

NEUTRINO OSCILLATIONS IN VERY DENSE MEDIA – PRODUCTION AND CHAOS

RASMUS SLOTH HANSEN

Ph.D. DISSERTATION



JULY 2015

SUPERVISOR: PROF. STEEN HANNESTAD

DEPARTMENT OF PHYSICS AND ASTRONOMY
AARHUS UNIVERSITY

NEUTRINO OSCILLATIONS
IN VERY DENSE MEDIA
– PRODUCTION AND CHAOS

A DISSERTATION
PRESENTED TO THE FACULTY OF SCIENCE AND TECHNOLOGY
OF AARHUS UNIVERSITY
IN PARTIAL FULFILMENT OF THE REQUIREMENTS FOR THE
PH.D. DEGREE

BY
RASMUS SLOTH HANSEN
JULY 31 2015

Figure 5.6 have been adapted with permission from a figure in *Physical Review D*.
Figure 5.7 have been adapted with permission from a figure in *Physical Review Letters*.
The specific source is cited in each case.

2015

Printed in Aarhus, Denmark.

Typeset in 11 pt. Latin Modern using L^AT_EX and the Memoir Class.

Preface

This thesis concludes my four years as a Ph.D. at Aarhus University where I have had the pleasure of studying the intricate physics of neutrino oscillations. My studies began with the introduction of a halfway finished project that Steen and Thomas had been working on for a while on chaotic neutrino oscillations. This forced me through a quick schooling both of solving numerical differential equations and of the mathematics of chaotic systems. The work concluded in a paper on the chaotic behaviour of active-sterile neutrino oscillations in the early Universe [1], but the idea of investigating the chaotic aspect of neutrino oscillations had caught on.

Instead of focusing on the early Universe, we did however turn our interest towards supernovae, where chaotic behaviour is omnipresent in recent simplified models of neutrino oscillations. Taking on one of the simplest models, it was possible to do a very extensive chaotic analysis [2], and we obtained interesting information about the dynamics of the system. In parallel, our interest in the early Universe had turned to methods for suppressing sterile neutrino production. Throughout the past two years, we have investigated models where the sterile neutrinos interact with each other. In our first approach we only considered a mechanism that could limit the production [3], but in a later work, we also included a rough analysis of implications for dark matter and Maria joined with a detailed analysis of the signal from the cosmological microwave background [4].

My research stay was in Sydney with Yvonne, where I looked for new challenging projects. The choice fell upon the messy dynamics of limiting sterile neutrino production by using a large asymmetry between particles and antiparticles, but in the end convergence problems pushed us in a different direction. The end result was a study of how detailed collisions affect neutrino oscillations [5], which hopefully brings us a small step closer to the solution of the original problem as well.

In total, the work for my thesis has resulted in the publications:

- [1] Steen Hannestad, **Rasmus Sloth Hansen**, and Thomas Tram. “Can active-sterile neutrino oscillations lead to chaotic behavior of the cosmological lepton asymmetry?” *JCAP* 1304 (2013), p. 032. DOI: [10.1088/1475-7516/2013/04/032](https://doi.org/10.1088/1475-7516/2013/04/032). arXiv: [1302.7279](https://arxiv.org/abs/1302.7279) [[astro-ph.CO](https://arxiv.org/archive/astro)].
- [2] **Rasmus Sloth Hansen** and Steen Hannestad. “Chaotic flavor evolution in an interacting neutrino gas”. *Phys.Rev.* D90.2 (2014), p. 025009. DOI: [10.1103/PhysRevD.90.025009](https://doi.org/10.1103/PhysRevD.90.025009). arXiv: [1404.3833](https://arxiv.org/abs/1404.3833) [[hep-ph](https://arxiv.org/archive/hep)].
- [3] Steen Hannestad, **Rasmus Sloth Hansen**, and Thomas Tram. “How Self-Interactions can Reconcile Sterile Neutrinos with Cosmology”. *Phys.Rev.Lett.* 112.3 (2014), p. 031802. DOI: [10.1103/PhysRevLett.112.031802](https://doi.org/10.1103/PhysRevLett.112.031802). arXiv: [1310.5926](https://arxiv.org/abs/1310.5926) [[astro-ph.CO](https://arxiv.org/archive/astro)].

