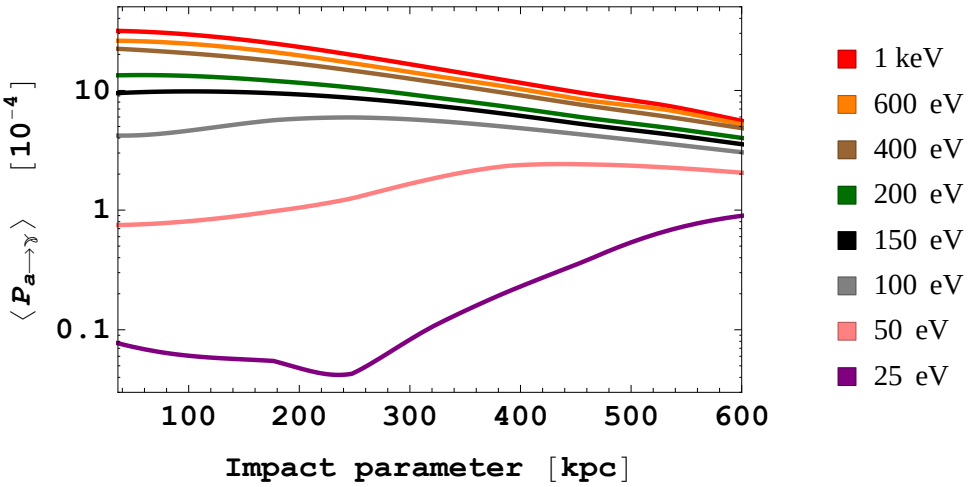

Figure 5.8: Mean conversion probability as a function of radius from the centre of Coma, for Model 3 of table 5.1 and with $M = 5.7 \times 10^{12}\,\mathrm{GeV}$.

variations occurred for larger energies. We thus conclude that our magnetic field model does not generate large fluctuations in conversion probabilities. However, to account for this when we plot comparisons between simulated and observed luminosities, we include a statistical error of 10% on our values.

5.7.2.2 Comparison with Observed Luminosities

We may now compare the predictions of ALP–photon conversion of a CAB in the Coma cluster to the actual observations of the soft excess with ROSAT, based on the analysis of Bonamente et al. [121]. Specifically, we will focus on the overall unabsorbed excess luminosity in the $0.2-0.4\,\mathrm{keV}$ band for various different annular regions around the centre of Coma.

Since the spectral shape of the soft excess is poorly known, the analysis of [121] quotes results for two different spectral fits to the excess emission: the first employs a power-law spectrum with photon index 1.75 (so that the excess flux is $\mathrm{d}N_\gamma/\mathrm{d}E \sim \nu^{-1.75}$), and the second is based on a thermal spectrum with $T = 80\,\mathrm{eV}$. These results — which differ from each other by an overall factor of ≈ 2.4 — are shown in table 5.3. While neither of these spectral models are in exact agreement with the shape of the photon spectrum obtained from CAB conversion, the thermal bremsstrahlung

Region (arcminute)	$L_{NT}(10^{41}\text{erg s}^{-1})$	$L_{thermal}(10^{41}\text{erg s}^{-1})$
0 - 3	11	4.6
3 - 6	22	9.1
6 - 9	26	10
9 - 12	25	10
12 - 18	47	21

Table 5.3: The results of [121] for excess luminosity from the Coma cluster.

spectrum has an exponentially decreasing tail in the $0.2 - 0.4\,\text{keV}$ range, as does the CAB spectrum for mean ALP energies $\langle E_{\text{CAB}} \rangle \lesssim 200\,\text{eV}$. We will therefore use the thermal excess luminosities of table 5.3 when comparing the predictions of our model to the data. We also note that the model uncertainty in the extraction of the soft excess mostly affects the overall luminosity and to a much smaller degree its spatial distribution. In our model, the overall luminosity has a simple dependence on the values of M and ΔN_{eff}, and the uncertainty in the overall luminosity translates into an uncertainty in $\Delta N_{\text{eff}}/M^2$.

We will now present our main results for the CAB explanation of the soft excess in Coma. We will first discuss luminosities from ALP–photon conversion in the Model 1 magnetic field (including 1σ variations of the model parameters η and B_0) and we will then turn to Model 2 and Model 3.

Model 1: Figure 5.9 shows the comparison between the observed (thermal) luminosity and that produced by the baseline model of [187] and two other models related by 1σ variations of the model parameters. The integrated luminosity in the $0.2 - 0.4\,\text{keV}$ range has been normalised to the total luminosity of the soft excess in the same range by varying M independently for each model. In all cases $\Delta N_{\text{eff}} = 0.5$ has been assumed, but we emphasise that alternate values for these parameters that normalise luminosities can be obtained by scaling $\Delta N_{\text{eff}} \to \lambda \Delta N_{\text{eff}}$ and $M \to \sqrt{\lambda} M$.

From figure 5.9, we note that while the luminosity in each bin is within observations by a factor of a few, there is a clear tendency to underproduce photons in the centre and overproduce them in the outskirts. ALP–photon conversion in the Model 1

magnetic field therefore does not provide a particularly good description of the Coma soft excess.

The three variations of Model 1 in figure 5.9 correspond to the magnetic field model parameters which best fit the Faraday rotation measures, here denoted by $\eta = 0.5$, and 1σ variations to $\eta = 0.4$ and $\eta = 0.7$. In all cases, ALP–photon conversion under- and over-produces photons in the inner and outer regions respectively. Understandably, increasing η so that the magnetic field falls off more rapidly with radius while simultaneously increasing B_0 to match Faraday rotation measures results in more luminosity to smaller radii relative to larger radii. We note that these variations are not large enough to make the CAB prediction of the morphology of the soft excess compatible with observations.

However, despite the poor fit of the ALP-converted photon luminosities to the soft excess for Model 1, it would be premature to conclude that the conversion of the CAB cannot explain the soft excess. Even before considering systematic differences between the magnetic field models of [187] and the actual magnetic field in Coma, as discussed in section 5.4.2, Faraday rotation measures only constrain the magnetic field model up to degeneracies in the spectral index of the vector potential, ζ, and the Fourier mode cut-off scale $\Lambda_{\rm max}$. In Model 3, we consider a magnetic field model which provides an equally good fit to Faraday rotation measures as Model 1, but with $\zeta = 0$ and $\Lambda_{\rm max} = 100\,{\rm kpc}$ (as opposed to $\zeta = 5/3$ and $\Lambda_{\rm max} = 34\,{\rm kpc}$ for Model 1). The stochastically generated magnetic field prior to modulation by $n_e(r)$ has $\tilde{B}(k)_{\rm (gen.)} \sim k^{(-\zeta-2)/2}$, so that for $\zeta = 0$ the power integral $\int {\rm d}k k^2 \tilde{B}(k)^2_{\rm (gen.)}$ has constant support from $k_{\rm min}$ to $k_{\rm max}$. Equivalently, magnetic fields with $\zeta > 0$ locate more power to smaller k-numbers and larger physical scales.

Thus, a key difference between Model 1 and Model 3 is the distribution of effective coherence lengths. In order to highlight the effect of concentrating more power of the magnetic field on smaller scales, we first consider a toy magnetic field model, which does not provide a good fit to Faraday rotation measures.[3] This is our Model 2, to

[3]Note that the small scales in this model are however not necessarily unphysical. Faraday rotation constrains the magnitude and coherence lengths of the parallel component of the magnetic field along

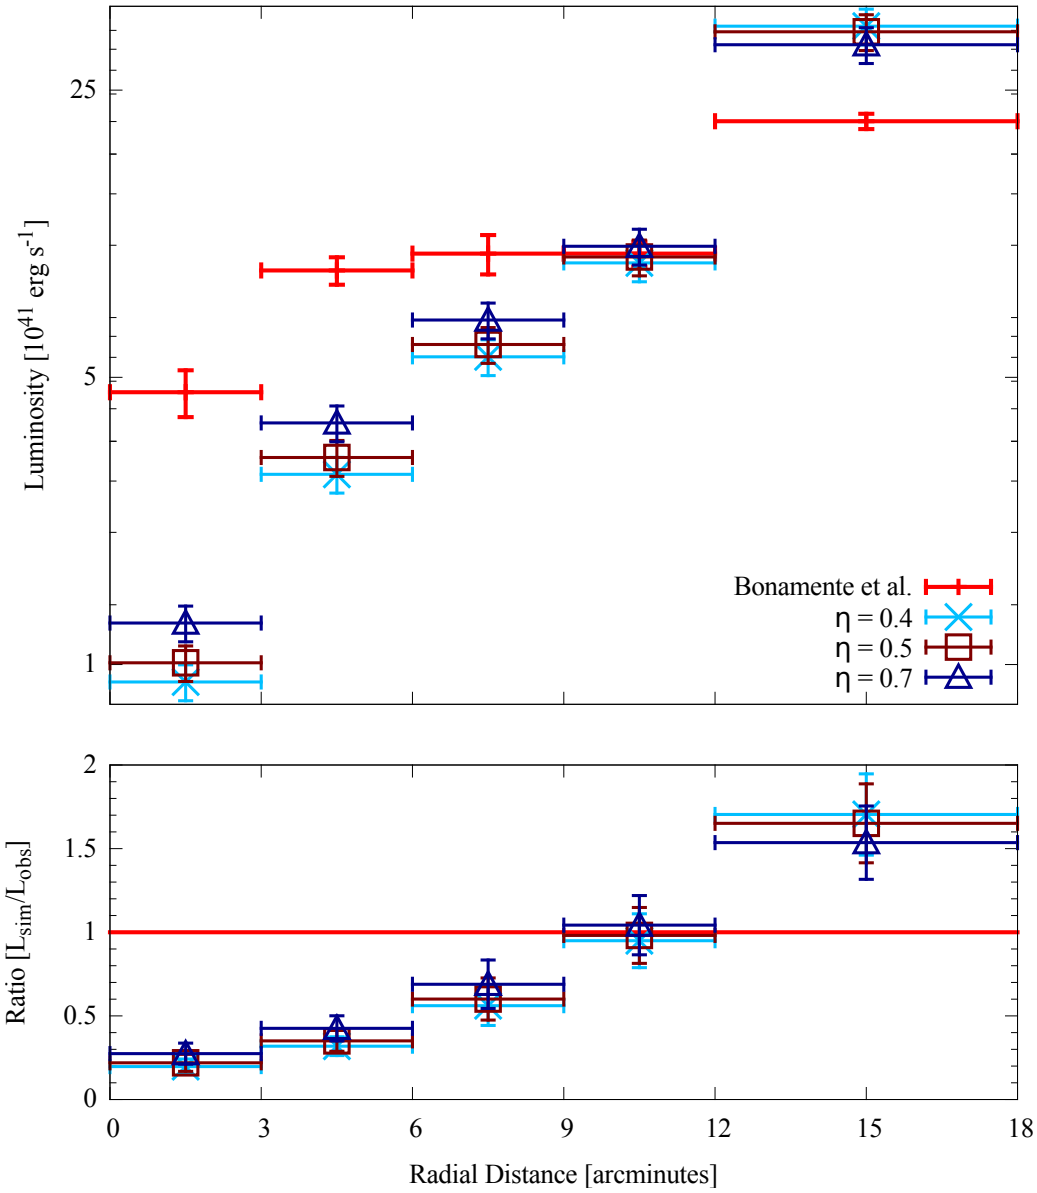

Figure 5.9: Luminosity comparison for Model 1 with different η values, compared to the 'thermal' excess data. For $\Delta N_{\rm eff} = 0.5$ and $\langle E_{\rm CAB} \rangle = 150\,{\rm eV}$, normalisation of the integrated luminosities give $M = 6.1 \times 10^{12}$, 6.7×10^{12} and $6.5 \times 10^{12}\,{\rm GeV}$ for $\eta = 0.4$, 0.5 and 0.7, respectively.

which we now turn.

Model 2: In this model the generated magnetic field, prior to modulation by a function of the electron density, only varies on scales between $2 - 5\,{\rm kpc}$. In this range, the magnetic field varies with $\zeta = 5/3$ and the modulation with electron density is

the line of sight, whereas ALP conversion involves the transverse components. The magnetic field models used here make these equal by assumption, but if the latter is actually smaller than the former by a factor of a few, this model could still be consistent with Faraday rotation.

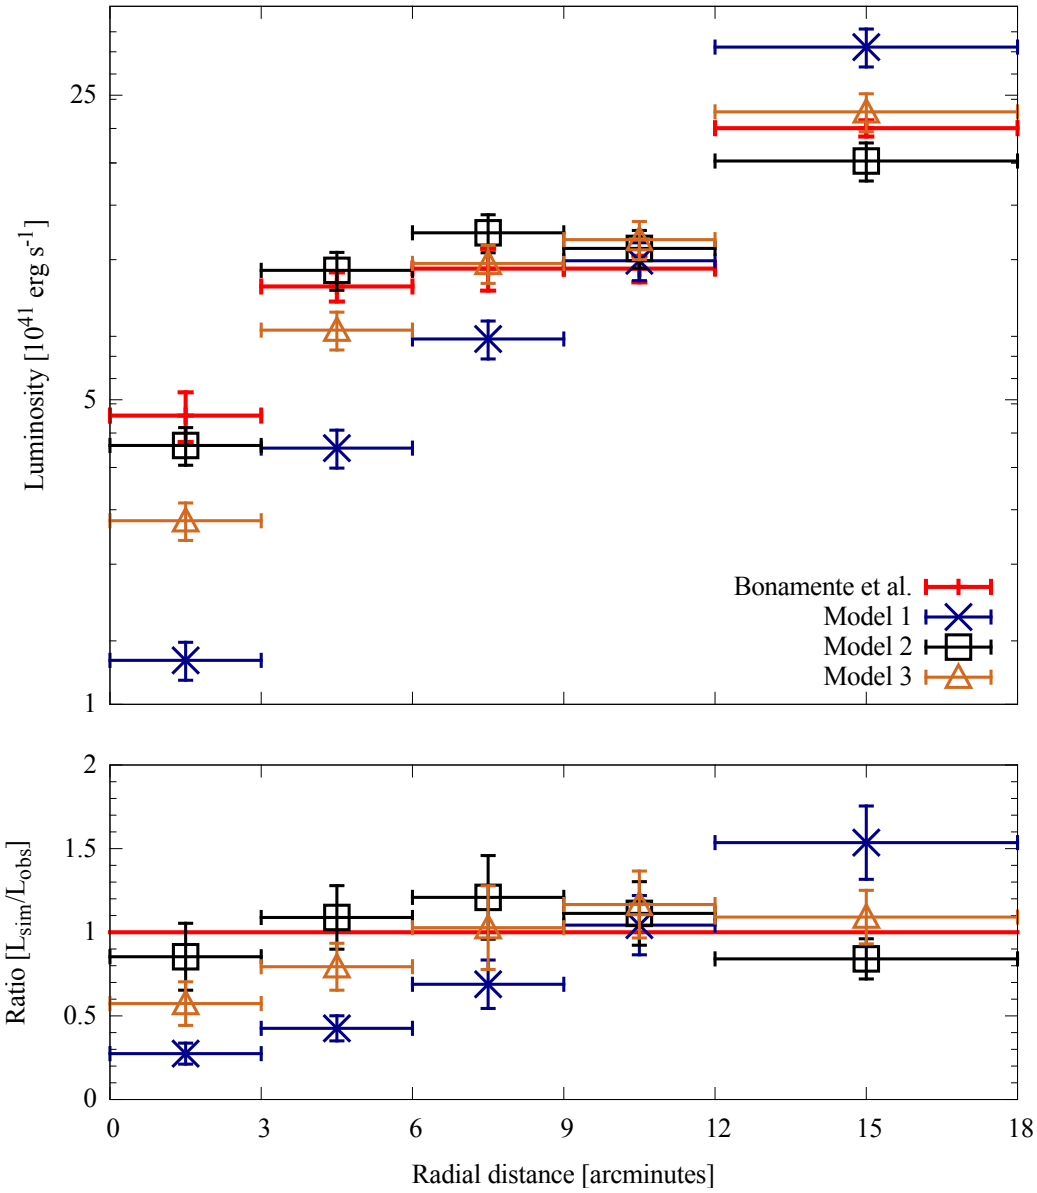

Figure 5.10: Luminosity comparison for the different models compared to the 'thermal' excess data. For $\Delta N_{\rm eff} = 0.5$ and $\langle E_{\rm CAB} \rangle = 150\,{\rm eV}$, normalisation of the integrated luminosities gives $M = 6.5 \times 10^{12}$, 5.2×10^{12} and $5.7 \times 10^{12}\,{\rm GeV}$ for Models 1 ($\eta = 0.7$), 2, and 3 respectively.

obtained with $\eta = 0.7$, setting $B_0 = 5.4\ \mu{\rm G}$. The simulated photon luminosities from this model match the observational data for the soft excess very well, as is shown in figure 5.10.

Based on our discussion in section 5.7.2.1 on the radial dependence of the simulated conversion probabilities, we may interpret the improved fit as due to a decreased 'effective coherence length', resulting in modes approaching the small-Δ attractor at smaller radii, c.f. figures 5.7a and 5.7b.

Model 3: We now return to magnetic field models consistent with observations of Faraday rotation measures in Coma, but focus on models which concentrate more power on smaller scales relative to Model 1. The effective degeneracy between values of ζ and $\Lambda_{\rm max}$ allows for reducing ζ by simultaneously increasing $\Lambda_{\rm max}$, as is illustrated in figure 16 of reference [187]. Here again, we take $\eta = 0.7$ and $B_0 = 5.4$ μG. The resulting photon luminosities from CAB conversion are shown in figure 5.10 and again show a good agreement with the observed soft excess.

The conclusions to draw from these are that an explanation of the soft excess via ALP–photon conversion appears to require the transverse components of the magnetic field to have more power on shorter scales than in the Kolmogorov spectra of [187]. This can be achieved either by allowing a flatter spectrum, so as to be consistent with Faraday rotation measures even for a Gaussian magnetic field, or by having different coherence lengths for transverse and parallel components of the magnetic field.