- [4] Maria Archidiacono, Steen Hannestad, **Rasmus Sloth Hansen**, and Thomas Tram. “Cosmology with self-interacting sterile neutrinos and dark matter - A pseudoscalar model”. *Phys.Rev.* D91.6 (2015), p. 065021. DOI: [10.1103/PhysRevD.91.065021](https://doi.org/10.1103/PhysRevD.91.065021). arXiv: [1404.5915](https://arxiv.org/abs/1404.5915) [[astro-ph.CO](#)].
- [5] Steen Hannestad, **Rasmus Sloth Hansen**, Thomas Tram, and Yvonne Y. Y. Wong. “Active-sterile neutrino oscillations in the early Universe with full collision terms”. *JCAP* (2015), forthcoming. arXiv: [1506.05266](https://arxiv.org/abs/1506.05266) [[hep-ph](#)].

Except for these papers, I have also visited Munich, where I continued the work on supernova neutrinos together with Georg, Ignacio and Sovan. This time, the topic was not directly related to chaos, but instead we used the related linear stability analysis to investigate if neutrino oscillations are unstable on small scales in a supernova. The work has resulted in a paper, which however does not form part of the thesis:

- [6] Sovan Chakraborty, Rasmus Sloth Hansen, Ignacio Izaguirre, and Georg Raffelt. “Self-induced flavor conversion of supernova neutrinos on small scales” (2015). arXiv: [1507.07569](https://arxiv.org/abs/1507.07569) [[hep-ph](#)].

The thesis is an attempt at uniting and binding together all these related topics, and for doing that, I have divided them into two general groups. In Part **I**, I have focused on the production of sterile neutrinos in the early Universe. That includes both the detailed calculation of collisional effects and the various attempts at limiting the production of sterile neutrinos. This leaves the two papers focused on chaos, which form the base of Part **II**, where I have also included a chapter introducing the topic of chaotic system analysis.

The papers have been incorporated in the different chapters rather than collecting them at the end of the thesis as it is often done. The motivation for this structure is to take advantage of the material covered in great detail in the papers for subsequent discussions. Hopefully, this approach also contributes to obtaining more coherence in the discussion of the different papers and adds to the perspective for each of them. The disadvantage is that the jumps between papers and material produced specifically for the thesis might take away a bit of the text flow, but bear with me if this happens. To help with the overview, I have marked sections of the thesis that are directly based on one of the papers with the symbol \mathfrak{P} before the section name, and each paper is ended by a large \mathfrak{P} . They have largely been left unchanged except for the formatting and certain sections that are covered in other parts of the thesis.

Acknowledgements

First of all I would like to thank my advisor Prof. Steen Hannestad for supervision and numerous interesting and enlightening discussions. I appreciate the opportunity that I have been given to test myself as a researcher and the many joyful experiences that come with it, and the value of an interested and dedicated supervisor cannot be understated. I would also like to thank Thomas Tram for introducing me to the exiting world of numerics and programming and for patiently explaining me the bliss of **LASAGNA**.

My thanks also goes to Yvonne Wong in Sydney, who hosted me during a research stay at University of New South Wales for almost half a year. The different perspective

on the world and on physics was invaluable, and the cheerful company was always inspiring.

A bit closer to home, I would like to thank Georg Raffelt for introducing me to the powers of linear stability analysis, and I would like to thank both him, Sovan Chakraborty and Ignacio Izaguirra for hospitality when I visited them last winter.

In Aarhus I have been part of a small but excellent group, and I am grateful for our many coffee breaks, delicious dinners and somehow infrequent Friday talks. My time here would not have been the same without Maria Archidiacono, Io Odderskov, Tobias Basse, Christian Schultz, Ole Eggers Bjælde, Jan Hamann, Sofie Koksang, Jeppe Mosgaard Dakin and naturally again Steen Hannestad and Thomas Tram. A thank you also goes to my many office mates for always providing a pleasant working environment.

Despite many prophecies of the opposite, I have somehow made it through my studies without becoming a coffee addict. Although there is no proof of a correlation, I would like to thank Kageklubben for ensuring a steady supply of cake through all my years in Aarhus.

Before I made it to university, my interests took a windy road until my plans for the future narrowed in on the hardcore natural sciences towards the end of high school, and I would like to thank my family for always being interested and for always supporting my projects. Even now, their interest is still there although I am sure that many of my ramblings must be incomprehensible.