5.7.2.3 Constraints on the ALP–photon Coupling

Having established that the CAB explanation of the cluster soft excess is in reasonable agreement with observations for magnetic field models motivated by observations of Faraday rotation measures, we will now discuss how additional observational constraints give rise to a relatively limited range of possible values for M and the mean CAB energy, $\langle E_{\rm CAB} \rangle$.

Strong bounds on the ALP–photon coupling have been obtained by laboratory experiments and astrophysical arguments, see Chapter 3 for more information. For light ALPs, the CAST search for solar ALPs has set a bound $M > 10^{10}\,\text{GeV}$. Proposed experiments looking either for light shining through a wall, such as ALPS-II [69], or for solar ALPs, such as IAXO [72], are expected to improve this bound by a factor of $10 - 20$.

The most stringent constraint for light ALPs, $m_a \lesssim 10^{-9}\,\text{eV}$, is from the absence of a γ ray burst in coincidence with neutrinos from Supernova 1987A. This bounds the coupling to be $M > 10^{11}\,\text{GeV}$ [45–47]. In addition to limits, an ALP with parameters

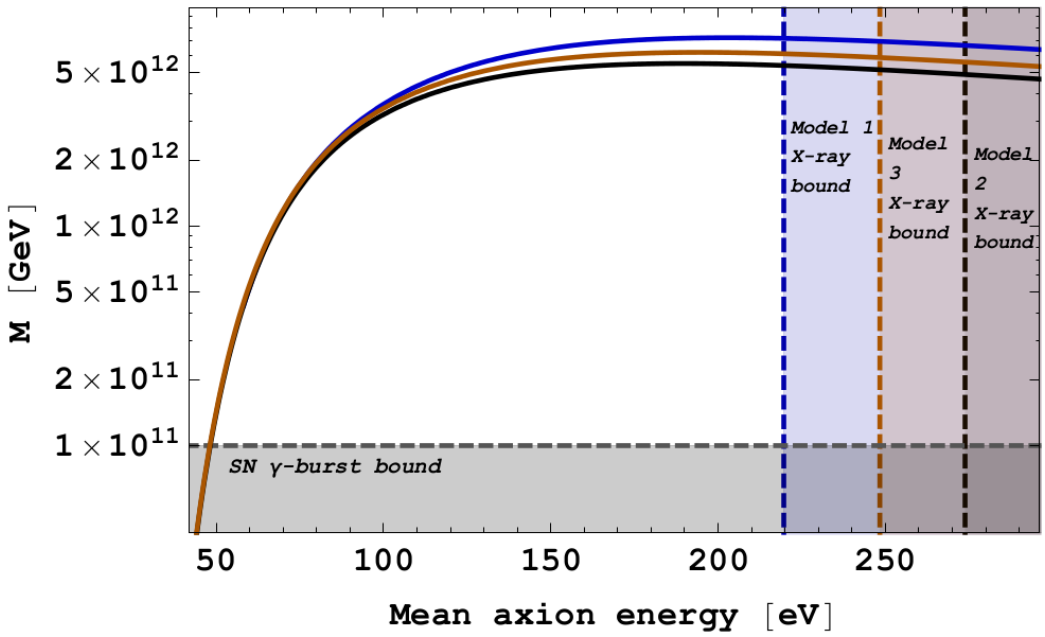

Figure 5.11: The values of M required to normalise the total soft excess from ALP–photon conversion in the $0.2-0.4\,\mathrm{keV}$ band to the observed total excess luminosity with the central 18 arcminutes of Coma as a function of $\langle E_{\mathrm{CAB}}\rangle$ for $\Delta N_{\mathrm{eff}} = 0.5$. Model 1 is represented by the blue solid curve, Model 2 by the black curve and Model 3 by the orange curve. The supernova γ-burst bound is indicated by a dashed grey line, and the bounds from overproduction of X-rays in the $0.5-0.6\,\mathrm{keV}$ range are indicated by a vertical dashed line for each model.

$m_a \lesssim 10^{-9}\,\mathrm{eV}$ and $M \approx 10^{11} - 10^{12}\,\mathrm{GeV}$ has been suggested to explain the anomalous transparency of the universe to gamma rays [60, 62–65].

We find that the CAB explanation for the soft excess in Coma is possible for light ALPs with $M \sim 10^{11} - 7\times 10^{12}\,\mathrm{GeV}$, which is in the same range as suggested by the gamma ray transparency hint.

Furthermore, while in our model ALP–photon conversion in the Coma magnetic field should explain the soft X-ray excess, strong restrictions on the model parameters can be obtained by noting that higher energy photons should not be abundantly produced from ALP–photon conversion: the excess is soft and does not survive to higher energies. This poses a restriction on the support of the CAB spectrum, as parametrised by the mean CAB energy. Here, we will impose that ALP–photon conversion in the $0.5-0.6\,\mathrm{keV}$ band should not contribute to more than 10% of the thermal luminosity in this range. More accurate — and quite possibly more stringent — bounds may be obtained by detailed extraction of the soft excess based on dedicated

templates for the CAB spectrum.

The astrophysical and laboratory bounds on M may then be translated into a lower bound on the mean CAB energy, and we find that, quite model-independently, this gives $\langle E_{\rm CAB}\rangle \gtrsim 45 - 50\,{\rm eV}$. Meanwhile, the bound from X-ray overproduction gives more model-dependent constraints, but allows for $\langle E_{\rm CAB}\rangle \lesssim 250\,{\rm eV}$ for the interesting Model 3.

In sum, we may then express the interesting values for the scale M as

$$10^{11}\,{\rm GeV} \lesssim M \lesssim 7\times 10^{12}\sqrt{\frac{\Delta N_{\rm eff}}{0.5}}\,{\rm GeV}\,. \tag{5.32}$$

5.7.3 Summary of Coma Results

Let us conclude this section by summarising our results from the study of Coma. This study has led us to three main conclusions. First, we have confirmed the assertion of [116] that the overall luminosity of the soft excess can easily be explained by ALP-photon conversion. For example, a CAB with mean energy of $\langle E_{\rm CAB}\rangle \approx 150\,{\rm eV}$ may explain the soft excess for an ALP-photon coupling of $M \sim \sqrt{\frac{\Delta N_{\rm eff}}{0.5}}\,6\times 10^{12}\,{\rm GeV}$.

Second, we have found that the success of the CAB explanation for the morphology of the cluster soft excess depends on some of the details of the cluster magnetic field, and in particular on the distribution of the coherence lengths of the magnetic domains. The morphology obtained from ALP-photon conversion is compatible with the observed soft excess for magnetic field models which predominantly have short (transverse) coherence lengths of a few kpc, as well as models which have uniformly distributed (transverse) coherence lengths from a few to 100 kpc. For Gaussian magnetic field models with transverse coherence lengths predominantly in the 10 kpc range, such as our Model 1 above, the CAB explanation does not provide a close match to the observed soft excess morphology in Coma. The CAB explanation of the soft excess then requires either a flatter power spectrum or a shorter coherence length for the transverse component of the magnetic field (note that Faraday rotation measures constrain only the magnetic field component parallel to the line of sight).

Third, the requirement that the cluster soft excess originates from a CAB strongly constrains the CAB properties. The absence of an X-ray excess at $E \gtrsim 0.5\,\mathrm{keV}$, and the astrophysical requirement that the ALP-photon coupling satisfies $M \gtrsim 10^{11}\,\mathrm{GeV}$, constrains the mean CAB energy to the range

$$50\,\mathrm{eV} \lesssim \langle E_{\mathrm{CAB}} \rangle \lesssim 250\,\mathrm{eV}\,, \tag{5.33}$$

so that the CAB explanation of the soft excess is viable for $10^{11}\,\mathrm{GeV} \lesssim M \lesssim 7 \times 10^{12}\sqrt{\frac{\Delta N_{\mathrm{eff}}}{0.5}}\,\mathrm{GeV}$. The ALP mass is similarly constrained to $m_a \lesssim 10^{-12}\,\mathrm{eV}$.

We now continue to study the soft X-ray excess in galaxy clusters, by studying the CAB in three more galaxy clusters, and comparing to soft X-ray observations in these clusters.

Following the publication of the work this chapter is based on, and prior to the following study of A665, A2199, and A2255, the authors of [212] continued to study the soft X-ray excess in the Coma cluster. They extrapolated this magnetic field model out to radii ~ 5 Mpc from the cluster centre, and found the observed excess at these large radii [157], could be fit by the inverse coupling 5×10^{12} GeV $\leq M \leq 3\times10^{13}$ GeV.

5.8 A665, A2199 and A2255

In this section we study the galaxy clusters A665, A2199, and A2255. There is no observed excess in the A665 cluster, thus it is an important testing ground for the CAB hypothesis. A CAB should not produce a large X-ray signal in this cluster. The other two clusters have hints of excesses, but at lower significances than in the Coma cluster.

5.8.1 Luminosity and Fractional Excess

We start by outlining the method we use to compare to data. The data is collected in the $0.1 - 0.28$ keV R1R2 band of ROSAT. In the case where an excess was observed, the luminosity in the $0.2 - 0.4$ keV band was computed by [121] for thermal and non-thermal emission models. Thus in the previous section on Coma, we could compare

directly to the observed luminosity. Here however, we are dealing with clusters where no excess was observed, thus we must compare instead to the number of counts observed. The analysis of [121] lists the expected and observed count rates in the R1 and R2 bands (R1R2, $0.1 - 0.28$ keV) of the ROSAT PSPC detector. To compare to these observations we use the NASA PIMMS software [215], we use this to predict the count rate in the ROSAT PSPC R1R2 band, given a flux and spectral shape.

The result of the simulation is a grid of conversion probabilities, which we compute for several ALP energies between 25 eV $-$ 1 keV. This grid is divided into ten concentric annuli and the conversion probability in each annulus is averaged. In this way we end up with an ALP$-$photon conversion probability distribution which is a function of radius and energy, $P(r, E)$. We use this ALP$-$photon conversion probability distribution to compute the soft X-ray luminosity from CAB conversion in the cluster. We do this by integrating the product of the probability function $P(r, E)$ and the CAB energy spectrum $E\,\frac{dN}{dE}$ over the area of the annulus of interest and over the energy interval $0.1 - 0.28$ keV. The CAB energy density is normalised relative to the CMB by setting $\Delta N_{\rm eff} = 0.5$.

As can be seen in [121] the assumed spectral shape of the emission can have a large impact on the luminosity for a given a count rate. This is both due to the different energy channels of the input and output, and because the detector response varies over the R1R2 band. To use the PIMMS tool we need to state what spectral shape the signal has—which would be a modified CAB spectrum to take into account the energy dependence of the conversion probabilities. Since we are converting between luminosity and count rate in the same energy channel, we checked to see the importance of the specific spectral model used. We use two models: a thermal blackbody spectrum which peaks at 150 eV, and a single temperature (0.08 keV) thermal bremsstrahlung MEKAL plasma model. This second model was the model used for calculating the luminosities in [121] which we compared the simulation results to in the previous section. We found little difference between the count rate prediction for the two models. The count rate is also slightly sensitive to the temperatures of

both the blackbody and the plasma model (which in principle should be changed when scanning over the mean CAB energy parameter later), the largest deviations only came when the temperatures were pushed much higher or lower than the values above, and thus we keep these temperatures fixed.

The largest uncertainty comes from the count rate prediction itself. There is likely an inherent systematic uncertainty in simulating the count rate in the ROSAT detector for a given luminosity. For instance, we took the excess luminosities given in [121] and converted them to count rates in the ROSAT R1R2 band using PIMMS to compare with observed count rates. We found that the count rates predicted by PIMMS were factors $0.5 - 2$ different than the observed counts, though these errors are probably larger due to the different energy ranges of the luminosity and count rates mentioned earlier. Thus it is important that we do not claim to be too precise with the count rate prediction from the simulation, as such we allow the count rates to vary by $\pm 50\%$ when constraining the CAB parameter space later.

5.8.2 Results

In this section we present the results of the simulations for the three individual clusters. We study the morphology and use the observed (lack of a) soft excess to constrain the CAB parameters M, the inverse coupling, and $\langle E_{CAB} \rangle$, the CAB spectrum mean energy. We then compare the best fit regions for the three clusters to that of the centre of Coma. Throughout we take $H_0 = 71$ km s^{-1}Mpc^{-1} and $\Omega_m = 0.27$, to be consistent with [206, 207], though we do not expect this choice to have a significant effect on the results compared to that from the systematic uncertainties on the magnetic fields.

5.8.2.1 A665

No excess is observed from A665, and it is thus an important check for the CAB hypothesis. Since the CAB will always convert to photons in magnetic fields, the conversion probability of an ALP to a photon needs to be small in A665 such that an

observable excess is not produced.

A665 will have smaller conversion probabilities than for instance the Coma cluster, where the CAB simulation agreed with the significant soft excess observed, for two reasons. Firstly, the central magnetic field is smaller (1.3 μG compared to 4.7 μG in Coma), whilst the electron density is roughly similar. Since the conversion probabilities always scale as $P(a \to \gamma) \propto B^2$ (see equations 5.20 and 5.21), there is a large $(B_{\rm Coma}/B_{\rm A665})^2$ reduction in conversion probabilities for A665. Secondly, A665 has larger coherence lengths ($\Lambda_{max,\ \rm Coma} \ll \Lambda_{max,\ \rm A665}$) roughly by a factor 10. With this increase, we see according to equation 5.19 the small Δ regime will be reached at ~ 10 times the impact parameter as for Coma—thus for the energies and impact parameters we are concerned with, the small angle approximation is never reached in A665. Outside of the small angle regime the conversion probabilities scale as

$$P(a \to \gamma) \propto \frac{1}{L}\left(\frac{B}{n_e}\right)^2, \tag{5.34}$$

since the conversion probability per domain is now independent of L, but the ALP passes through D/L domains, this should lead to another large reduction in conversion probabilities. However, the conversion in Coma was mostly in the small angle regime, and since probabilities in the small angle formula are naturally smaller (large Δ implies $\sin^2\Delta = 1/2$, whereas small Δ implies $\sin^2\Delta \approx \Delta^2 \ll 1$) the net result is a factor two drop in conversion probabilities for A665 compared to Coma.

From the above two factors the conversion probabilities are a large factor smaller in A665 compared to the previously considered Coma cluster. However, the increased size of A665, both in terms of the cross-sectional area entering the luminosity calculation, and increased propagation distance,[4] results in a luminosity that is a factor of ~ 9 smaller than that for Coma. The morphology is also completely different. In Coma the conversion probabilities were in the small-Δ regime and were thus larger at low radii, in A665 the opposite is true, see equation 5.21. The CAB conversion is

[4]We note that the finite volume of the simulation naturally affects the results. A larger propagation distance, assuming the magnetic field is still non-zero, will naturally lead to larger conversion probabilities due to the dependence on D of equation 5.19.

thus largest in the region where the thermal emission from the ICM is smallest.

As the canonical CAB parameter values we set the mean CAB energy to 150 eV, and the inverse coupling to $M = 7 \times 10^{12}$ GeV. For these values the total fractional excess across the whole cluster is $\xi = 0.025$, which is unobservably small, however, let us now discuss the morphology. Figure 5.12a shows the simulated and observed fractional excesses upon dividing the cluster into three annuli—note 6 arcminutes corresponds to 1090 kpc at the distance of A665. The errors on the observed data correspond to 68% confidence levels on the observed counts. The hot ICM emission at large radii is very small (9.1×10^{-3} s^{-1}) and thus despite the average excess being insignificant, the CAB does predict an excess at > 3 arcminutes.

The shaded regions for the simulation data points only take into account the statistical uncertainties on the random nature of the field. It was shown in the simulation of Coma in the previous section that upon generating several stochastic magnetic fields with the same power spectra and radial behaviour, the conversion probabilities at most varied by 10%. As discussed earlier, there is an additional uncertainty in predicting the counts in the ROSAT detector for a given luminosity. We will take this into account when constraining the CAB parameter space next.

The inverse coupling M, and the mean CAB energy $\langle E_{CAB} \rangle$, are unknown parameters in the model. Thus we now redo the analysis scanning over these two parameters, the resultant parameter space is shown in figure 5.12b. The red regions signify the parameter space where a significant excess is simulated—i.e. the simulated fractional excess would be distinguishable from zero at 95% confidence in [121], this corresponds to a count rate in the outer annulus of 2.6×10^{-3} s^{-1}. The light red region corresponds to allowing the count rate to be 50% larger than PIMMS predicts, the darker region corresponds to allowing the PIMMS prediction to be smaller by 50%, as discussed in section 5.8.1. The shape of the excluded region is easy to understand, for a given value of M, the largest luminosity occurs when most of the CAB spectrum is in the R1R2 energy band, hence forcing the tightest bounds on M. Note that we cannot push the CAB energy up arbitrarily high in order to not produce an unob-

served excess in energy bands above 0.4 keV. We see that the red regions include parts (all) of the best fit region found from simulations of the centre of the Coma cluster. The simulation thus predicts a significant soft excess, in the sense defined above, if $M \lesssim 6 - 10 \times 10^{12}$ GeV. Although we take time to point out that averaged across the whole cluster, the simulation agrees with observations that no excess would be observed.

The best fit regions from Coma and disfavoured region from A665 thus overlap each other, though we caution that the favoured or disfavoured values of M for both Coma and A665 have an additional uncertainty due to systematics on the magnetic field models. For instance as discussed earlier these turbulent, stochastic field models are likely to be simplified realisations of the true magneto-hydrodynamical ICM. In addition, the parameters of the model for A665 are constrained only by comparing simulated to observed radio halos, this requires several assumptions about the form of the relativistic electron population within the cluster. As a result, the observed anisotropies and shift in the radio halo centre and X-ray centre are assumed to be caused purely by the stochastic magnetic field (note a field with very large coherence lengths across the cluster will naturally be more anisotropic than one with very small coherence lengths) and not due to the form of the relativistic electrons. These assumptions allow the field parameters to be constrained, however it is likely that the resultant best fit parameters are highly sensitive to the form of the relativistic electron population. The slight overlap of parameter spaces is within the expected uncertainty arising from these concerns, especially given the dependence of the morphology on the magnetic field parameters, as seen in the Coma simulation.

5.8.2.2 A2199

A2199 is the first cool-core cluster considered in these simulations. The electron density in the centre of the cluster peaks at 0.1 cm^{-3}, which is $\sim$ 30 times that in Coma, A665 and A2255, and follows a double β-model profile, given in equation 5.12. The double β-model encompasses the extra, cooler (therefore high density) central

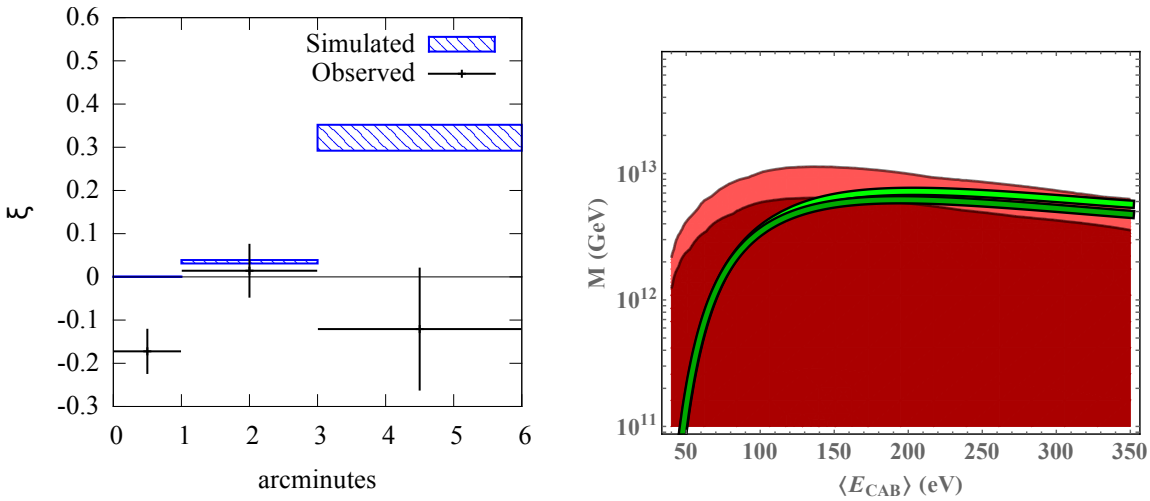

(a) (b)

Figure 5.12: a) The simulated and observed fractional excess in A665, the simulation takes $M = 7 \times 10^{12}$ GeV and $\langle E_{CAB} \rangle = 150$ eV. The observed data error bars are 68% limits on the observed counts, the shaded region accounts for statistical uncertainty on simulation. b) The M and E_{CAB} parameter space, shown in red is the region where a non-zero soft excess in the outer annulus would be observable at 95% C.L., overlaid in green are the best fit regions from the centre of Coma.

component of the electron density. We can see from equations 5.17 and 5.18 that outside of the small-Δ approximation

$$P(a \to \gamma) \propto n_e^{-2}, \tag{5.35}$$

and thus high electron densities suppress conversion. The cluster however has a very small core radius ($\mathcal{O}$(20 kpc) compared to $\mathcal{O}$(300 kpc) for Coma, A665 and A2255), and thus the electron density drops very rapidly with increasing radius, and thus we do not expect much suppression of conversion probabilities with respect to non-cool-core clusters. As we have argued before, large electron densities imply large magnetic fields, and thus A2199 has a large central magnetic field value of $B_0 \sim 12$ μG, much larger than the two other cluster's magnetic fields. However again this will drop rapidly from the centre due to the small core radius.

We show the simulated excess in A2199 in figure 5.13a—15 arcminutes at A2199 corresponds to $\sim$ 540 kpc. The blue points are the simulated fractional excesses for the magnetic field parameter values $(B_0,\ \eta) = (11.9\ \mu\text{G},\ 0.9)$. In this plot $M =$

7×10^{12} GeV and $\langle E_{CAB} \rangle = 150$ eV as usual, and again the shaded region takes into account the statistical uncertainty on the stochastic field. The CAB-generated luminosity in the $0.2 - 0.4$ keV band between $12 - 15$ arcminutes is simulated to be 4.7×10^{39} erg s^{-1}, which is more than an order of magnitude smaller than the observed luminosity (assuming thermal emission) of the soft excess, in the same energy range, of 1.7×10^{41} erg s^{-1} [121]. Thus for the Coma best fit parameters the CAB predicts no observable excess in A2199 for the canonical magnetic field parameters. This is mostly due to the incredibly fast fall of the magnetic field magnitude with radius when $\eta = 0.9$, and the fact that we simulate out to many times the core radius.

We first discuss changes to the magnetic field parameters before scanning over the CAB parameter space. As we mentioned earlier, the magnetic field in A2199 is poorly constrained. To be specific, the central magnetic field magnitude and the radial parameter η are constrained to be $B_0 = (11.9 \pm 9.0)\,\mu$G and $\eta = 0.9 \pm 0.5$. In the following we will consider varying these two parameters. These two parameters are degenerate with respect to Faraday rotation measurements, a larger central magnetic field value with a steeper fall with radius (higher η) produces a magnetic field with the same average value. Thus these two parameters must be changed in tandem. We do not consider any changes in the spectral index or the largest lengthscale $\Lambda_{\rm max}$, although we note that these also have a range of allowed values.

Within the uncertainties on $(B_0,\ \eta)$ we also show the simulated soft excess for the parameter choice $(B_0,\ \eta) = (6\ \mu\text{G},\ 0.5)$ plotted as the orange points in figure 5.13a. We choose the value $\eta = 0.5$ because as mentioned earlier, this corresponds to the case where the energy density in the magnetic field is proportional to the energy density of the thermal gas. The value of the parameter η effects the rate at which the magnetic field magnitude falls with radius, a smaller η then results in a larger magnetic field magnitude at large cluster radii, thereby increasing the conversion probabilities at large radii. We see this as an increased luminosity and hence increased soft excess at large radii. For these magnetic field parameters we see that the CAB parameters that fit Coma match the observed morphology (given the amount of scatter in the data)

and magnitude of the soft excess in A2199 well.

We next fit the magnitude of the soft excess by scanning over the (M, E_{CAB}) parameter space. We neglect the data in the $0 - 1$ arcminutes annulus due to the presence of a large deficit of soft X-rays. We show the region of CAB parameter space where the simulation fits the data, for both magnetic field parameters mentioned above, in figure 5.13b, where the thickness of the bands takes into account the count rate uncertainty discussed earlier. The dark regions are the regions where the total excess is reproduced, shown in lighter colours are the regions where the magnitude of the excess in the $12 - 15$ arcminutes annulus only is reproduced. We see then that given the uncertainty on the magnetic field strength and radial profile, the excess can be fit for a large range of inverse couplings, M. The CAB converting to photons in A2199 is consistent with Coma and A665 if the magnetic field magnitude does not fall as quickly with radius as the best fit field.

The current soft X-ray data contains a large amount of scatter and uncertainty. The total excess over the whole of the cluster is small (several %) and not significant, however there are radial regions where there is an observed excess, and there is a general trend that it increases with radius. We have shown that if the magnetic field falls less steeply with radius, the magnitude and general morphology are fit well, for a CAB with $M = 6 - 12 \times 10^{12}$ GeV. However there is large uncertainty in both the magnetic field and the soft excess observations. Clearly a more constrained magnetic field model, and hopefully future soft X-ray observations will shed more light on the CAB in A2199.

5.8.2.3 A2255

The simulated magnetic field in A2255 is an unrealistic 'stitching' together of two fields with different coherence lengths. For radii less than the core radius, the field has a flat power spectrum with the full range of scales available. For radii greater than the core radius, the field has a power spectrum index of $\zeta = 2$ and a minimum lengthscale of 64 kpc. A larger spectral index puts more of the power in smaller

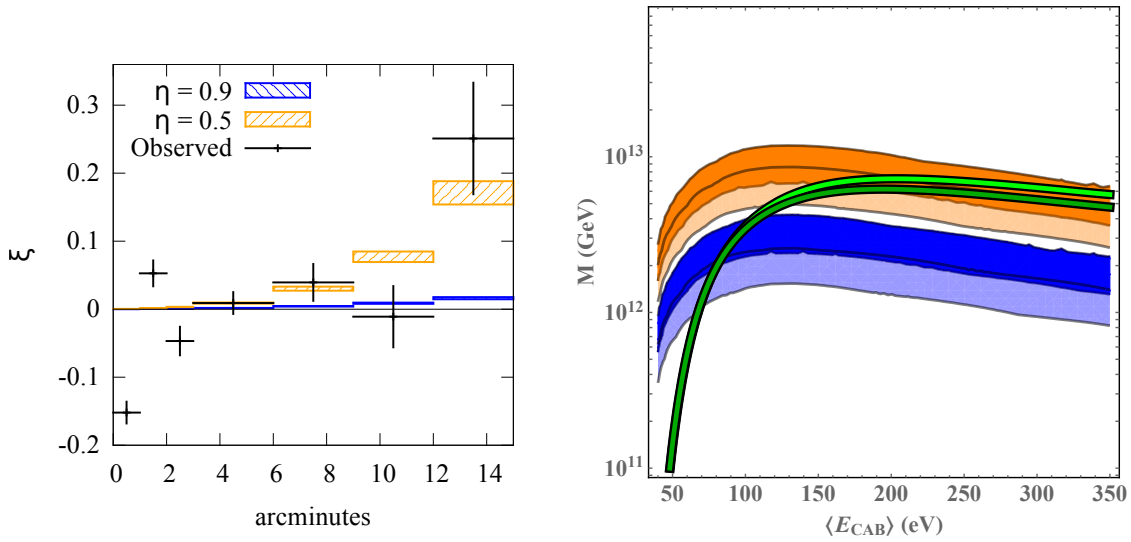

(a) (b)

Figure 5.13: a) The fractional excess in A2199. Both simulated data sets are taken with $M = 7 \times 10^{12}$ GeV and $\langle E_{CAB} \rangle = 150$ eV. b) The M v $\langle E_{CAB} \rangle$ parameter space which fits the data for the two models, defined as when the total observed fractional excess is equal to the total simulated fractional excess between $1 - 15$ arcminutes. The lighter colours indicate the regions where the excess is fixed to that in the $12 - 15$ arcminutes annulus. For comparison, overlaid in green are the best fit regions from the centre of Coma.

momentum modes. As a result the field in the centre of the cluster will on average have much shorter coherence lengths than in the outer parts of the cluster, which will typically have very large coherence lengths. We use a smoothing function at the core radius to interpolate between the two fields.

This magnetic field model then gives a unique morphology to the ALP−photon conversion probabilities. For ALPs whose impact parameter is small, $r < r_c$, the ALPs will pass through a large region where the field has small coherence lengths. The ALP conversion is then in the small-angle regime for a large proportion of the ALP's propagation, thus we expect conversion probabilities that decrease with increasing impact parameter, until $r \sim r_c$. For impact parameters larger than the core radius the ALPs will propagate only through the $\zeta = 2$ magnetic field, and the large coherence lengths here result in conversion probabilities which are out of the small Δ regime (although high energies and large impact parameter will reduce the angle Δ back into the small angle regime), and thus the conversion probabilities increase with radius.

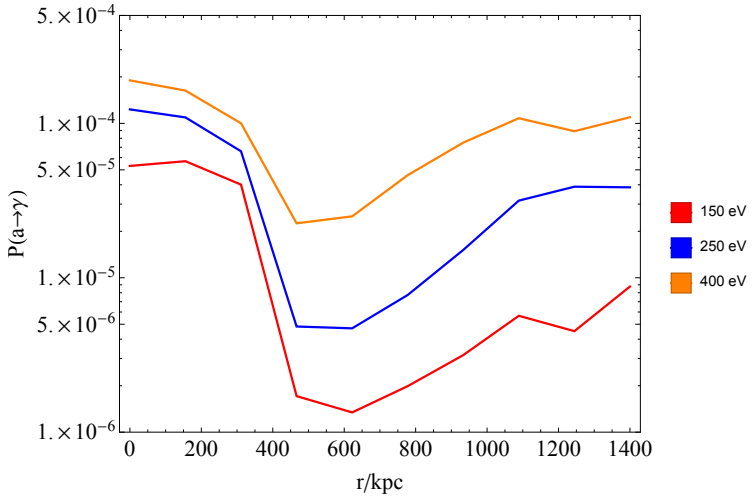

Figure 5.14: The ALP-photon conversion probabilities in the cluster A2255 for an ALP of impact parameter r, shown for three different ALP energies. We see the change in behaviour between ALPs that propagate mostly through the flat spectrum field at low impact parameter, and those that propagate only through the $\zeta = 2$ field at large impact parameter. The conversion probabilities are computed with $M = 7 \times 10^{12}$ GeV.

We show this feature of the conversion probabilities for three ALP energies in figure 5.14. We see the clear feature at $r_c \sim 400$ kpc where the ALPs are no longer passing through any of the smaller coherence length field. The conversion probabilities in the outer field are smaller relative to the centre due to the large coherence lengths resulting in the ALPs passing through few magnetic field domains.

In figure 5.15a we plot the simulated soft excess for the two radial parameters $\eta = 0.5, 0.7$.[5] Since the value of the radial parameter η is not well constrained we choose these two values as representative of the two interesting cases discussed earlier. In the plot we have again taken $M = 7 \times 10^{12}$ GeV and $\langle E_{CAB} \rangle = 150$ eV, and the shaded region is the statistical uncertainty. 15 arcminutes at A2255 corresponds to ~ 1.4 Mpc. We see that both field parameter choices reproduce the soft excess magnitude, and the morphology at radii < 9 arcminutes. The low conversion probabilities around ~ 500 kpc translate into a small, unobservable excess, exactly where the observations reveal no excess between $6 - 9$ arcminutes.

Both of the models then have large excesses at large radii due to the radial profile

[5]Note the magnetic field magnitude normalisation in A2255 is chosen slightly differently to that of A2199. According to [205] we normalise the field such that the average magnetic field inside the central 1 Mpc^3 is equal to 1.2 μG, which automatically ensures the value of B_0 changes according to the value of η chosen.

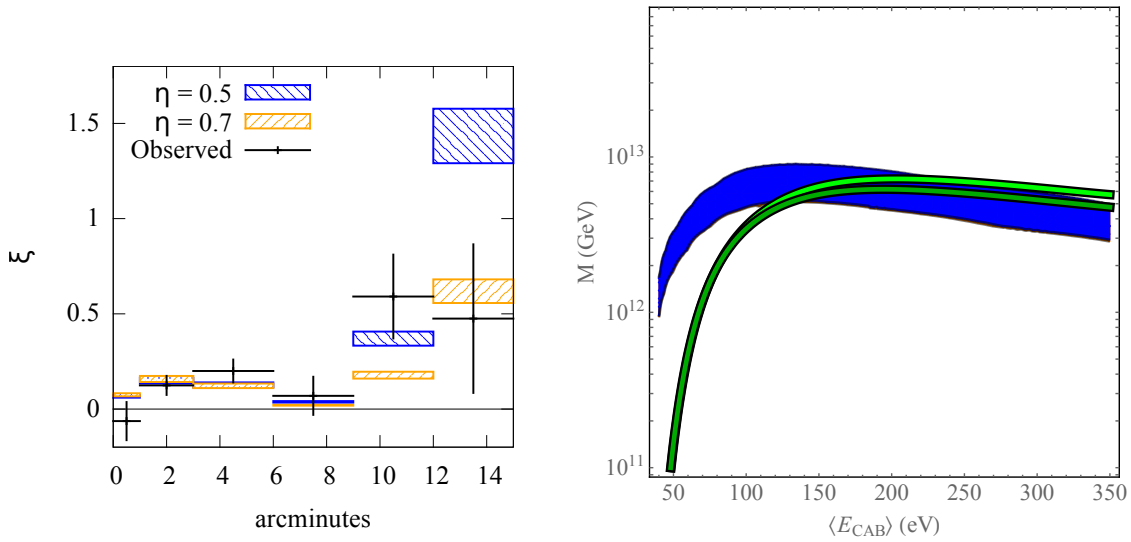

(a) (b)

Figure 5.15: a) The simulated and observed fractional excess in A2255. Plotted are the magnetic field models with radial parameter $\eta = 0.5$ in blue, and $\eta = 0.7$ in orange. For the simulations we take $M = 7 \times 10^{12}$ GeV and $\langle E_{CAB} \rangle = 150$ eV. b) The M v $\langle E_{CAB} \rangle$ parameter space which fits the data in the central radial zones (excluding > 9 arcminutes where the data is uncertain) for the two models (which overlap almost completely). For comparison, overlaid in green are the best fit regions from the centre of Coma.

of the conversion probabilities and the low expected ICM emission. The data at > 9 arcminutes have a large uncertainty and thus the morphology is not a concern, however the magnitude of the simulated excess for $\eta = 0.5$ at the outskirts is much higher than that observed. Increasing the radial parameter η results in the magnetic field dropping off more rapidly with radius, lower field values necessarily produce lower conversion probabilities and thus increasing to $\eta = 0.7$ does not overproduce soft X-rays and fits the data at > 9 arcminutes well.

In figure 5.15b we illustrate the fact that the CAB parameters that best fit the observed soft excess between $1 - 9$ arcminutes overlap the best fit CAB parameter region from the central part of Coma. Again the thick band takes into account varying the count rate prediction from PIMMS by $\pm 50\%$. We do not fit the large radius data due to the signal uncertainties mentioned earlier. The CAB morphology of the soft excess in A2255 clearly prefers a field which falls with radius more steeply than the canonical $B^2 \propto n_e$ choice.