Finally, I would like to thank Mia, my wife, for support and for sharing my interest for physics and astronomy. Thanks for our many interesting discussions and for always taking the time to listen when I needed someone to complain or talk to.

Summary

Neutrinos are some of the lightest and most elusive particles that we know of, and still, they can play an important role in certain extreme environments. On top of this, they have the interesting property that they can change from one flavour to another while they move through space because of mixing between flavour and mass. A feature which gives rise to many interesting oscillation phenomena.

This thesis is based on my work and publications through 4 years as a Ph.D. student. As we make our way through it, we move from relatively simple production of sterile neutrinos in the early Universe through a series of models that can prevent this same production to chaotic phenomena that can occur in models of neutrino oscillations for both the early Universe and for supernovae.

For the production of sterile neutrinos, we find the expected overall results, but there are some details of the commonly used approximations that should be noticed. Most prominently, all previous approximations fail to account correctly for scattering and/or annihilation processes in simplified collision terms. This shortcoming can be corrected by considering the two contributions separately and ensuring that only annihilations can change the neutrino number density.

When attempting to suppress the sterile neutrino production using an asymmetry between particles and antiparticles, one encounter severe numerical problems that are still unsolved. This has no influence on the qualitative conclusions, but stable quantitative results are hard to obtain and will be left for future work. Conversely, we find that the production can be suppressed by introducing interactions between the sterile neutrinos themselves in a numerically stable fashion. This mechanism works both for a vector- and a pseudoscalar force mediator although the obtained cosmologies are somewhat different.

Coming back to the particle asymmetry, we look at a long standing question regarding the existence of chaos for a very small mixing between active and sterile neutrinos. As several other groups have found, we detect chaotic behaviour using the momentum averaged approach, but conversely to previous results, we find no chaos when using the full momentum dependent equations.

A system that on the other hand turns out to reveal chaotic behaviour is a simple two-beam model for supernova neutrinos that we analyse extensively. Apart from the chaoticity, we find that some directions in parameter space are much more unstable than others, while others again are marginally stable. This might have implications for the stability of supernova neutrinos, although more realistic models are needed before any hard conclusions can be drawn.

Resumé

Neutrinoer er nogle af de letteste og mest flygtige partikler, vi kender til, og alligevel kan de spille en væsentlig rolle i visse ekstreme miljøer. Bortset fra dette, har de den interessante egenskab, at de kan skifte fra den ene type til den anden, mens de bevæger sig igennem rummet på grund af en blanding mellem type og masse. En egenskab, der giver anledning til mange interessante oscillationsfænomener.

Denne afhandling er baseret på mit arbejde og mine udgivelser igennem 4 år som ph.d. studerende. Mens vi gennemgår dette, vil vi bevæge os fra en relativt simpel produktion af sterile neutrinoer i det tidlige Univers, gennem en serie af modeller der kan undertrykke denne produktion, til kaotiske fænomener der kan opstå i modeller for neutrino oscillationer i både det tidlige Univers og i supernovaer.

For den sterile neutrinoproduktion i det tidlige Univers finder vi generelt set de forventede resultater, men der er nogle detaljer omkring de approksimationer, som ofte bliver anvendt, der bør bemærkes. Den mest fremtrædende mangel er, at ingen tidligere approksimationer giver en passende behandling af annihilations- og/eller spredningsprocesser i deres simplificerede kollisionsled. Denne fejl kan udbedres ved at behandle de to bidrag separat og sikre sig, at det kun er annihilationerne, der kan ændre på antalstætheden af neutrinoer.

Undertrykkelse af den sterile neutrinoproduktion med en partikel-antipartikel asymmetri giver anledning til alvorlige numeriske problemer, som ingen hidtil har løst på tilfredsstillende vis. Dette har ikke nogen betydning for de kvalitative konklusioner, men det er svært at opnå stabile kvantitative resultater, hvilket derfor er udskudt til fremtidigt arbejde. Omvendt finder vi, at produktionen kan undertrykkes på en numerisk stabil måde, hvis vi introducerer vekselvirkninger de sterile neutrinoer imellem. Denne mekanisme virker både med en vektor boson og en pseudoskalar som vekselvirkningspartikel, selvom de øvrige konsekvenser for kosmologi er noget forskellige.