Finally we note that the simulated magnetic field is not a realistic field. In addition to the usual systematic uncertainty due to assuming this stochastic model for the cluster magnetic field, the simulated field for A2255 is produced by stitching together two stochastic fields produced with different power spectra. Thus while the morphology of the simulated soft excess is in very good agreement with the observed excess, it is clear that this is mostly a feature of the field model used. The magnetic field transforms from a field with small coherence lengths to one with large coherence lengths over a short distance, and thus the conversion probabilities drop suddenly beyond the core radius. A better implementation of the field would be to have increasing coherence lengths as a function of radius, something that currently cannot be implemented as part of the simulation. A magnetic field with a spectral index which increases with radius and thus has coherence lengths which smoothly increase with radius would not have such striking distinction between small and large impact parameter as in figure 5.14. It is likely that the CAB-produced excess between $6-12$ arcminutes would not be as small, but would still provide a good fit to the observed morphology. In [212] the outskirts of Coma were analysed using a magnetic field whose coherence lengths increase with radius using a semi-analytical approach, we expect such a method is better suited to studying A2255.

5.8.3 Comparison and Summary

In this section we have continued to study the soft excess from a CAB in three more galaxy clusters. Here we compare the CAB parameter space between the clusters, and summarise our results.

We calculated the expected luminosities from CAB conversion by simulating the cluster magnetic fields and numerically calculating the conversion probability for an ALP of mass $m_a < 10^{-13}$ eV travelling through the field. We compared the simulation results to ROSAT PSPC count rates in the $0.1-0.28$ keV R1R2 band.

Observations of the cluster A665 have shown no excess, thus we can use this to bound the CAB parameter space. Averaged across the whole of the cluster, the

simulation predicts an unobservably small soft excess. However, a large excess is predicted at large radii, since here the conversion probabilities are highest, and the ICM emission is lowest. Stipulating that the simulated excess from a CAB at these large radii should not be larger than twice the error on the observed count rate, results in the disfavoured region

$$M < 6 - 10 \times 10^{12} \text{ GeV}, \tag{5.36}$$

for CAB mean energies $\langle E_{CAB} \rangle \sim 100 - 250$ eV. The range takes into account the uncertainty on where the bound lies due to the uncertainty in predicting count rates in the ROSAT PSPC detector. This disfavoured region overlaps with the preferred region from the centre of Coma ($M = 5 - 8 \times 10^{12}$ GeV), and thus the simulation of a CAB in A665 is in slight disagreement with the simulations of the Coma cluster.

Across the whole cluster there is no significant excess predicted, in agreement with the Coma simulation, however the morphology of the excess predicts a large excess at large radii. The morphology of the soft excess from a CAB is crucially dependent on the details of the magnetic field. However there are large uncertainties on the magnetic field model of A665, due to the assumptions about the relativistic electron population used to constrain the model. Thus further study of the magnetic field in A665 is needed to determine whether observations of A665 rule out the CAB explanation of the soft excess.

For the cluster A2199, a CAB produces a soft X-ray excess which would be very small (a factor of ~ 20 smaller than observed) for the best fit Coma values, for the canonical radial scaling parameter value of $\eta = 0.9$. We can alleviate this issue by decreasing this parameter within its allowed uncertainty, to $\eta = 0.5$. This results in a magnetic field magnitude which falls with radius more slowly and thus has a much more rapidly increasing soft excess with radius. The choice $\eta = 0.5$ then gives a good fit to the morphology of the observed excess. The magnitude of the excess is fit well for the lower η field for the parameters $M = 6 - 12 \times 10^{12}$ GeV and thus agrees with the analysis of Coma. Although we caution that there is a large amount of scatter in the A2199 soft excess data, and that the observed excess is at very low significance.

The simulations of A2255 also fit the observed excess well. The canonical choice of the radial scaling parameter $\eta = 0.5$, overproduces X-rays in the outer regions, but changing this value to $\eta = 0.7$ (which is allowed within observational uncertainties) produces an excess with the right morphology and magnitude for $M = 5 - 9 \times 10^{12}$ GeV, which agrees with A2199 and again with the Coma cluster. The morphology is fit very well, but this is mostly a feature of the magnetic field model used. The magnetic field model has an unrealistic sharp transition around the core radius from small coherence lengths to large coherence lengths. Where this transition happens the predicted excess dips due to a drastic reduction in conversion probabilities, and this is exactly where the observed excess is small.

In figure 5.16 we show the regions of the CAB parameters which fit the soft X-ray observations of the three clusters considered here, for comparison we also show the CAB parameter space which fits the soft excess in the centre and the outskirts of the Coma cluster. We show the good agreement between the clusters by showing in figure 5.16c the best fit values of M for all the clusters considered so far, setting (arbitrarily) $\langle E_{CAB} \rangle = 150$ eV. From the simulations of these three clusters we find the preferred range of the coupling M to be,

$$M = 6 - 12 \times 10^{12} \text{ GeV } \sqrt{\frac{\Delta N_{\text{eff}}}{0.5}}. \tag{5.37}$$

Which overlaps well with the parameter space from the centre of Coma, $M = 5 - 8 \times 10^{12}$ GeV, and from the outskirts of the Coma cluster, which requires $M = 5 - 30 \times 10^{12}$ GeV.

It has been observed that the cluster soft excess is preferentially seen at large radii [121]. In all three clusters we have studied in this section, the fractional excess produced from CAB conversion is largest at large radii. This is because the thermal emission from the ICM falls off faster than the conversion probabilities from the magnetic field. Thus we see that the CAB reproduces this observed morphological trend.

The study of these three clusters has been an important consistency check for

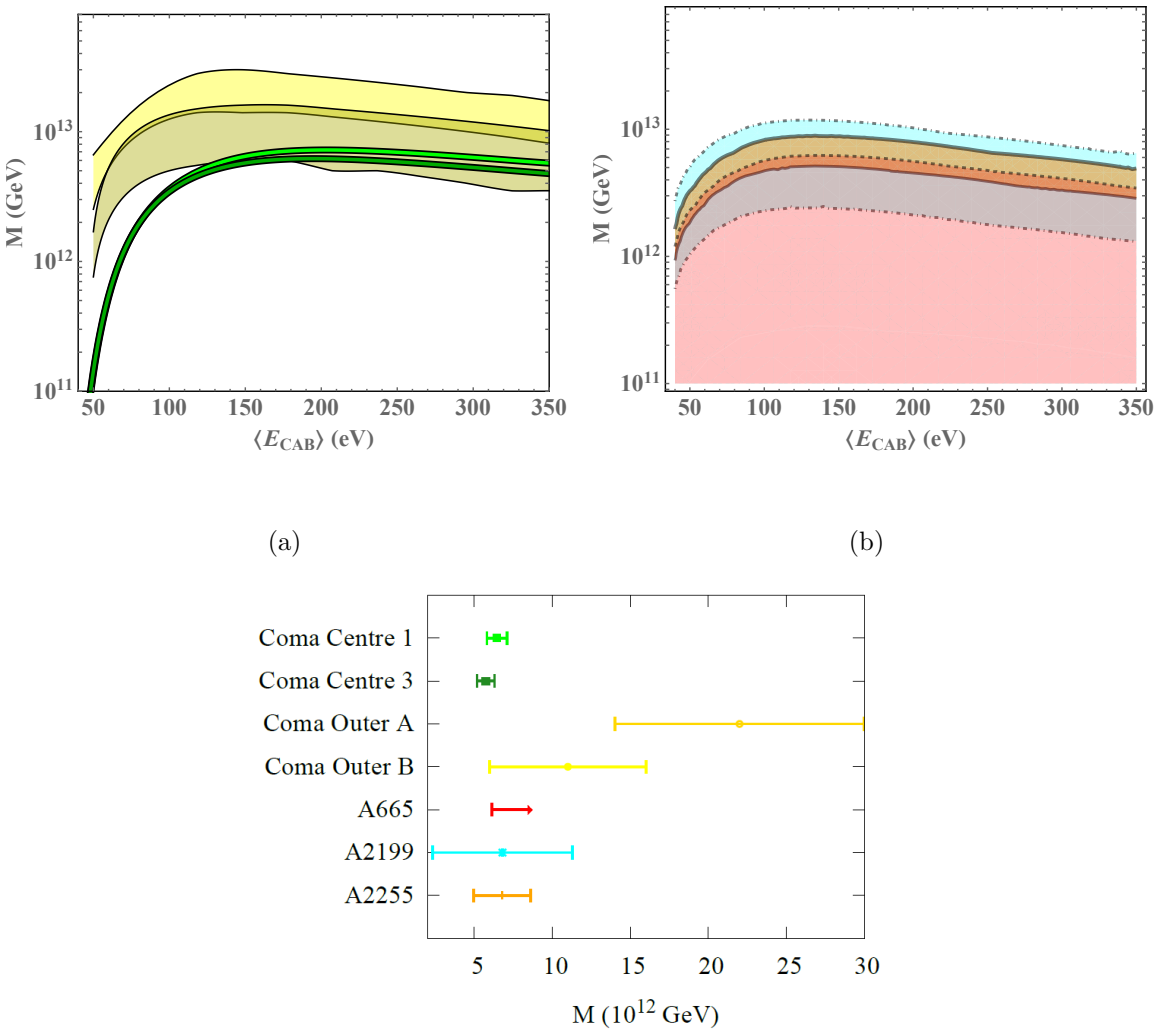

Figure 5.16: a) The best fit regions of the M vs $\langle E_{CAB} \rangle$ parameter space from the analyses of the centre (green) and the outskirts (yellow) of the Coma cluster. The light green and dark green regions correspond to Models 1 and 3, of the centre, respectively. Similarly the light yellow and dark yellow regions are Model A and Model B from the outskirts of Coma (which correspond to Models 1 and 3 from the centre). b) The best fit regions from the three clusters studied here: the disfavoured region from A665 is shown in red (dashed line); the best fit region for A2199 is shown in blue (dot-dashed line); and the same for A2255 is in orange (solid line). c) A comparison of the best fit values for M between simulations of the three clusters studied here, along with the best fit values from the Coma studies. The values of M have been taken corresponding to (the arbitrary choice) $\langle E_{CAB} \rangle = 150$ eV.

the CAB explanation of the soft excess. The best-fit CAB parameters agree fairly well between the four clusters. The three clusters studied here do not have as much X-ray observation time as the Coma cluster, nor are the magnetic fields as well determined, thus there are still large uncertainties in the results for these clusters. However

the consistency between the four clusters studied so far adds evidence to the CAB explanation of the soft excess.

5.9 Conclusions

In this Chapter I have used simulations to test the idea that an excess of soft X-ray emisson from galaxy clusters can be explained by a primordially-generated background of ALPs. These ALPs are produced by modulus decay at reheating, as studied in Chapter 4. Since they are very weakly-interacting they form a present-day homogeneous background. This background is probed indirectly by its contribution to dark radiation, but can also be probed directly by the ALP's coupling to photons. As the background passes through galaxy cluster magnetic fields the ALPs can convert to photons and produce a soft X-ray signal.

We have simulated the magnetic field in four clusters. The model we use for the magnetic field is a stochastically-generated, power-law model, that results in a tangled magnetic field. The magnetic field falls with radius with the same profile as the thermal electron density. We then numerically calculated ALP–photon conversion probabilities.

We started by studying the soft excess in the Coma cluster. Here a soft excess has been found at very high statistical significance, and the magnetic field properties are also well studied and constrained. We chose three sets of model parameters for the Coma cluster, two of which are a good fit to Faraday rotation observations, the other was used for illustrative purposes. The simulations reproduce well the magnitude of the observed excess in Coma, for an ALP–photon coupling $M \sim 5 - 8 \times 10^{12}$ GeV. We also studied the morphology of the ALP–photon conversion probabilities for the models and found that the morphology of the excess can be fit well if the magnetic field has a flat power spectrum, or predominantly short coherence lengths.

We have also studied the clusters A665, A2199, and A2255 as an important consistency check. A665 has no observed excess and we find the exclusion region from this cluster agrees well with the values of M which reproduce the observed (low sig-

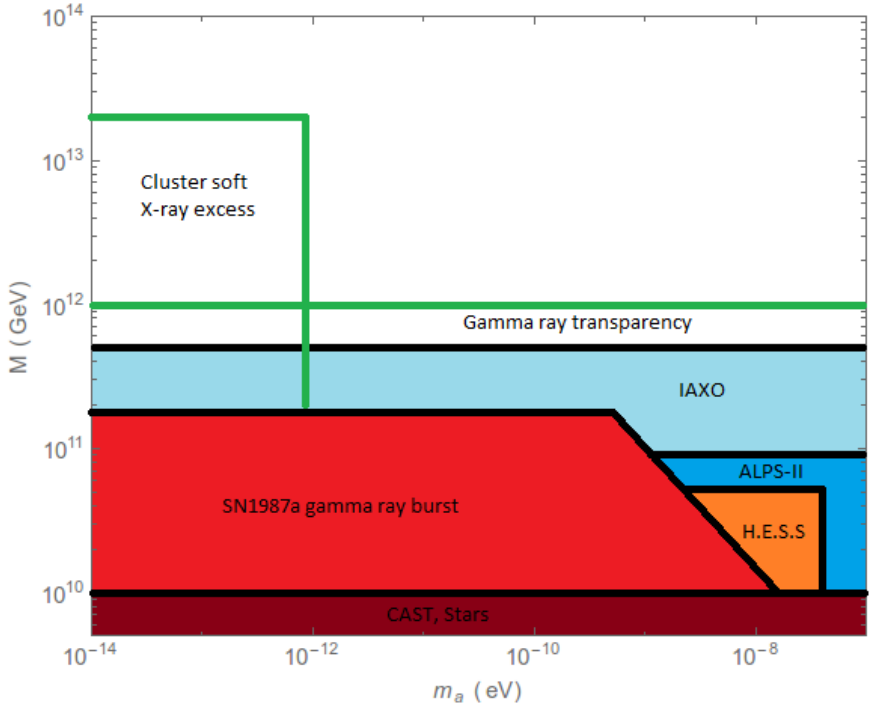

Figure 5.17: The low mass ALP parameter space showing the bounds on the ALP mass m_a and ALP−photon inverse coupling M. The strongest laboratory bound comes from the experiment CAST (dark red) [71], an even stronger astrophysical bound is from the lack of a gamma ray burst associated with SN1987a (red) [45–47], we have also included the bound from the H.E.S.S. study of PKS 2155-304 [55]. Indicated in green are the favoured regions for an ALP explanation to the soft X-ray excess, from this work and [212], and to the anomalous gamma ray transparency of the universe [56, 57, 60, 63, 65], bound from [62]. We have also indicated in blue the regions expected to be probed by the laboratory experiments ALPS-II [69] and IAXO [72].

nificance) excesses in A2199 and A2255. There is slight tension between the exclusion region derived from A665 and the results from the Coma simulation, however better knowledge of the magnetic field in A665 is needed to test this further.

The best fit regions across the four clusters studied here are in agreement with each other, and indicate an ALP–photon coupling $M = 6 - 12 \times 10^{12}$ GeV. We show how this fits in with current bounds in figure 5.17. We note these values are obtained assuming $\Delta N_{\rm eff} = 0.5$, the best fit M scales as $M \sim (\Delta N_{\rm eff})^{1/2}$, and we include this in the figure.

While the simulations conducted so far have shown good agreement between CAB parameters across the clusters studied, the uncertainties on the magnetic fields of clusters limit further analysis. The four clusters studied so far exhaust the list of clusters which have both soft X-ray excess observations and constrained magnetic field

power spectra and radial profiles. On top of this it is possible that these magnetic fields models do not capture the full magneto-hydrodynamical structure of the real cluster magnetic fields. Thus it would be interesting to take magnetic field values found from a numerical MHD simulation of cluster formation and see whether there are any qualitative differences in the resulting ALP-photon conversion probabilities.

While the study of the soft excess from a CAB is limited at the moment, this should improve with time as cluster magnetic fields are determined for more clusters and with greater accuracy. Currently, for some of the clusters we studied, to constrain the magnetic field, assumptions about the relativistic electron population that produces the observed radio halo need to be made. In the future it is hoped that increased Faraday rotation data in many clusters can reduce the need for such assumptions. If information can be obtained about the magnetic field and electron density in many clusters from the sample [121], it would be interesting to see whether this can be correlated with the presence or absence of a soft excess. Furthermore, we have seen that the predicted soft X-ray morphology is intimately linked to the magnetic field parameters. Thus whilst we have shown a good match to the observed morphologies, better knowledge of the cluster magnetic fields will allow us to study the soft X-ray excess in more detail.

One of the most interesting results from this section is that galaxy cluster magnetic fields are very efficient at converting X-ray ALPs into photons. In addition we have shown that the X-ray regime is exactly in a crossover region between the large and small Δ regimes, which give qualitatively different conversion probabilities. In the next chapter we will propose a way to use galaxy clusters to search directly for ALPs, using only the X-ray cluster emission. Photon conversion to ALPs will lead to significant distortions of the X-ray spectrum of a galaxy cluster.

The effects of a cosmic ALP background have since been studied in various other astrophysical and cosmological settings. In [216] a detailed study of CAB conversion in the Milky Way magnetic field showed that a CAB would give a signal several orders of magnitude below the cosmic X-ray background, rendering it unobservable

from the Milky Way. A strong constraint on the properties of the ALP that make up the CAB was found in [217]. Conversion of CAB ALPs into MeV photons in the early universe would cause a significant amount of reionisation, before the reionisation epoch ($z \sim 11$). If there existed a primordial magnetic field during the reionisation period, whose current magnitude is $B \sim 1$ nG, coherent over Mpc scales, then values of M almost up to the Planck scale are ruled out. However as discussed earlier, it is not clear whether these assumptions about the cosmological magnetic field hold.

6 Constraining ALPs using Galaxy Clusters

In this chapter I use what we have learnt in the previous chapter about ALP-photon conversion in galaxy clusters to study the opposite phenomenon. The X-ray photons emitted by the galaxy cluster itself may convert to ALPs in the cluster's magnetic field, this effect is energy-dependent and can be searched for by looking at modifications to the galaxy cluster's thermal bremsstrahlung spectrum. This is independent of any string theory model, as it does not require the production of ALPs in the early universe. Here I present the first study of this effect, and its prospects for constraining ALP parameters, which was first published in [4], written in collaboration with Joseph P. Conlon and David Marsh.