Tilbage ved partikelasymmetrien undersøger vi et spørgsmål, der har været uløst længe vedrørende eksistensen af kaos for en meget lille blanding af aktive og sterile neutrinoer. Ligesom flere andre grupper tidligere har konstateret, ser vi kaotisk opførsel, når vi bruger en tilgang med midlet impuls, men modsat tidligere resultater finder vi ikke kaos, når vi benytter de fulde impulsaafhængige ligninger.

Et system, der derimod udviser kaotisk opførsel, er en simpel to-stråle model for supernovaneutrinoer, som vi analyserer indgående. Ud over den kaotiske natur, finder vi ud af, at nogle retninger i variabelrummet er mere ustabile end andre, mens nogle helt tredje er marginalt stabile. Dette kan have betydning for stabiliteten af supernovaneutrinoer, men før sådanne konklusioner kan drages, vil det være nødvendigt at analysere langt mere realistiske modeller.

Contents

Preface	i
Summary	v
Resumé	vii
1 Introduction	1
Part I Production of sterile neutrinos	3
2 General theoretical aspects	7
2.1 Introduction to neutrino oscillations	7
2.1.1 The MSW-resonance	9
2.2 Quantum kinetic equations	10
2.2.1 Liouville operator	11
2.2.2 Vacuum oscillations	11
2.2.3 Forward scattering on the background	12
2.2.4 Collision term	15
2.2.5 Two neutrino case	16
2.3 The expanding Universe	18
2.3.1 Deriving a variable	19
2.4 Numerical methods	20
3 Full collision term	23
3.1 \mathfrak{P} Active–sterile oscillations in the early Universe with full collision terms	24
3.1.1 Quantum kinetic equations	24
3.1.1.1 Repopulation and damping	25
3.1.1.2 Approximation schemes	26
3.1.1.3 Numerical implementation	30
3.1.2 Numerical results	32
3.1.2.1 Numerical convergence	33
3.1.2.2 Comparison of approximation schemes	35
3.1.2.3 Electron mass and Pauli blocking	36
3.1.2.4 Impact on distribution functions	36
3.1.2.5 Dependence on mixing parameters	39
3.1.3 Conclusions	42
3.1.A Derivation of the full collision terms	43
3.1.A.1 Neutrino–electron scattering in the s -channel	43

Contents

3.1.A.2	The massive case	46
3.1.A.3	The full collision terms	47
3.1.B	Repopulation and damping coefficients in the A/S approximation	50
3.2	Lepton asymmetries and active-active oscillations	52
4	Large lepton asymmetry	55
4.1	Quantum kinetic equations	56
4.2	Convergence problems	58
4.2.1	Solving the QKEs using RADAU5	61
4.2.2	Fixed minimal step size	63
4.2.3	Future improvements	65
5	Sterile interactions	67
5.1	Vector boson mediator	68
5.2	⌘ How self-interactions can reconcile sterile neutrinos with cosmology	70
5.3	BBN bounds and coupling to dark matter	77
5.4	Late time production	79
5.5	Pseudoscalar boson mediator	81
5.6	⌘ Cosmology with self-interacting sterile neutrinos and dark matter	83
5.7	Current state of sterile interactions and self-interacting dark matter	92
5.8	Dark matter production	93
5.8.1	Vector boson mediator	94
5.8.1.1	Averaged adiabatic approximation	94
5.8.1.2	Adiabaticity	95
5.8.1.3	Simulations	96
5.8.2	Improvements and a pseudoscalar mediator	97
Part II	Chaotic neutrino oscillations	101
6	Stability analysis	105
6.1	Linear stability analysis	105
6.2	Lyapunov analysis	106
6.2.1	Lyapunov exponents	106
6.2.2	Lyapunov vectors	108
6.2.3	Transient systems	110
7	Chaotic growth of the lepton asymmetry	111
7.1	⌘ Chaotic behaviour of the cosmological lepton asymmetry	113
7.1.1	Equations of motion	113
7.1.1.1	Quantum kinetic equations	113
7.1.1.2	Parametrisation of momentum space	115
7.1.1.3	Quantum rate equations	117
7.1.2	Quantifying chaos	118
7.1.3	Results	119
7.1.3.1	Quantum rate equations	119
7.1.3.2	Quantum kinetic equations	120