6.1 Current Bounds

As we have already seen in Chapter 5, galaxy clusters are highly efficient at inter-converting light axion-like particles (ALPs) and photons [2, 3, 116] (see [64, 218, 219] for some earlier work). In this chapter, we show that the absence of large distortions of the cluster thermal X-ray bremsstrahlung spectrum may be used to derive the strongest bounds to date on the ALP–photon coupling for light ALPs. These bounds

are stricter than the current bound $M \gtrsim 2 \times 10^{11}\,\mathrm{GeV}$ [45–47] from SN1987a, which holds for $m_a \lesssim 10^{-10}$ eV, and better than the planned next generation earth-based experiments ALPS-II and IAXO, $M > 3 \times 10^{11}\,\mathrm{GeV}$. For a review of the bounds on the ALP photon coupling, see section 3.3.

We also note that the method used here is similar to that used in previous studies. For instance, a recent paper used the absence of CMB distortions through clusters to produce strong bounds on the ALP photon coupling $M \gtrsim 10^{11}$ GeV [220], although as this relies on resonance effects it is only relevant for a small range of ALP masses around $m_a \sim 10^{-13}$ eV. A similar method was used in [221], which constrained ALPs by looking for distortions to an AGN spectrum shining through the Hydra A cluster. The resulting bounds were slightly less constraining than from SN1987a. Here we use the cluster emission itself to get bounds which are more stringent. We argue that current and future data from X-ray observations of galaxy clusters can significantly improve the bounds on M for the entire small m_a region ($\lesssim 1 \times 10^{-12}\,\mathrm{eV}$), potentially reaching to around $M \sim 10^{12}$ GeV.

I start in section 6.2 by restating photon–ALP conversion in a language more suitable to a discussion of the effects we consider. Then, after revising X-ray emission from galaxy clusters, I discuss in section 6.3 the expected distortions to this signal arising from photon–ALP conversion. In section 6.4, I look at a few more properties of the distorted spectrum, such as its dependence on the size of the extracted signal region and potential polarisation signals. Finally, I conclude in section 6.5.

6.2 Photon-ALP Conversion in Galaxy Clusters

Firstly, let us review the relevant theory of ALP–photon conversion for this chapter.

We found in the previous chapter that even in complicated magnetic field models, the conversion can be understood well by looking at the probabilities in constant magnetic field domains. For photons travelling in a medium with plasma frequency

$$\omega_{\mathrm{pl}} = (4\pi\alpha n_e/m_e)^{1/2} = \left(\frac{n_e}{10^{-3}\mathrm{cm}^{-3}}\right)^{0.5} 1.2 \times 10^{-12}\ \mathrm{eV}, \tag{6.1}$$

and a constant magnetic field in a domain of length L, the photon-to-ALP conversion probability for unpolarised light can be written as[1],

$$P_{\gamma\to a} = \frac{1}{2}\frac{\Theta^2}{1+\Theta^2}\sin^2\left(\Delta\sqrt{1+\Theta^2}\right), \tag{6.2}$$

where $\Theta = \frac{2B_\perp\omega}{Mm_{eff}^2}$, $\Delta = \frac{m_{eff}^2 L}{4\omega}$, and $m_{eff}^2 = m_a^2 - \omega_{pl}^2$. Here $B_\perp$ denotes the component of the magnetic field that is perpendicular to the ALP/photon wave vector. The factor of $\frac{1}{2}$ accounts for the fact that only one polarisation state of light participates in the mixing.

The typical values of Θ and Δ in galaxy clusters, for $m_a = 0$, are

$$\begin{aligned}\Theta &= 0.28\left(\frac{B_\perp}{1\,\mu\mathrm{G}}\right)\left(\frac{\omega}{1\,\mathrm{keV}}\right)\left(\frac{10^{-3}\,\mathrm{cm}^{-3}}{n_e}\right)\left(\frac{10^{11}\,\mathrm{GeV}}{M}\right),\\ \Delta &= 0.54\left(\frac{n_e}{10^{-3}\,\mathrm{cm}^{-3}}\right)\left(\frac{L}{10\,\mathrm{kpc}}\right)\left(\frac{1\,\mathrm{keV}}{\omega}\right).\end{aligned} \tag{6.3}$$

There are two different regimes for the conversion probabilities in clusters, which we saw in our discussion of the radial dependence of the conversion probabilities in Coma in Chapter 5. For $\Theta, \Delta \ll 1$, conversion is quadratic in both size and coherence length of the magnetic field,

$$P_{\gamma\to a} = \frac{1}{2}\Theta^2\Delta^2 = \frac{B_\perp^2 L^2}{8M^2}, \tag{6.4}$$

and ALP-photon conversion is energy-independent. However, for $\Theta \ll 1$ and $\Delta \gg 1$, the conversion probability is energy-dependent and progressively suppressed at lower energies,

$$P_{\gamma\to a} = \frac{1}{2}\Theta^2\sin^2(\Delta) \sim \frac{m_e^2 B_\perp^2}{\pi^2\alpha^2 n_e^2 M^2}\omega^2\sin^2\left(\frac{\pi\alpha}{m_e}\frac{n_e L}{\omega}\right). \tag{6.5}$$

Keeping the other parameters fixed, we can define a critical energy marking the crossover between the regimes,

$$\frac{\omega_\Delta}{\mathrm{keV}} = 0.54\left|\frac{n_e}{10^{-3}\,\mathrm{cm}^{-3}} - \left(\frac{m_a}{1.2\cdot 10^{-12}\,\mathrm{eV}}\right)^2\right|\left(\frac{L}{10\,\mathrm{kpc}}\right), \tag{6.6}$$

and thus for typical clusters and for $M \gtrsim 10^{11}\mathrm{GeV}$[2], the crossover occurs within

[1]Note this is just a rewriting of the single domain formula from Chapter 5, noting the different definitions between θ and Θ, with $\tan 2\theta = \Theta$, making the physics relevant to this chapter more obvious.

[2]If M is taken too small then we actually enter the large Θ regime, in this case the conversion is

the range of observed X-ray photon energies. This effect can be easily seen in, for example, figure 5.6 for the Coma cluster. Looking at the conversion probabilities as a function of energy for zero impact parameter, there is a clear transition from small to large conversion probabilities as a function of energy.

It can also be seen using equation 6.5 that the energy dependence of the conversion probabilities will lead to very striking signatures. The conversion is oscillatory with a wavelength proportional to $1/\omega$. Thus, at small energies the conversion is rapidly oscillatory and stochastic, along each different sightline the photon/ALP wave vector will sample different magnetic field values and electron densities, leading to very different conversion probabilities.

We now show that this effect can lead to substantial distortions of the thermal spectrum of the intra-cluster medium, using the Coma cluster as our prime example.

6.3 Cluster Spectral Distortions from ALPs

The intra-cluster medium (ICM) permeating galaxy clusters is a hot thermal plasma with temperatures, depending on the cluster, of between 2 and 10 keV. The cluster is visible in X-rays through the thermal bremsstrahlung of the ICM, with both continuum and line emission.

Figures 6.1 and 6.2 show fits to emission from the Coma cluster (adapted from [222] and [223]), which hosts an approximately isothermal ICM with a temperature of 8.1 keV. Note the excellent quality of the fits, with residuals all below the 10% level. As a further illustration of the attained precision, we note that the possible 3.55 keV dark matter line reported in [224] is observed as a per cent level effect above the background, thereby requiring the ability to characterise the background thermal emission to the same level of accuracy.

Thus, galaxy clusters provide an intrinsic, diffuse, bright and well characterised source of X-ray photons. If a significant proportion of the thermal photons were to convert to ALPs, the resulting X-ray spectrum would be distorted. The photon-

maximal across the entire X-ray regime.

to-ALP conversion probability is determined largely by the structure of the cluster magnetic field.

As we saw in Chapter 5, galaxy clusters support magnetic fields of $\mathcal{O}(\mu\mathrm{G})$ strength that are coherent over 1–10 kpc scales (for reviews see [188, 189]). For the Coma cluster, we constructed a detailed model of the cluster magnetic field in the previous chapter.

Using this simulation of the Coma magnetic field, we have simulated photon-to-ALP conversion of the ICM by numerically solving the equation of motion, as in Chapter 5, along a set of sightlines that sample a given field of view. We calculate conversion probabilities for photons originating at each point in the cluster along the line of sight. The net conversion probability for each sightline is given by an average weighted by the thermal photon emissivity at each starting point. This emissivitiy scales only as $\sim n_e^2$, assuming there are no temperature gradients. For a given field of view we also perform the same weighted average of conversion probabilities over all sightlines with the field.

The distorted X-ray spectrum is then given by,

$$f_{\mathrm{distorted}}(\omega) = (1 - \langle P_{\gamma\to a}\rangle(\omega))\, f_{\mathrm{intrinsic}}(\omega)\,, \tag{6.7}$$

where $\langle P_{\gamma\to a}\rangle(\omega)$ denotes the average conversion probability for the sightlines within the field of view.

For $M \lesssim 10^{11}\,\mathrm{GeV}$, the conversion probability saturates for the entire X-ray range, leading to a uniform reduction of the cluster luminosity to 2/3 of its original amount. However, as shown in equation 6.6, for precisely the most observationally interesting range of $10^{11} < M/\mathrm{GeV} \lesssim 10^{12}$, the photon-to-ALP conversion probability is large at higher X-ray energies, while being suppressed at lower energies. Thus, such photon-to-ALP conversion induces spectral distortions of the thermal ICM spectrum.[3] In figure 6.3 we illustrate this by plotting $\langle P_{\gamma\to a}\rangle(\omega)$ obtained from the simulations with $m_a = 0$, $M = 4\times10^{11}$ GeV and a field of view of size $(100\,\mathrm{kpc})^2$ at the centre of the

[3]We note that this effect is unlikely to generate the cluster soft X-ray excess, as that excess is observed at lower energies.

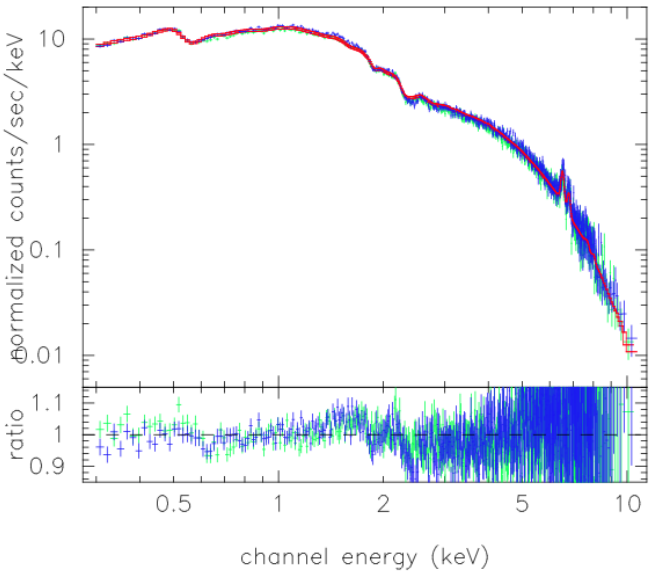

Figure 6.1: XMM-Newton best fit flux model from Coma, within a radius of $\sim 300\,\mathrm{kpc}$ [222].

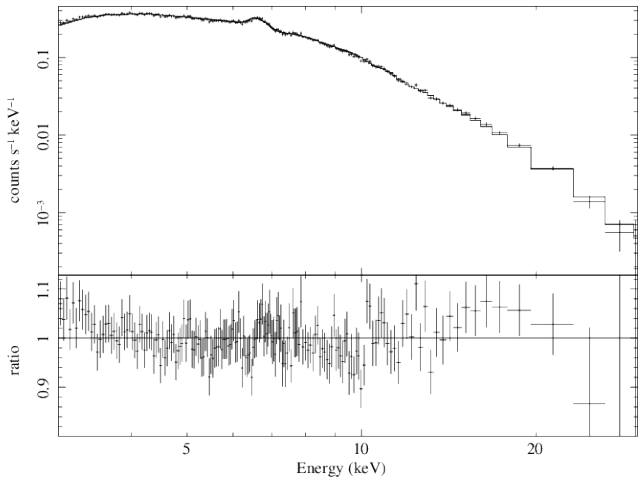

Figure 6.2: NuSTAR spectrum and best-fit single temperature model for the central $12' \times 12'$ of the Coma cluster [223].

Coma cluster. For such a large field of view, the oscillatory and stochastic behaviour mentioned earlier gets averaged out to the step-like distortion we see in figure 6.3.

Given the quality of the fits in figures 6.1 and 6.2, the simulated photon-to-ALP conversion probabilities can be used to constrain the ALP parameters M and m_a. We can estimate the ability of cluster spectra to probe the ALP–photon coupling by assuming deviations in the spectrum of 10% or greater would be observable, we plot the results in Figure 6.4. For $m_a \lesssim 5 \times 10^{-12}\,\mathrm{eV}$ we estimate that this method is sensitive to ALP couplings up to $M \sim 7 \times 10^{11}\,\mathrm{GeV}$, while for $m_a \gtrsim 5 \times 10^{-12}\,\mathrm{eV}$, the conversion probability starts getting suppressed by the ALP mass, and for $m_a \gtrsim$

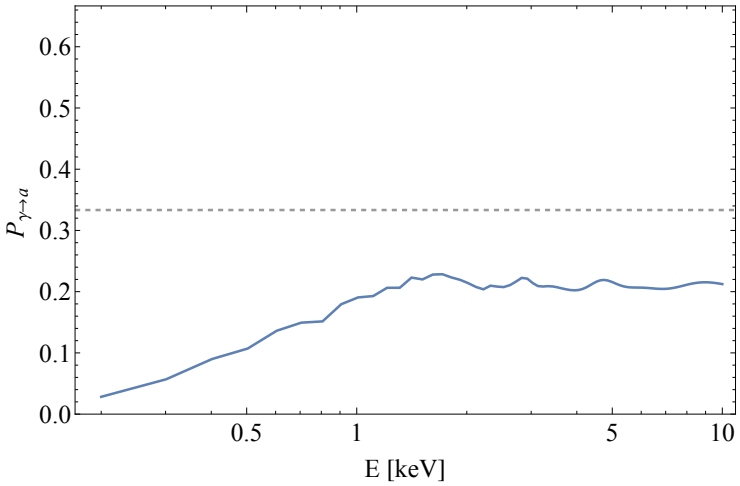

Figure 6.3: Averaged conversion probabilities for $M = 4 \times 10^{11}\,\mathrm{GeV}$, $m_a = 0\,\mathrm{eV}$, for a 100 kpc field of view. The dashed line represents the saturation value of 1/3.

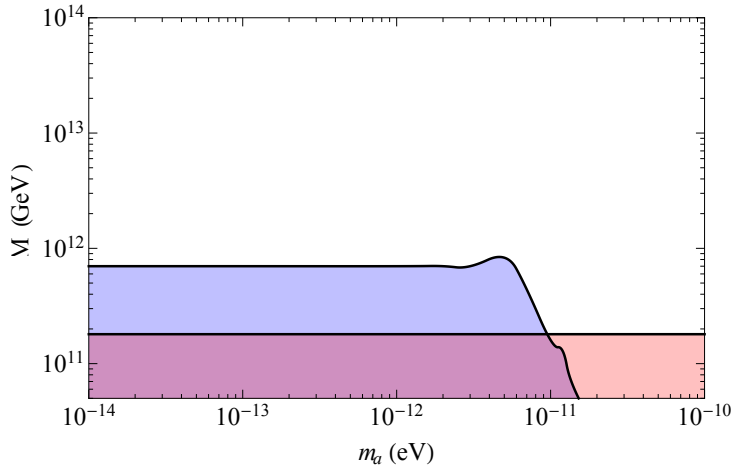

Figure 6.4: Astrophysical bounds from the lack of an observation of a gamma ray burst with SN1987a in red, bounds on the parameter space from this study in blue. Because of uncertainties in the cluster magnetic field, the precise values of M excluded may be uncertain by a factor of two.

$1 \times 10^{-11}\,\mathrm{eV}$ the estimated sensitivity to the coupling becomes weaker than existing astrophysical constraints.

We therefore conclude that constraints on spectral deviations of the cluster thermal bremsstrahlung spectrum can provide the strongest astrophysical bounds to date on the ALP-photon coupling, excluding an interesting range of parameters that will be probed in upcoming laboratory experiments.

6.4 Further Properties of the Distorted Spectrum

There are several further properties, and thus potentially correlated signals, of an ALP-distorted cluster spectrum.

First, as discussed earlier, along a single line of sight photon-to-ALP conversion is highly stochastic and energy-dependent. This energy dependence varies between each line of sight and gets averaged over large fields of view. Thus there are additional potential signatures observable only for small fields of view. In figure 6.5 we illustrate this by showing the averaged conversion probability for a $(5\,\mathrm{kpc})^2$ field of view. Since the magnetic field is coherent over $\mathcal{O}(10\,\mathrm{kpc})$ scales, the photon–ALP conversion varies little across such a small field of view. This allows the energy-dependence of the conversion probabilities to be seen. Note that since the simulated field is random in nature, the exact energy dependence of the real field will differ, but the presence of these sinusoidal oscillations at X-ray energies is a characteristic feature of ALP-photon conversion, with the frequency of these oscillations decreasing with increasing energy.

When averaging over a large field of view of size far greater than the coherence length of the cluster magnetic field (as in figure 6.3), these variations wash out, due to the differing energy denpendence of the different sightlines. However, by extracting the cluster X-ray spectrum across a small region it may be possible to directly observe such oscillations. In fact such a study would be more sensitive to the ALP–photon coupling as it is hard to see what astrophysical effect could mimic such oscillatory features, and for single sightlines these features are more pronounced.

The best candidates for this kind of study are very nearby clusters, where the fixed telescope angular resolution corresponds to the smallest physical scales. For example, the Virgo cluster is at a redshift of $z = 0.004$ and there is approximately 500 ks of observation time on Virgo with the Chandra telescope, whose superb imaging capabilities gives it arcsecond resolution. By considering spectra extracted from small angular regions instead of the full field of view, one could search for the presence (or

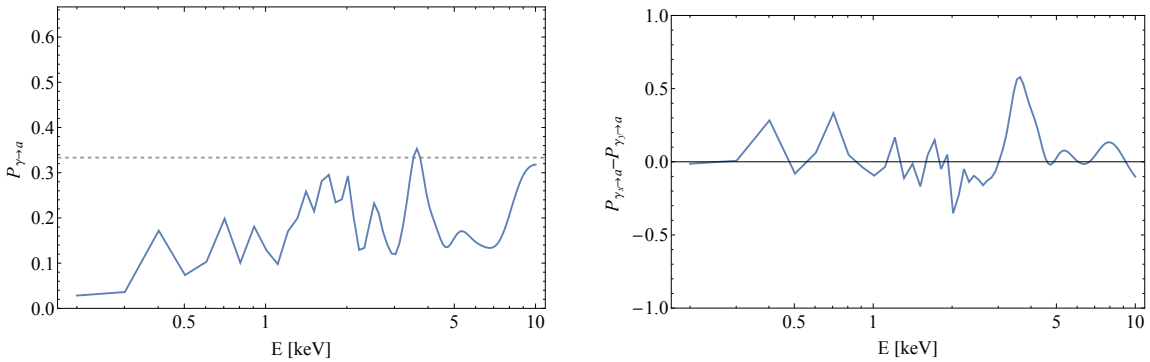

Figure 6.5: *Left*—The conversion probabilities averaged over both polarisations. *Right*—The difference between conversion probabilities for the two polarisations of X-rays, showing the induced polarisation caused by ALP–photon oscillations.

absence) of such oscillatory features.

There is another approach to searching for these small scale features. Clusters are large Mpc-scale objects, and gradients in their internal structure will generally also be large-scale. In the absence of photon-ALP conversion, the bremsstrahlung photon count from nearby pixels on a detector is expected to be identical, subject to the variations due to Poisson statistics. Photon-ALP conversion can lead to significant variations in conversion probabilities from nearby sightlines, and so in the presence of ALPs the photon count from nearby pixels should have a greater variation than would be expected simply from Poisson statistics.

Furthermore, along a single line of sight the X-ray emission becomes highly polarised. In figure 6.5 we show that with a small field of view this induced polarisation would be observable, but again for larger field of views this effect gets washed out. The result is that for $10^{11}\,\mathrm{GeV} \lesssim M \lesssim 10^{12}\,\mathrm{GeV}$, the diffuse thermal ICM emission extracted from a small region would be polarised. If an X-ray telescope with polarisation capability, such as the European LOFT or Chinese XTP missions, is (finally, after forty years) flown, the presence or absence of polarised emission from the ICM could lead either to the detection of ALPs or strong bounds on their coupling. While optical polarisation induced by ALP–photon oscillations has been considered before (see for example [225] studied the induced polarisation in the local super-cluster magnetic field), the advantage of a galaxy cluster environment is efficient X-ray conversion and better studied magnetic fields.

The small-scale spectral oscillations can also be sought by considering X-ray point sources that are fortuitously observed through a cluster. Such a search has been carried out for the Hydra A cluster in [221] using a central AGN source. However this is not an ideal source. Hydra A is at a redshift of $z = 0.06$, more than ten times further away than e.g. the Virgo cluster, and furthermore the AGN is optically thick, with significant intrinsic absorption at lower X-ray energies, making it harder to separate photon-ALP leakage from this effect.

Another interesting effect would be that the spectral distortions would have a radial dependence within a cluster. The free electron density decreases with radius, and observations indicate the magnetic field strength also decreases with radial distance from the centre of the cluster. Photons travelling along off-centred sightlines will then convert at a reduced rate at higher energies (see equation 6.4), and the 'step' in the conversion probability appear at lower energies, as implied by equation 6.6. This behaviour would be characteristic of ALP-induced distortions.

6.5 Conclusions

In this chapter we have built on the knowledge of ALP–photon conversion in galaxy cluster magnetic fields gained from the previous chapter.

Photon-to-ALP conversion is energy-dependent, and thus leads to distortions of X-ray spectra. The thermal emission from galaxy clusters provides a well characterised X-ray source, arising within a magnetic field environment that would get significantly distorted for ALP masses $m_a \lesssim 1 \times 10^{-12}\,\mathrm{eV}$ and ALP-photon couplings $10^{11}\,\mathrm{GeV} \lesssim M \lesssim 7 \times 10^{11}\,\mathrm{GeV}$. Thus dedicated searches for these distortions should be able to place highly competitive, and potentially the most stringent, bounds on the ALP parameter space. We should of course also caution that bounds on M can only ever be as good as knowledge of the astrophysical magnetic field, which for clusters is probably currently uncertain up to a factor of two.

We have used the absence of any reported large deviations from a thermal spectrum to estimate the ability of current X-ray satellites to probe the ALP parameter space.

Existing analyses however are not optimised for ALP searches, since the spectra is taken from a large field of view. In this case the large spectral distortions are averaged out due to the different energy dependence of different line of sights. The result is a step-like feature in the photon spectrum. We can however use this to estimate sensitivity to the ALP–photon coupling. An optimised search strategy would be to extract spectra over a smaller field of view. Here the distortions are larger and have a striking oscillatory energy dependence.

An attractive reason to study X-ray emission from clusters and its connection to ALPs is that the data is only going to get better. The three current satellites XMM-Newton, Chandra, and Suzaku, all use CCD detectors with an intrinsic resolution of $\Delta E \sim 100\,\mathrm{eV}$, and these satellites all continue to make observations of clusters. Over the long term, ATHENA has been approved by ESA for a 2028 launch. It was hoped that the new satellite Hitomi would provide increased sensitivity to ALPs. The satellite had a far superior energy resolution of 7 eV. However the satellite has now been lost, after making only one observation of the Perseus galaxy cluster. The spatial resolution of Hitomi is poorer than the three satellites mentioned above, for Perseus this corresponds to ~ 20 kpc, which is probably larger than the typical coherence length of the magnetic field, making it sub-optimal for searching for these distortions. Currently the only data is a full field of view spectrum, roughly $50\ \mathrm{kpc} \times 50\ \mathrm{kpc}$. The full data has not yet been released, and depending on the size of the effects, it may still be possible to see something by studying the spectra from individual pixels.

Knowledge of cluster magnetic fields will also be enhanced over the short, medium and longer term. Through observing (or not) upscattered inverse Compton-CMB photons, hard X-ray imaging telescopes will either produce direct measurements of the magnetic field, or more stringent lower bounds on the field strength. Over the next decade, the Square Kilometer Array will also come online. With a large increase in the number of available radio sources, this will improve Faraday rotation measurements of galaxy clusters' magnetic fields.

In the short term however, it should be possible to use data from the current

generation of satellites to perform a systematic search for ALPs in galaxy clusters. Here we have only estimated the bound based on observations of the Coma cluster with a large field of view. Looking for distortions in the spectra from small regions of the cluster can be combined with the fact that there are many such regions within each cluster to provide a very sensitive probe of ALPs, and a much more robust determination of the bounds on the ALP–photon coupling. In addition, since the energy dependence is most striking for single sightlines, looking for distortions in point source spectra is also a promising method to search for the presence of ALPs.

Indeed, we have used this method to look for ALPs in the spectra of a point source, the AGN in the centre of the Perseus cluster. This was published in [6], and was qualitatively the same as the method used by [221]. Hints of distortions in the spectrum caused by ALPs were found, hinting at an ALP with coupling to photons of $M \sim 5 \times 10^{11}$ GeV.

7 Summary and Conclusions

In this thesis I have studied various cosmological and astrophysical implications of axion-like particles (ALPs). ALPs are a generic prediction of string theory compactification models, and thus their study provides a window into string phenomenology. I have looked at their production via modulus decay in the early universe as a way to test string theory models. ALPs can be searched for using their potential coupling to electromagnetism, which means that in an external magnetic field they can convert into photons, and vice versa. I looked at two possible astrophysical signatures of ALPs, the soft X-ray excess in galaxy clusters, and spectral distortions in the thermal X-ray emission from galaxy clusters.

In Chapter 4 we studied ALP production in the early universe in a class of string theory compactification models called the Large Volume Scenario (LVS). Displacement of string moduli fields during inflation leads to periods where the modui dominate the energy density of the universe. Thus in string models the moduli are typically responsible for reheating. The LVS is tractable to analyse as it has one unique lightest modulus, and thus reheating is driven solely by this modulus' decay. We started by reviewing the tree level computation of the decay rates to different sectors. The two important decay modes are to Higgs pairs, which is the only visible sector decay and

is thus responsible for reheating, and to ALPs. The ALPs produced in this decay act as dark radiation, additional relativistic energy density not predicted by the Standard Model. We showed that in the minimal LVS (MLVS), with only one ALP, an MSSM visible sector, and a coupling of the modulus to the Higgses of $Z = 1$, dark radiation is over-produced compared to observations.

In view of the large hierarchy between the string scale, where $Z = 1$ is defined, and the modulus mass, the coupling Z can receive potentially large radiative corrections. We computed these analytically and found that the contributions from gauge boson loops and top loops tend to cancel, and thus change Z by only a small amount. We confirmed this numerically, scanning over MSUGRA parameter space, and matching on to the observed Higgs mass, and found that even in the most optimistic case, dark radiation is produced in excess of observations. Thus effectively ruling out the MLVS.

Several groups have looked at dark radiation production in other string compactification models, and have found that the bounds can be evaded for certain sets of parameter choices. However, in all of these models it is true that dark radiation observations provide a very stringent constraint on the model, and in some cases can completely rule it out. In addition constraints on dark radiation are likely to become even tighter, thus studies of dark radiation production can be a powerful probe of high energy physics and string theory model building.

It is clear from this study that dark radiation production at reheating appears to be a generic prediction in string theory compactification models. If these ALPs are produced, then because they are very weakly interacting, in the present day they will form a homogeneous background of relativistic particles. The energy of such a cosmic ALP background (CAB) is dependent on the mass of the modulus and can plausibly be in the X-ray regime. In Chapter 5 we studied in detail a proposal that conversion of this CAB into X-ray photons in galaxy cluster magnetic fields could explain the cluster soft X-ray excess. The cluster soft X-ray excess is a significant excess seen with several satellites at energies around ~ 200 eV, in about a third of all galaxy clusters. Conventional astrophysical explanations such as warm gases and inverse Compton

scattering of CMB photons struggle to explain the excess across many clusters.

To study the soft X-ray excess we constructed simulations of the magnetic fields inside four galaxy clusters: Coma, A665, A2199, and A2255, and calculated ALP–photon conversion probabilities. Galaxy cluster magnetic fields typically have magnitudes $\mathcal{O}(1 - 10\ \mu\text{G})$ and coherence lengths $\mathcal{O}(10\ \text{kpc})$. The magnetic fields were modelled as random, tangled fields, with power-law power spectra. We showed that a CAB can explain the X-ray observations in each of the four clusters for a consistent set of CAB parameters, given current uncertainties on the magnetic fields in these objects. We find the best fit CAB parameters to be $M \sim 6 - 12 \times 10^{12}$ GeV and $E_{CAB} \sim 150$ eV.

We also studied the morphology of the excess in each galaxy cluster. We found in Coma that the morphology was fit better for a magnetic field with a flat power spectrum, or for models with very small coherence lengths. We found the morphology of the soft excess in A2199 and A2255 is reproduced by the simulations. We find the lack of an observed excess in A665 puts constraints on the CAB in slight disagreement with the best fit parameters from the Coma cluster analysis, although this disagreement is not strong considering uncertainties on the magnetic fields. Observations have shown a trend for increasing excess with radii, a feature which is reproduced by our simulations of CAB conversion. The conversion probabilities do not fall with radius as fast as the cluster X-ray emission, leading to larger excesses at large cluster radii.

The soft X-ray excess observations are now around a decade old. The current generation of X-ray satellites are not as well suited to studying the soft X-ray excess, due to their small fields of view. The new satellite Hitomi was hoped to be able to study the soft excess, since it had energy coverage down to 0.3 keV and superb energy resolution. Unfortunately it was only able to make one observation before breaking apart, during this observation of the Perseus cluster the satellite had no sensitivity below 2 keV, so there is now no hope for studying the soft excess with Hitomi. The future satellite eROSITA should however be able to study the soft excess.

Current studies of the CAB are hindered due to very limited knowledge of galaxy

cluster magnetic fields. The four clusters we have studied are the only four clusters for which there are both soft X-ray observations and detailed models of the magnetic field, and even still there are systematic uncertainties on the magnetic field models. In addition, we found that the predicted morphology of the excess is dependent on the magnetic field parameters, thus better knowledge of these is crucial for future studies of the soft excess. The SKA (square kilometre array) will increase knowledge of astrophysical magnetic fields in many clusters, allowing for a more systematic study of the CAB. The CAB can also be searched for in other astrophysical objects, such as in the Milky Way or other galaxies [216], and through its effects on the evolution of the universe [217]. Different models for cluster magnetic fields, for example using MHD simulations of cluster formation, could also be used to study the CAB and the cluster soft X-ray excess.

More indirectly, this explanation can be probed by the CABs contribution to dark radiation. Should the bounds on dark radiation become more severe, then the energy density of the CAB must be smaller. Thus in order to get the same soft excess signals, conversion probabilities of ALPs to photons would need to be increased. This can only be accomplished by a larger coupling between photons and ALPs (or smaller inverse coupling M). Bounds on M and on dark radiation are becoming ever more constraining and thus it seems plausible that in the long term the parameter region needed to explain the soft excess can be squeezed from both above and below.

In the last section we looked at constraining the ALP–photon coupling using galaxy clusters. This constraint is independent of any early universe physics as it does not require a population of ALPs to be produced. If ALPs exist, then galaxy cluster X-ray spectra will get distorted as the photons convert to ALPs in the cluster's magnetic field. This effect is energy-dependent and so will lead to energy-dependent distortions to the X-ray emission. We showed that, averaged over a large field of view, the effect of ALPs is to impart a step like feature into the X-ray spectrum, as the conversion probabilities are small at low energies, and larger at high X-ray energies. We also showed that along single photon sight-lines, or for small fields of view, ALP conversion

leads to large, quasi-oscillatory spectral distortions. There would also be an additional correlated polarisation signal caused by photon–ALP conversion.

The Coma cluster has been studied with a large field of view with several satellites, including XMM-Newton. Since these satellites have residuals at the 10% level, and no deviations from a thermal spectrum have been reported we showed that this leads to an estimated sensitivity to the ALP–photon coupling of up to $M \sim 7 \times 10^{11}$ GeV. The conversion probabilities were calculated using the same model for the magnetic field as in the soft X-ray excess study. A more precise bound would require a full analysis, fitting the data with the thermal bremsstrahlung model and the photon–ALP conversion probabilities. It is for this reason that we do not claim from this analysis to have obtained a bound better than the bound from SN1987a, we merely showed that this method is capable to doing so with a more thorough analysis.

We did however show a better bound can be found by instead looking for spectral distortions from small fields of view studies, or by looking at the spectra of point sources either in or behind a galaxy cluster. In both of these cases the resulting large oscillatory distortions would give a striking signal for the presence of ALPs. We think this is a much more promising place to search for ALPs, and can provide more precise bounds. We have carried out such an analysis [6] for the point source NGC1275 in the centre of the Perseus galaxy cluster (although the Perseus magnetic field is poorly constrained compared to Coma), finding hints for the existence of an ALP. The study of NGC1275 was also able to push the bounds on the ALP–photon coupling past the limits from the lack of a SN1987a gamma ray signal to around $2 - 5 \times 10^{11}$ GeV. Again, increased knowledge of galaxy cluster magnetic fields will allow for even more precise and stringent constraints.

Let us finish by summarising the state of axion-like particle physics in light of this thesis. We have shown that the galaxy cluster soft X-ray excess can be explained if there is a primordially-generated background of relativistic ALPs. We showed that the explanation holds for ALPs that couple to photons with an inverse coupling of $M \sim 7{\times}10^{12}$ GeV. This coupling holds for a CAB energy density equal to $\Delta N_{\rm eff} = 0.5$,

and for a CAB spectrum with an average energy around 150 eV. The best fit M in this scenario is far beyond the reach of current and planned laboratory experiments, and current astrophysical bounds.

We have shown that looking for distortions to the thermal spectra of galaxy clusters or the spectra of point sources either located in or behind a cluster should lead to much improved astrophysical bounds, although even still they are unlikely to get near the best fit CAB parameters. In a subsequent study we have found indications for distortions in the spectra of a point source in the centre of the Perseus galaxy cluster, this hints at an inverse coupling of photons to ALPs around the range $5-10\times10^{11}$GeV [6], which is somewhat lower than the best fit CAB value, thus it appears at first sight that these results are in some tension with each other.

Let us however discuss some ways that the two could co-exist. Firstly its important to remember that string theory generically predicts several ALPs, therefore it just may happen that the two effects are created by two different ALPs with differing couplings to electromagentism. The ALP responsible for the cluster distortions need not be produced in large quantities at reheating (e.g. by a suppressed branching ratio of the volume modulus to these ALPs), and thus the effects are completely decoupled, in this case there is little hope of testing the CAB ALP directly.

On the other hand the best fit value for the coupling from the CAB analysis is dependent on the unknown quantity $\Delta N_{\rm eff}$, lowering this value results in a lower CAB density and thus lower best fit M. Future experiments could reveal this number to be much smaller, bringing the best fit CAB inverse coupling in line with the coupling from the cluster distortions. In this case there is hope of potentially detecting these ALPs at future laboratory experiments.