7.1.3.3	Lyapunov analysis of quantum kinetic equations	127
7.1.4	Conclusions	129
7.1.A	LASAGNA	129
7.1.A.1	Overview	129
7.1.A.2	ODE-solvers	130
7.1.A.3	Linear algebra solvers	130
7.2	Perspectives on the chaotic lepton asymmetry	132
8	Chaotic supernova neutrinos	133
8.1	Instabilities in collective supernova neutrino oscillations	133
8.2	⌘ Chaotic flavor evolution in an interacting neutrino gas	136
8.2.1	The two beam model	136
8.2.2	N momentum modes	136
8.2.3	Two momentum modes	137
8.2.4	Numerical results	139
8.2.4.1	Stationary solutions	141
8.2.4.2	Bipolar solutions	144
8.2.4.3	Non-periodic solutions	146
8.2.5	Conclusions	149
8.2.A	Numerical calculation of Lyapunov exponents and covariant Lyapunov vectors	150
8.3	Further developments	155
Part III	Future developments	157
	Experiments elucidating the elusive neutrino	159
	Other particular experiments	161
	Megaparsec milestones	162
	The next galactic supernova	163
	Bibliography	165
	List of Corrections	
	Note: There should be a paragraph on this in Chapter 8.	17

Introduction

One of the most abundant particles in the Universe is also one of the most weakly interacting. As such it is not strange that the neutrino was not detected until 1953, and that even today some of the most fundamental properties are unknown. The neutrino was first proposed by Pauli in the early 1930s as a desperate way of ensuring energy conservation in beta-decay, and while he did not have much confidence in the idea from the start, the evidence gradually confirmed the existence of the new particle. The first theoretical foundation came with the Fermi theory a few years later, and today we describe the neutrino and the electron as a $SU(2)$ -doublet in the electro-weak theory proposed by Weinberg and Salam. At low energies the $SU(2) \times U(1)$ symmetry of the theory is broken by the Higgs mechanism, and this results in the very different properties we observe for the electron and the neutrino. Like the electron has two siblings in the second and third generation of fermions in form of the muon and the tau particle, the neutrino also comes in three flavours as electron-, muon-, and tau neutrinos.

In the Standard Model of particle physics, which is partly based on the Weinberg-Salam theory, the neutrino is massless, and it has only a left-handed component while its antiparticle has only a right-handed component. In the 1990s this picture was proven incomplete with the realisation that at least some of the neutrinos must have a mass despite that a direct mass measurement has not been performed even today. Instead it was proven using observations of neutrinos from the atmosphere where the number of muon neutrinos depended on the baseline length even after taking geometric factors into account. The only satisfying explanation of this is that muon neutrino oscillate into tau neutrinos on their way from the production in the atmosphere to the detection just below the surface of the Earth.

The tension between the Standard Model and neutrino oscillations originates in the fact that oscillations are only possible for a non-zero mass difference between different propagation states as we will see in Section 2.1 on neutrino oscillations. This implies that some of the neutrinos are massive, and it also means that neutrino oscillations are easy to include ad hoc in the Standard Model by including a small mass term.

Another prominent source of neutrinos is the Sun, and since the first solar electron neutrinos were measured in the 1960s, the rate has been significantly lower than what was predicted by solar models. The explanation for this deficit was not clear until the Sudbury Neutrino Observatory in 2001 managed to measure not only the electron neutrinos, but also the muon and tau neutrinos. This established that the missing electron neutrinos do in fact turn up in the other flavour channels, and thus the difference can be explained by oscillations between the neutrinos. Most recently, the Borexino experiment has measured the electron neutrino flux from the pp -processes which produce the major part of the solar power output [7]. In accordance with neutrino oscillation theory, the relative result is higher than what was found for higher energy neutrinos,

1 Introduction

which again confirms our understanding of both neutrino oscillations and the Sun.

Neutrino oscillations were first suggested by Pontecorvo in 1957 [8], and it was this basic idea that was used for describing the atmospheric neutrino observations more than 40 years later. However, these vacuum oscillations fail to give a good description of solar neutrinos if they only describe the propagation from surface of the Sun to the Earth. Instead one needs to include the effect of the high electron density inside the Sun that was pointed out by Wolfenstein [9]. The many electrons in the background affect neutrino propagation and shift the effective masses differently for electron neutrinos compared to muon- and tau neutrinos. When this effect is exactly cancelled by the vacuum mass difference, a resonance arises as pointed out by Mikheev and Smirnov [10], which give rise to efficient MSW-conversion of neutrino flavour. For the highest energy solar neutrinos, this conversion is complete before they exit the Sun, and we find the very low electron neutrino flux that was measured already in the 1960s. The lower energy neutrinos, on the other hand, experience an incomplete conversion, and the result is that a relatively larger number of electron neutrinos can be detected on Earth.

Although the Sun is a very hot and dense environment, the Universe presents and has presented other regions of space and time that are even more extreme. Black holes are of course the prime example, but for neutrinos they are not that interesting as they are, well, black. The supernova collapse that can lead to a black hole on the other hand, is quite interesting along with the early Universe. Both cases present a large abundance of neutrinos and a high density of matter, which means that MSW-resonances are important for the neutrino flavour evolution. As it turns out, it is not only the electrons in the media that are important. For the early Universe one can imagine several variants of exotic physics that can modify the neutrino propagation just as the electrons does in the Sun. This will be our main focus in the first part of this thesis where we consider how sterile neutrino production can be suppressed or enhanced in the early Universe. For the second part of the thesis we will move our attention towards the more mathematically interesting non-linearity of neutrino oscillations. When the neutrinos are very abundant, they play an important role in setting the stage for their own propagation. This happens in supernovae and, if the conditions are correct, in the early Universe. Such a coupling of the neutrino flavour evolution back to the background of neutrino flavour, means that the equations become non-linear, and in some cases we find that this leads to both chaotic and collective behaviour of the neutrino oscillations.

A good understanding of these effects in supernovae and in the early Universe is important for two different reasons. In the early Universe, we need to understand neutrino oscillations well in order to properly describe the effects of various extensions to the Standard Model of particle physics. Furthermore, the rich phenomenology of neutrino oscillations might be the solution to some of the puzzles of modern particle physics and cosmology. These puzzles reach all the way from Earth-based neutrino experiments that find anomalies to small inconsistencies between different cosmological observations of dark matter. The motivation for supernovae is slightly more practical in nature. We expect a galactic supernova about every 30 years, and today we have the detector capacity to see a quite detailed signal. However, we are still trying to grasp the consequences of neutrino oscillations between the three neutrino flavours that are part of the Standard Model, and while some of the predictions seem robust, others depend a lot on the model used. This will hopefully improve significantly over the coming years as we learn more about the, at times chaotic, topic of collective neutrino oscillations.

Part I

Production of sterile neutrinos

To be, or not to be, that is the question.

— *Shakespeare's Hamlet*

One of the simplest extensions to the Standard Model of particle physics is to introduce one or more additional neutrinos. From precision measurements of the Z -resonance [11], we know that there can only be three active neutrinos with masses below m_Z , and hence a new neutrino would need to be a sterile $SU(2) \times U(1)$ singlet. This role could conveniently be played by a righthanded neutrino. The Standard Model only contains lefthanded neutrinos, and the addition of the righthanded would restore a sense of symmetry between right- and lefthanded particles.

One of the interesting roles that a sterile neutrino can have is to explain the tiny mass of the Standard Model neutrinos. This can be done using the seesaw mechanism, where the small mass of the normal neutrino is balanced by a large mass for the sterile neutrino. Otherwise, if we wish to have light sterile neutrinos, such a mechanism cannot be involved for all the sterile neutrinos. Although the simplest models involve one sterile neutrino for every active, models can be constructed which give three light active neutrinos from two heavy steriles and leave the last sterile with a small mass. Other alternatives for including light sterile neutrinos could be mirror theories, where every particle has a mirror partner which is mostly decoupled from the Standard Model, but also more general hidden sector theories can accommodate a light sterile neutrino.

Currently there are two main motivations for considering light sterile neutrinos. The first one is a number of hints from Earth-based experiments that all point to a mass difference of $\sim 1\text{eV}$, while the second motivation comes from dark matter where sterile neutrinos may constitute either warm or cool dark matter. The problem for both these cases is that the simple non-resonant production mechanism is unable to account for the correct cosmological abundance of the sterile neutrino. The eV-scale sterile tend to be overproduced while the dark matter candidate fails to meet its required abundance, given the constraints on mixing parameters.