In addition, the ALP responsible for the soft excess need not comprise the full CAB and thus may have an effective $\Delta N_{\rm eff}$ smaller than any observed value (if indeed there is an observation). Indeed we know this would be the case in the LVS model since the volume ALP couples too weakly to SM fields to give a CAB, thus a second ALP population would also have to be produced with a small energy density and have

a much stronger coupling to electromagetism. In this case it is perfectly plausible that the two effects can be driven by the same ALP. The main difficulty in this scenario is the one we came across in Chapter 4, that generically even with one ALP dark radiation is over-produced, and adding more ALPs may just make this situation worse.

Finally we note that sliding the CAB energy changes the best fit value of M, and could again bring it in agreement with the best fit value from the cluster distortions. However the energy cannot be changed too much otherwise excesses would be observed in other energy channels, for instance in the UV or the keV X-ray regime, and thus this does not seem a particularly appealing scenario.

Overall, studies of ALPs using distortions to either galaxy cluster X-ray spectra, or to point source spectra either in or behind a galaxy cluster are the most promising place to look for low mass ALPs. In addition any potential detections should be probeable in a laboratory setting (although maybe only in the long term). Future measurements of $\Delta N_{\rm eff}$ will be important in determining whether the CAB ALPs can be connected to these cluster distortions in some way. In the mean time only indirect studies using magnetic field models can be used to study these ALPs.

References

[1] S. Angus, J. P. Conlon, U. Haisch, and A. J. Powell, JHEP **12**, 061 (2013), arXiv:1305.4128 [hep-ph]

[2] S. Angus, J. P. Conlon, M. C. D. Marsh, A. J. Powell, and L. T. Witkowski, JCAP **1409**, 026 (2014), arXiv:1312.3947 [astro-ph.HE]

[3] A. J. Powell, JCAP **1509**, 017 (2015), arXiv:1411.4172 [astro-ph.CO]

[4] J. P. Conlon, M. C. D. Marsh, and A. J. Powell(2015), arXiv:1509.06748 [hep-ph]

[5] J. P. Conlon and A. J. Powell, JCAP **1501**, 019 (2015), arXiv:1406.5518 [hep-ph]

[6] M. Berg, J. P. Conlon, F. Day, N. Jennings, S. Krippendorf, A. J. Powell, and M. Rummel(2016), arXiv:1605.01043 [astro-ph.HE]

[7] G. Grilli di Cortona, E. Hardy, and A. J. Powell, JHEP **08**, 014 (2016), arXiv:1606.07090 [hep-ph]

[8] P. A. R. Ade *et al.* (Planck)(2015), arXiv:1502.01589 [astro-ph.CO]

[9] D. Baumann and L. McAllister, *Inflation and String Theory* (Cambridge University Press, 2015)

[10] S. Riemer-Sørensen, D. Parkinson, and T. M. Davis, Publications of the Astronomical Society of Australia **30**, e029 (Mar. 2013), arXiv:1301.7102 [astro-ph.CO]

[11] S. Sarkar, Rept. Prog. Phys. **59**, 1493 (1996), arXiv:hep-ph/9602260 [hep-ph]

[12] P. Candelas and X. de la Ossa, Nucl. Phys. **B355**, 455 (1991)

[13] S. Dimopoulos and G. F. Giudice, Phys. Lett. **B379**, 105 (1996), arXiv:hep-ph/9602350 [hep-ph]

[14] C. D. Hoyle, D. J. Kapner, B. R. Heckel, E. G. Adelberger, J. H. Gundlach, U. Schmidt, and H. E. Swanson, Phys. Rev. **D70**, 042004 (2004), arXiv:hep-ph/0405262 [hep-ph]

[15] M. Grana, Phys. Rept. **423**, 91 (2006), arXiv:hep-th/0509003 [hep-th]

[16] J. P. Conlon, Fortsch. Phys. **55**, 287 (2007), arXiv:hep-th/0611039 [hep-th]

[17] S. B. Giddings, S. Kachru, and J. Polchinski, Phys. Rev. **D66**, 106006 (2002), arXiv:hep-th/0105097 [hep-th]

[18] G. D. Coughlan, W. Fischler, E. W. Kolb, S. Raby, and G. G. Ross, Phys. Lett. **B131**, 59 (1983)

[19] M. S. Turner, Phys. Rev. **D28**, 1243 (1983)

[20] S. Davidson and S. Sarkar, JHEP **11**, 012 (2000), arXiv:hep-ph/0009078

[21] T. Banks, D. B. Kaplan, and A. E. Nelson, Phys. Rev. **D49**, 779 (1994), arXiv:hep-ph/9308292

[22] B. de Carlos, J. A. Casas, F. Quevedo, and E. Roulet, Phys. Lett. **B318**, 447 (1993), arXiv:hep-ph/9308325

[23] M. Cicoli, M. Goodsell, and A. Ringwald, JHEP **10**, 146 (2012), arXiv:1206.0819 [hep-th]

[24] C. A. Baker *et al.*, Phys. Rev. Lett. **97**, 131801 (2006), arXiv:hep-ex/0602020

[25] R. D. Peccei, *Axions: Theory, cosmology, and experimental searches*, Lect. Notes Phys. **741**, 3 (2008)

[26] R. D. Peccei and H. R. Quinn, Phys. Rev. Lett. **38**, 1440 (1977)

[27] R. D. Peccei and H. R. Quinn, Phys. Rev. **D16**, 1791 (1977)

[28] S. Weinberg, Phys. Rev. Lett. **40**, 223 (1978)

[29] F. Wilczek, Phys. Rev. Lett. **40**, 279 (1978)

[30] K. A. Olive *et al.* (Particle Data Group), Chin. Phys. **C38**, 090001 (2014)

[31] L. F. Abbott and P. Sikivie, Phys. Lett. **B120**, 133 (1983)

[32] P. Svrcek and E. Witten, JHEP **06**, 051 (2006), arXiv:hep-th/0605206 [hep-th]

[33] A. Arvanitaki, S. Dimopoulos, S. Dubovsky, N. Kaloper, and J. March-Russell, Phys. Rev. **D81**, 123530 (2010), arXiv:0905.4720 [hep-th]

[34] B. S. Acharya, K. Bobkov, and P. Kumar, JHEP **11**, 105 (2010), arXiv:1004.5138 [hep-th]

[35] P. Sikivie, Phys. Rev. Lett. **51**, 1415 (1983), [Erratum: Phys. Rev. Lett.52,695(1984)]

[36] P. Sikivie, Phys.Rev. **D32**, 2988 (1985)

[37] G. Raffelt and L. Stodolsky, Phys. Rev. **D37**, 1237 (1988)

[38] G. G. Raffelt, Phys. Rev. **D33**, 897 (1986)

[39] G. G. Raffelt and D. S. P. Dearborn, Phys. Rev. **D37**, 549 (1988)

[40] A. Ayala, I. Domínguez, M. Giannotti, A. Mirizzi, and O. Straniero, Phys. Rev. Lett. **113**, 191302 (2014), arXiv:1406.6053 [astro-ph.SR]

[41] M. Giannotti, I. Irastorza, J. Redondo, and A. Ringwald(2015), arXiv:1512.08108 [astro-ph.HE]

[42] K. Hirata *et al.* (Kamiokande-II), Phys. Rev. Lett. **58**, 1490 (1987), [,727(1987)]

[43] R. M. Bionta *et al.*, Phys. Rev. Lett. **58**, 1494 (1987)

[44] E. L. Chupp, W. T. Vestrand, and C. Reppin, Phys. Rev. Lett. **62**, 505 (1989)

[45] J. W. Brockway, E. D. Carlson, and G. G. Raffelt, Phys.Lett. **B383**, 439 (1996), arXiv:astro-ph/9605197

[46] J. Grifols, E. Masso, and R. Toldra, Phys.Rev.Lett. **77**, 2372 (1996), arXiv:astro-ph/9606028

[47] A. Payez, C. Evoli, T. Fischer, M. Giannotti, A. Mirizzi, and A. Ringwald, JCAP **1502**, 006 (2015), arXiv:1410.3747 [astro-ph.HE]

[48] J. R. Ellis and K. A. Olive, Phys. Lett. **B193**, 525 (1987)

[49] G. Raffelt and D. Seckel, Phys. Rev. Lett. **60**, 1793 (1988)

[50] M. S. Turner, Phys. Rev. Lett. **60**, 1797 (1988)

[51] R. Mayle, J. R. Wilson, J. R. Ellis, K. A. Olive, D. N. Schramm, and G. Steigman, Phys. Lett. **B203**, 188 (1988)

[52] R. Mayle, J. R. Wilson, J. R. Ellis, K. A. Olive, D. N. Schramm, and G. Steigman, Phys. Lett. **B219**, 515 (1989)

[53] H.-T. Janka, W. Keil, G. Raffelt, and D. Seckel, Phys. Rev. Lett. **76**, 2621 (1996), arXiv:astro-ph/9507023

[54] W. Keil, H.-T. Janka, D. N. Schramm, G. Sigl, M. S. Turner, and J. R. Ellis, Phys. Rev. **D56**, 2419 (1997), arXiv:astro-ph/9612222

[55] A. Abramowski *et al.* (HESS), Phys. Rev. **D88**, 102003 (2013), arXiv:1311.3148 [astro-ph.HE]

[56] D. Horns and M. Meyer, JCAP **1202**, 033 (2012), arXiv:1201.4711 [astro-ph.CO]

[57] G. I. Rubtsov and S. V. Troitsky, JETP Lett. **100**, 397 (2014), arXiv:1406.0239 [astro-ph.HE]

[58] G. Galanti, M. Roncadelli, A. De Angelis, and G. F. Bignami(2015), arXiv:1503.04436 [astro-ph.HE]

[59] J. Biteau and D. A. Williams, Astrophys. J. **812**, 60 (2015), arXiv:1502.04166 [astro-ph.CO]

[60] A. De Angelis, M. Roncadelli, and O. Mansutti, Phys.Rev. **D76**, 121301 (2007), arXiv:0707.4312 [astro-ph]

[61] A. De Angelis, G. Galanti, and M. Roncadelli, Phys. Rev. **D84**, 105030 (2011), [Erratum: Phys. Rev.D87,no.10,109903(2013)], arXiv:1106.1132 [astro-ph.HE]

[62] M. Meyer, D. Horns, and M. Raue, Phys.Rev. **D87**, 035027 (2013), arXiv:1302.1208 [astro-ph.HE]

[63] M. Simet, D. Hooper, and P. D. Serpico, Phys.Rev. **D77**, 063001 (2008), arXiv:0712.2825 [astro-ph]

[64] D. Horns, L. Maccione, M. Meyer, A. Mirizzi, D. Montanino, and M. Roncadelli, Phys. Rev. **D86**, 075024 (2012), arXiv:1207.0776 [astro-ph.HE]

[65] M. Sanchez-Conde, D. Paneque, E. Bloom, F. Prada, and A. Dominguez, Phys.Rev. **D79**, 123511 (2009), arXiv:0905.3270 [astro-ph.CO]

[66] S. Troitsky, Phys. Rev. **D93**, 045014 (2016), arXiv:1507.08640 [astro-ph.HE]

[67] R. Ballou *et al.* (OSQAR), Phys. Rev. **D92**, 092002 (2015), arXiv:1506.08082 [hep-ex]

[68] K. Ehret *et al.*, Phys. Lett. **B689**, 149 (2010), arXiv:1004.1313 [hep-ex]

[69] R. Bähre *et al.*, JINST **8**, T09001 (2013), arXiv:1302.5647 [physics.ins-det]

[70] S. Andriamonje *et al.* (CAST), JCAP **0704**, 010 (2007), arXiv:hep-ex/0702006

[71] S. Aune *et al.* (CAST), Phys.Rev.Lett. **107**, 261302 (2011), arXiv:1106.3919 [hep-ex]

[72] E. Armengaud, F. Avignone, M. Betz, P. Brax, P. Brun, *et al.*, JINST **9**, T05002 (2014), arXiv:1401.3233 [physics.ins-det]

[73] S. J. Asztalos *et al.* (ADMX), Phys. Rev. **D69**, 011101 (2004), arXiv:astro-ph/0310042

[74] S. J. Asztalos *et al.* (ADMX), Phys. Rev. Lett. **104**, 041301 (2010), arXiv:0910.5914 [astro-ph.CO]

[75] D. Budker, P. W. Graham, M. Ledbetter, S. Rajendran, and A. Sushkov, Phys. Rev. **X4**, 021030 (2014), arXiv:1306.6089 [hep-ph]

[76] A. Arvanitaki and A. A. Geraci, Phys. Rev. Lett. **113**, 161801 (2014), arXiv:1403.1290 [hep-ph]

[77] M. Cicoli, J. P. Conlon, and F. Quevedo, Phys. Rev. **D87**, 043520 (2013), arXiv:1208.3562 [hep-ph]

[78] T. Higaki and F. Takahashi, JHEP **11**, 125 (2012), arXiv:1208.3563 [hep-ph]

[79] A. G. Riess *et al.*, Astrophys. J. **730**, 119 (2011), [Erratum: Astrophys. J.732,129(2011)], arXiv:1103.2976 [astro-ph.CO]

[80] G. Hinshaw *et al.* (WMAP), Astrophys. J. Suppl. **208**, 19 (2013), arXiv:1212.5226 [astro-ph.CO]

[81] J. L. Sievers *et al.* (Atacama Cosmology Telescope), JCAP **1310**, 060 (2013), arXiv:1301.0824 [astro-ph.CO]

[82] Z. Hou *et al.*, Astrophys. J **782**, 74 (Feb. 2014), arXiv:1212.6267 [astro-ph.CO]

[83] R. Adam *et al.* (Planck)(2015), arXiv:1502.01582 [astro-ph.CO]

[84] A. G. Riess *et al.*(2016), arXiv:1604.01424 [astro-ph.CO]

[85] R. Cooke, M. Pettini, R. A. Jorgenson, M. T. Murphy, and C. C. Steidel, Astrophys. J. **781**, 31 (2014), arXiv:1308.3240 [astro-ph.CO]

[86] E. Aver, K. A. Olive, R. Porter, and E. D. Skillman, JCAP **1311**, 017 (2013), arXiv:1309.0047 [astro-ph.CO]

[87] V. Balasubramanian, P. Berglund, J. P. Conlon, and F. Quevedo, JHEP **03**, 007 (2005), arXiv:hep-th/0502058

[88] J. P. Conlon, F. Quevedo, and K. Suruliz, JHEP **08**, 007 (2005), arXiv:hep-th/0505076

[89] R. Blumenhagen, J. P. Conlon, S. Krippendorf, S. Moster, and F. Quevedo, JHEP **09**, 007 (2009), arXiv:0906.3297 [hep-th]

[90] A. Hebecker, A. K. Knochel, and T. Weigand, JHEP **06**, 093 (2012), arXiv:1204.2551 [hep-th]

[91] A. Hebecker, A. K. Knochel, and T. Weigand, Nucl. Phys. **B874**, 1 (2013), arXiv:1304.2767 [hep-th]

[92] T. Higaki, K. Nakayama, and F. Takahashi, JHEP **1307**, 005 (2013), arXiv:1304.7987 [hep-ph]

[93] G. F. Giudice and A. Masiero, Phys. Lett. **B206**, 480 (1988)

[94] M. T. Grisaru, W. Siegel, and M. Rocek, Nucl. Phys. **B159**, 429 (1979)

[95] S. J. Gates, M. T. Grisaru, M. Rocek, and W. Siegel, Front. Phys. **58**, 1 (1983), arXiv:hep-th/0108200

[96] M. E. Machacek and M. T. Vaughn, Nucl. Phys. **B222**, 83 (1983)

[97] M. E. Machacek and M. T. Vaughn, Nucl. Phys. **B236**, 221 (1984)

[98] D. M. Capper and D. R. T. Jones, Nucl. Phys. **B252**, 718 (1985)

[99] S. P. Martin, Adv. Ser. Direct. High Energy Phys. **18**, 1 (1997), arXiv:hep-ph/9709356

[100] B. C. Allanach, Comput. Phys. Commun. **143**, 305 (2002), arXiv:hep-ph/0104145

[101] D. M. Pierce, J. A. Bagger, K. T. Matchev, and R.-j. Zhang, Nucl. Phys. **B491**, 3 (1997), arXiv:hep-ph/9606211

[102] T. Aaltonen *et al.* (CDF, D0), Phys. Rev. **D86**, 092003 (2012), arXiv:1207.1069 [hep-ex]

[103] M. Beneke, P. Falgari, S. Klein, and C. Schwinn, Nucl. Phys. **B855**, 695 (2012), arXiv:1109.1536 [hep-ph]

[104] J. Beringer *et al.* (Particle Data Group), Phys. Rev. **D86**, 010001 (2012)

[105] G. Aad *et al.* (ATLAS), Phys. Lett. **B716**, 1 (2012), arXiv:1207.7214 [hep-ex]

[106] S. Chatrchyan *et al.* (CMS), Phys. Lett. **B716**, 30 (2012), arXiv:1207.7235 [hep-ex]

[107] G. Aad *et al.* (ATLAS, CMS), Phys. Rev. Lett. **114**, 191803 (2015), arXiv:1503.07589 [hep-ex]

[108] B. C. Allanach, A. Djouadi, J. L. Kneur, W. Porod, and P. Slavich, JHEP **09**, 044 (2004), arXiv:hep-ph/0406166

[109] S. Angus, JHEP **10**, 184 (2014), arXiv:1403.6473 [hep-ph]

[110] A. Hebecker, P. Mangat, F. Rompineve, and L. T. Witkowski, JHEP **09**, 140 (2014), arXiv:1403.6810 [hep-ph]

[111] N. Arkani-Hamed and S. Dimopoulos, JHEP **06**, 073 (2005), arXiv:hep-th/0405159

[112] G. F. Giudice and A. Romanino, Nucl. Phys. **B699**, 65 (2004), [Erratum: Nucl. Phys.B706,65(2005)], arXiv:hep-ph/0406088

[113] M. Cicoli and F. Muia, JHEP **12**, 152 (2015), arXiv:1511.05447 [hep-th]

[114] B. S. Acharya and C. Pongkitivanichkul(2015), arXiv:1512.07907 [hep-ph]

[115] J. P. Conlon and M. C. D. Marsh, JHEP **1310**, 214 (2013), arXiv:1304.1804 [hep-ph]

[116] J. P. Conlon and M. C. D. Marsh, Phys. Rev. Lett. **111**, 151301 (2013), arXiv:1305.3603 [astro-ph.CO]

[117] F. Zwicky, Astrophys. J. **86**, 217 (Oct. 1937)

[118] R. Lieu, J. P. D. Mittaz, S. Bowyer, F. J. Lockman, C.-Y. Hwang, and J. H. M. M. Schmitt, Astrophys. J. **458**, L5 (Feb. 1996)

[119] R. Lieu, J. P. D. Mittaz, S. Bowyer, J. O. Breen, F. J. Lockman, E. M. Murphy, and C.-Y. Hwang, Science **274**, 1335 (Nov. 1996)

[120] S. Bowyer, M. Lampton, and R. Lieu, Science **274**, 1338 (Nov. 1996)

[121] M. Bonamente, R. Lieu, M. K. Joy, and J. H. Nevalainen, Astrophys. J. **576**, 688 (Sep. 2002), arXiv:astro-ph/0205473

[122] F. Durret, J. S. Kaastra, J. Nevalainen, T. Ohashi, and N. Werner, Space Sci. Rev. **134**, 51 (Feb. 2008), arXiv:0801.0977

[123] M. Murgia, F. Govoni, L. Feretti, G. Giovannini, D. Dallacasa, R. Fanti, G. B. Taylor, and K. Dolag, Astron. Astrophys. **424**, 429 (Sep. 2004), astro-ph/0406225

[124] T. Higaki, K. Kamada, and F. Takahashi, JHEP **1209**, 043 (2012), arXiv:1207.2771 [hep-ph]

[125] T. Higaki, K. Nakayama, and F. Takahashi, JCAP **1309**, 030 (2013), arXiv:1306.6518 [hep-ph]

[126] M. Fairbairn, Phys.Rev. **D89**, 064020 (2014), arXiv:1310.4464 [astro-ph.CO]

[127] R. F. Malina and S. Bowyer, NASA STI/Recon Technical Report A **93**, 25725 (Jul. 1991)

[128] J. P. D. Mittaz, R. Lieu, and F. J. Lockman, Astrophys. J. **498**, L17 (May 1998)

[129] R. Lieu, M. Bonamente, and J. P. D. Mittaz, Astrophys. J. **517**, L91 (Jun. 1999), astro-ph/9902300

[130] S. Bowyer, T. W. Berghöfer, and E. J. Korpela, Astrophys. J. **526**, 592 (Dec. 1999), astro-ph/9911001

[131] T. W. Berghöfer, S. Bowyer, and E. Korpela, Astrophys. J. **535**, 615 (Jun. 2000)

[132] T. W. Berghöfer, S. Bowyer, and E. Korpela, Astrophys. J. **545**, 695 (Dec. 2000), astro-ph/0011274

[133] S. Bowyer, E. Korpela, and T. Berghöfer, Astrophys. J. **548**, L135 (Feb. 2001), astro-ph/0101300

[134] R. Lieu, M. Bonamente, J. P. D. Mittaz, F. Durret, S. Dos Santos, and J. S. Kaastra, Astrophys. J. **527**, L77 (Dec. 1999), astro-ph/9910298

[135] M. Bonamente, R. Lieu, and J. P. D. Mittaz, Astrophys. J. **547**, L7 (Jan. 2001), astro-ph/0011186

[136] F. Durret, E. Slezak, R. Lieu, S. Dos Santos, and M. Bonamente, Astron. Astrophys. **390**, 397 (Aug. 2002), astro-ph/0204345

[137] U. Briel et al., *ROSAT User's Handbook*

[138] M. Bonamente, R. Lieu, J. Nevalainen, and J. S. Kaastra, Astrophys. J. **552**, L7 (May 2001), astro-ph/0103331

[139] M. Bonamente, R. Lieu, and J. Kaastra, Astron. Astrophys. **443**, 29 (Nov. 2005), astro-ph/0512591

[140] M. Bonamente, R. Lieu, and J. P. D. Mittaz, Astrophys. J. **561**, L63 (Nov. 2001), astro-ph/0107379

[141] M. Bonamente and J. Nevalainen, Astrophys. J. **738**, 149 (Sep. 2011), arXiv:1107.0932 [astro-ph.HE]

[142] J. Nevalainen, R. Lieu, M. Bonamente, and D. Lumb, Astrophys. J. **584**, 716 (Feb. 2003), astro-ph/0210610

[143] J. S. Kaastra, R. Lieu, T. Tamura, F. B. S. Paerels, and J. W. den Herder, Astron. Astrophys. **397**, 445 (Jan. 2003), astro-ph/0210684

[144] J. S. Kaastra, R. Lieu, T. Tamura, F. B. S. Paerels, and J. W. A. den Herder, Astrophys. Space Sci. Libr. **309**, 37 (2004), arXiv:astro-ph/0305424

[145] J. S. Kaastra, Journal of Korean Astronomical Society **37**, 375 (Dec. 2004)

[146] Y. Takei *et al.*, Publ. Astron. Soc. Jpn **59**, 339 (Jan. 2007), astro-ph/0608072

[147] N. Werner, J. S. Kaastra, Y. Takei, R. Lieu, J. Vink, and T. Tamura, Astron. Astrophys. **468**, 849 (Jun. 2007), arXiv:0704.0475

[148] T. Lehto, J. Nevalainen, M. Bonamente, N. Ota, and J. Kaastra, Astron. Astrophys. **524**, A70 (Dec. 2010), arXiv:1010.3840 [astro-ph.CO]

[149] G. E. Bulbul, R. K. Smith, A. Foster, J. Cottam, M. Loewenstein, R. Mushotzky, and R. Shafer, Astrophys. J. **747**, 32 (Mar. 2012), arXiv:1110.4422 [astro-ph.CO]

[150] Y. Takei, J. P. Henry, A. Finoguenov, K. Mitsuda, T. Tamura, R. Fujimoto, and U. G. Briel, Astrophys. J. **655**, 831 (Feb. 2007), astro-ph/0610424

[151] A. Finoguenov, U. G. Briel, and J. P. Henry, Astron. Astrophys. **410**, 777 (Nov. 2003), astro-ph/0309019

[152] R. Lieu and J. P. D. Mittaz, in *The Identification of Dark Matter*, edited by N. J. C. Spooner and V. Kudryavtsev (2005) pp. 18–24, astro-ph/0501007

[153] J. N. Bregman and E. J. Lloyd-Davies, Astrophys. J. **644**, 167 (Jun. 2006), astro-ph/0602527

[154] J. Nevalainen, M. Bonamente, and J. Kaastra, Astrophys. J. **656**, 733 (Feb. 2007), astro-ph/0610461

[155] M. Bonamente, J. Nevalainen, and R. Lieu, Astrophys. J. **668**, 796 (Oct. 2007), arXiv:0707.0992 [astro-ph]

[156] M. Bonamente, M. K. Joy, and R. Lieu, Astrophys. J. **585**, 722 (Mar. 2003), astro-ph/0211439

[157] M. Bonamente, R. Lieu, and E. Bulbul, Astrophys. J. **696**, 1886 (May 2009), arXiv:0903.3067 [astro-ph.CO]

[158] J. S. Kaastra, R. Lieu, J. P. D. Mittaz, J. A. M. Bleeker, R. Mewe, S. Colafrancesco, and F. J. Lockman, Astrophys. J. Lett. **519**, L119 (Jul. 1999), astro-ph/9905209

[159] R. Lieu, M. Bonamente, and J. P. D. Mittaz, Astron. Astrophys. **364**, 497 (Dec. 2000), astro-ph/0010610

[160] T. W. Berghöfer and S. Bowyer, Astrophys. J. Lett. **565**, L17 (Jan. 2002), astro-ph/0112337

[161] J. S. Kaastra, R. Lieu, J. A. M. Bleeker, R. Mewe, and S. Colafrancesco, Astrophys. J. Lett. **574**, L1 (Jul. 2002), astro-ph/0206235

[162] J. Mittaz, R. Lieu, R. Cen, and M. Bonamente, Astrophys. J. **617**, 860 (Dec. 2004), arXiv:0409661

[163] L.-M. Cheng *et al.*, Astron. Astrophys. **431**, 405 (Feb. 2005), astro-ph/0409707

[164] R. Lieu and M. Bonamente, Astrophys. J. **698**, 1301 (Jun. 2009), arXiv:0903.3066 [astro-ph.CO]

[165] R. Cen and J. P. Ostriker, Astrophys. J **514**, 1 (Mar. 1999), astro-ph/9806281

[166] A. C. Fabian, Science **271**, 1244 (1996)

[167] W. V. D. Dixon, M. Hurwitz, and H. C. Ferguson, Astrophys. J. **469**, L77 (Oct. 1996), astro-ph/9607126

[168] W. V. D. Dixon, S. Sallmen, M. Hurwitz, and R. Lieu, Astrophys. J. **550**, L25 (Mar. 2001), astro-ph/0102093

[169] C.-Y. Hwang, Science **278**, 1917 (Dec. 1997)

[170] T. A. Ensslin and P. L. Biermann, Astron. Astrophys. **330**, 90 (Feb. 1998), astro-ph/9709232

[171] C. L. Sarazin and R. Lieu, Astrophys. J. **494**, L177 (Feb. 1998), astro-ph/9712049

[172] S. Bowyer and T. W. Berghöfer, Astrophys. J. **506**, 502 (Oct. 1998), astro-ph/9804310

[173] G. Brunetti, L. Feretti, G. Giovannini, and G. Setti, in *Diffuse Thermal and Relativistic Plasma in Galaxy Clusters*, edited by H. Boehringer, L. Feretti, and P. Schuecker (1999) p. 263, astro-ph/9907125

[174] C. L. Sarazin(1999), arXiv:astro-ph/9906272

[175] G. Brunetti, G. Setti, L. Feretti, and G. Giovannini, Mon.Not.Roy.Astron.Soc. **320**, 365 (2001), arXiv:astro-ph/0008518

[176] G. Brunetti, G. Setti, L. Feretti, and G. Giovannini, New Astron. **6**, 1 (Feb. 2001), astro-ph/0011301

[177] V. Petrosian, Astrophys. J. **557**, 560 (Aug. 2001), astro-ph/0101145

[178] Y. Fujita and C. L. Sarazin, Astrophys. J. **563**, 660 (Dec. 2001), astro-ph/0108369

[179] A. M. Atoyan and H. J. Völk, Astrophys. J. **535**, 45 (May 2000), astro-ph/9912557

[180] M. Y. Tsay, C.-Y. Hwang, and S. Bowyer, Astrophys. J. **566**, 794 (Feb. 2002)

[181] M. Ackermann *et al.*, Astrophys. J. **717**, L71 (Jul. 2010), arXiv:1006.0748 [astro-ph.HE]

[182] S. Ando and D. Nagai, JCAP **1207**, 017 (2012), arXiv:1201.0753 [astro-ph.HE]

[183] J. Han *et al.*, Mon.Not.Roy.Astron.Soc. **427**, 1651 (2012), arXiv:1207.6749 [astro-ph.CO]

[184] M. Ackermann *et al.* (Fermi-LAT), Astrophys. J. **787**, 18 (2014), arXiv:1308.5654 [astro-ph.HE]

[185] B. Huber *et al.*, Astron. Astrophys. **560**, A64 (Dec. 2013), arXiv:1308.6278 [astro-ph.HE]

[186] F. Zandanel and S. Ando, Mon.Not.Rot.Astron.Soc. **440**, 663 (May 2014), arXiv:1312.1493 [astro-ph.HE]

[187] A. Bonafede *et al.*, Astron. Astrophys. **513**, A30 (Apr. 2010), arXiv:1002.0594 [astro-ph.CO]

[188] F. Govoni and L. Feretti, Int.J.Mod.Phys. **D13**, 1549 (2004), arXiv:astro-ph/0410182 [astro-ph]

[189] L. Feretti, G. Giovannini, F. Govoni, and M. Murgia, Astron. Astrophys. Rev. **20**, 54 (May 2012), arXiv:1205.1919 [astro-ph.CO]

[190] A. Brandenburg and K. Subramanian, Phys. Rep. **417**, 1 (Oct. 2005), astro-ph/0405052

[191] A. A. Schekochihin and S. C. Cowley, Physics of Plasmas **13**, 056501 (May 2006), astro-ph/0601246

[192] R. Durrer and A. Neronov, Astron. Astrophys. Rev. **21**, 62 (Jun. 2013), arXiv:1303.7121 [astro-ph.CO]

[193] F. Tavecchio *et al.*, Mon.Not.Rot.Astron.Soc. **406**, L70 (Jul. 2010), arXiv:1004.1329 [astro-ph.CO]

[194] P. Ade *et al.* (Planck)(2015), arXiv:1502.01594 [astro-ph.CO]

[195] J. Donnert, K. Dolag, H. Lesch, and E. Muller, Mon.Not.Roy.Astron.Soc. **392**, 1008 (2009), arXiv:0808.0919 [astro-ph]

[196] A. Brandenburg, K. Enqvist, and P. Olesen, Phys.Rev. **D54**, 1291 (1996), arXiv:astro-ph/9602031

[197] K. Jedamzik and G. Sigl, Phys. Rev., 103005(May 2011), arXiv:1012.4794 [astro-ph.CO]

[198] T. Kahniashvili, A. G. Tevzadze, A. Brandenburg, and A. Neronov, Phys.Rev. **D87**, 083007 (2013), arXiv:1212.0596 [astro-ph.CO]

[199] K. Dolag, M. Kachelriess, S. Ostapchenko, and R. Tomàs, Astrophys. J. Lett. **727**, L4 (Jan. 2011), arXiv:1009.1782 [astro-ph.HE]

[200] K. Dolag, M. Bartelmann, and H. Lesch, Astron. Astrophys. **348**, 351 (Aug. 1999), astro-ph/0202272

[201] K. Dolag, M. Bartelmann, and H. Lesch, Astron. Astrophys. **387**, 383 (May 2002)

[202] M. Thierbach, U. Klein, and R. Wielebinski, Astron. Astrophys. **397**, 53 (Jan. 2003), astro-ph/0210147

[203] C. M. H. Chen, D. E. Harris, F. A. Harrison, and P. H. Mao, Mon.Not.Roy.Astron.Soc. **383**, 1259 (Jan. 2008), arXiv:0801.0288

[204] L. Rudnick and K. M. Blundell, Astrophys. J. **588**, 143 (May 2003), astro-ph/0301260

[205] F. Govoni, M. Murgia, L. Feretti, G. Giovannini, K. Dolag, and G. B. Taylor, Astron. Astrophys. **460**, 425 (Dec. 2006), astro-ph/0608433

[206] V. Vacca, M. Murgia, F. Govoni, L. Feretti, G. Giovannini, *et al.*(2010), arXiv:1001.1058 [astro-ph.CO]

[207] V. Vacca *et al.*, Astron. Astrophys. **540**, A38 (Apr. 2012), arXiv:1201.4119 [astro-ph.CO]

[208] A. Bonafede, F. Vazza, M. Brüggen, M. Murgia, F. Govoni, L. Feretti, G. Giovannini, and G. Ogrean, Mon.Not.Roy.Astron.Soc. **433**, 3208 (Aug. 2013), arXiv:1305.7228 [astro-ph.CO]

[209] D. S. Hudson, R. Mittal, T. H. Reiprich, P. E. Nulsen, H. Andernach, *et al.*(2009), arXiv:0911.0409 [astro-ph.CO]

[210] U. G. Briel, J. P. Henry, and H. Boehringer, Astron. Astrophys. **259**, L31 (Jun. 1992)

[211] K.-T. Kim, P. P. Kronberg, P. E. Dewdney, and T. L. Landecker, Astrophys. J. **355**, 29 (May 1990)

[212] D. Kraljic, M. Rummel, and J. P. Conlon, JCAP **1501**, 011 (2015), arXiv:1406.5188 [hep-ph]

[213] M. Frigo and S. G. Johnson, Proceedings of IEEE **93**, 216–231 (2005)

[214] E. L. Wright, Publ. Astron. Soc. Pac. **118**, 1711 (Dec. 2006), astro-ph/0609593

[215] http://heasarc.gsfc.nasa.gov/docs/software/tools/pimms.html

[216] F. V. Day, Phys. Lett. **B753**, 600 (2016), arXiv:1506.05334 [hep-ph]

[217] C. Evoli, M. Leo, A. Mirizzi, and D. Montanino(2016), arXiv:1602.08433 [astro-ph.CO]

[218] M. Fairbairn, T. Rashba, and S. V. Troitsky, Phys. Rev. **D84**, 125019 (2011), arXiv:0901.4085 [astro-ph.HE]

[219] C. Burrage, A.-C. Davis, and D. J. Shaw, Phys. Rev. Lett. **102**, 201101 (2009), arXiv:0902.2320 [astro-ph.CO]

[220] M. Schlederer and G. Sigl, JCAP **1601**, 038 (2016), arXiv:1507.02855 [hep-ph]

[221] D. Wouters and P. Brun, Astrophys. J. **772**, 44 (2013), arXiv:1304.0989 [astro-ph.HE]

[222] M. Arnaud *et al.*, Astron.Astrophys. **365**, L67 (2001), arXiv:astro-ph/0011086

[223] F. Gastaldello *et al.*, Astrophys. J. **800**, 139 (2015), arXiv:1411.1573 [astro-ph.HE]

[224] E. Bulbul, M. Markevitch, A. Foster, R. K. Smith, M. Loewenstein, and S. W. Randall, Astrophys. J. **789**, 13 (2014), arXiv:1402.2301 [astro-ph.CO]

[225] A. Payez, J. R. Cudell, and D. Hutsemékers, JCAP **7**, 041 (Jul. 2012), arXiv:1204.6187 [astro-ph.CO]