The hints from experiments come from various sources. The first anomaly comes from the Liquid Scintillator Neutrino Detector (LSND) where an excess of electron anti-neutrinos were observed in a muon anti-neutrino beam [12]. This result is confirmed by MiniBooNE [13], where they also see a small excess of electron neutrinos in a beam of muon neutrinos. The second anomaly comes from reactor neutrinos where a reevaluation of the expected rate [14] of anti-neutrinos has given a deficit of 6% in all short baseline measurements. The third earth-based anomaly comes from calibration of gallium experiments where the GALLEX and SAGE experiments see fewer events than expected when they calibrate the detectors using radioactive sources [15, 16].

All of these experiments can be explained as additional oscillations into one or several sterile neutrinos with a mass of $\sim 1\text{eV}$, but none of the results are significant enough to confirm the hypothesis. Also, the various results are in tension if interpreted as the same signal. The MiniBooNE and LSND signals for antineutrinos are compatible, but the MiniBooNE neutrino data are only marginally compatible with a simple oscillation interpretation involving only one sterile neutrino. Additionally, the hints from reactor and gallium experiments point to a slightly different mass, although still around the same order of magnitude, and other accelerator experiments are excluding a significant part of the parameter space. The muon neutrino disappearance experiment KARMEN [17] has been in tension with the oscillation interpretation of the LSND anomaly for several years, but recently the ICARUS experiment published measurements that put the sterile neutrino interpretation under even more pressure [18]. Given the current state of contradicting results, the best approach probably is to consider the hints separately

and accept that at least some of them will be resolved by better understanding of the experiments and existing theory rather than new physics. Nonetheless, there are three anomalies that point to the same sterile neutrino mass range and if just one of them turn out to be a genuine signal, it would be a huge step forward in our understanding of particle physics. This is also the motivation behind the many experimental efforts towards testing the eV-scale sterile neutrino hypothesis [19, 20], and hopefully the issue will be settled in the near future.

On top of the internal tensions within and between the various anomalies, there are also constraints from cosmology as we already mentioned briefly. A 1eV sterile neutrino is fully thermalised in the early Universe as we will see in Chapter 3 if it has a mixing angle compatible with the neutrino anomalies, and the latest results from Planck [21] excludes such an addition to the usual cosmological model with cold dark matter and a cosmological constant (Λ CDM) at roughly 3σ . Hence, there has been an interest in models that can suppress the production of eV-scale sterile neutrinos, and in Chapter 4 and 5, we will describe and discuss some of the mechanisms that can give rise to such a suppression.

The second motivation — sterile neutrino warm dark matter — has been partly motivated by small scale problems for dark matter in the galaxy [22–24]. The promise of warm dark matter is to possibly resolve some of these problems due to its higher velocities compared to cold dark matter, but constraints from large scale structures are beginning to eliminate the relevant parameter space [25]. Furthermore, recent results suggest that the problems could be resolved by the complicated dynamics of baryonic feedback, an issue which will be discussed further in Chapter 5 on describing interacting sterile neutrinos.

Although this first motivation for warm dark matter seems to be fading, new and exiting observations have revived the interest even beyond the former level. The observed signal is a X-ray line at 3.5keV found in both various galaxy clusters [26] as well as in the Andromeda galaxy [27] and the Milky Way centre [28]. Such an X-ray line could come from the decay of a 7keV sterile neutrino, but both the signal and its origin has been questioned by several different groups [29–35]. Hopefully the issue will be settled once better dedicated data become available [36], but given the messy astrophysics that are involved in identifying such lines, there can be given no guarantee. Whether this is actually a sterile neutrino, some other exotic particle or astrophysics, the sterile neutrino remains a good candidate for dark matter. The main problem is that one needs a production mechanism different from the usual non-resonant oscillations as these cannot account for the full dark matter abundance given the constraints on mixing parameters from X-ray observations. In the end of Chapter 5, we consider a new mechanism that could potentially give rise to the correct dark matter abundance.

Before we get to the point of discussing all these current and interesting physics, we will however have to learn a bit about the equations that govern the neutrino oscillations in the early Universe in Chapter 2, where the quantum kinetic equations will be introduced along with a short discussion of the Friedmann equation and the numerical methods necessary for solving these.

General theoretical aspects

2.1 Introduction to neutrino oscillations

The phenomena of neutrino oscillations originates in the fact that different neutrinos are mixed. This means that a flavour eigenstate does not correspond directly to a single mass eigenstate, but is rather a linear combination of several mass states. While the flavour eigenstates are denoted by letters that indicate their family, the mass eigenstates are labelled with numbers. The mixing can be described using a mixing matrix, and for the three active neutrinos, the Pontocorvo-Maki-Nakagawa-Sakata matrix, which can be written as a product of three simple matrices, is the most commonly used

$$\begin{pmatrix} \nu_e \\ \nu_\mu \\ \nu_\tau \end{pmatrix} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos \theta_{23} & \sin \theta_{23} \\ 0 & -\sin \theta_{23} & \cos \theta_{23} \end{bmatrix} \begin{bmatrix} \cos \theta_{13} & 0 & \sin \theta_{13} e^{-i\delta_{\text{CP}}} \\ 0 & 1 & 0 \\ -\sin \theta_{13} e^{-i\delta_{\text{CP}}} & 0 & \cos \theta_{13} \end{bmatrix} \begin{pmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \end{pmatrix}. \quad (2.1)$$

With this parametrisation, the mixing is described neatly using three mixing angles, θ_{12} , θ_{23} , and θ_{13} and one CP-violating phase, δ_{CP} ¹.

For simplicity, it is often desirable to consider the mixing between two neutrino flavours at a time even though we know that three flavours exist. In this case, we can focus our attention to one of the submatrices, and taking the mixing of the electron- and muon neutrino as an example, we find that

$$\begin{pmatrix} \nu_e \\ \nu_\mu \end{pmatrix} = U \begin{pmatrix} \nu_1 \\ \nu_2 \end{pmatrix} = \begin{bmatrix} \cos \theta_{12} & \sin \theta_{12} \\ -\sin \theta_{12} & \cos \theta_{12} \end{bmatrix} \begin{pmatrix} \nu_1 \\ \nu_2 \end{pmatrix}, \quad (2.2)$$

where ν_1 and ν_2 are the propagating mass eigenstates. If there are no neutrino masses, all the mass eigenstates are degenerate, and the transformation in Equation (2.2) would have no physical meaning. This becomes even clearer when considering the time evolution. In presence of mixing, the Hamiltonian for the flavour states can be approximated as²

$$H \approx U \frac{1}{2E} \begin{bmatrix} m_1^2 & 0 \\ 0 & m_2^2 \end{bmatrix} U^\dagger, \quad (2.3)$$

using the ultra-relativistic approximation $E_i = \sqrt{p^2 + m_i^2} \approx p + m_i^2/2E$, and ignoring the mass independent term proportional to p . Calculating the matrix products, subtracting

¹For Majorana neutrinos, there are two additional phases.

²This approach is not valid if the mixing matrix changes as a function of time as it is for instance the case if there is a matter background. The derivation by Fukugita and Yanagida [37] is more general.

2 General theoretical aspects

the constant $(m_1^2 + m_2^2)/4E$, and using double-angle formulas, the final Hamiltonian is

$$\begin{aligned} H &\approx \frac{1}{2E} \begin{bmatrix} \cos^2 \theta_{12} m_1^2 + \sin^2 \theta_{12} m_2^2 & \cos \theta_{12} \sin \theta_{12} (m_2^2 - m_1^2) \\ \cos \theta_{12} \sin \theta_{12} (m_2^2 - m_1^2) & \sin^2 \theta_{12} m_1^2 + \cos^2 \theta_{12} m_2^2 \end{bmatrix} - \frac{m_1^2 + m_2^2}{4E} \\ &= \frac{\delta m_{21}^2}{4E} \begin{bmatrix} -\cos 2\theta_{12} & \sin 2\theta_{12} \\ \sin 2\theta_{12} & \cos 2\theta_{12} \end{bmatrix}, \end{aligned} \quad (2.4)$$

where $\delta m_{21}^2 = m_2^2 - m_1^2$. The time evolution of the flavour eigenstates can now be formulated as a Schrödinger equation

$$i \frac{d}{dt} \begin{pmatrix} \nu_e \\ \nu_\mu \end{pmatrix} = \frac{\delta m_{21}^2}{4E} \begin{bmatrix} -\cos 2\theta_{12} & \sin 2\theta_{12} \\ \sin 2\theta_{12} & \cos 2\theta_{12} \end{bmatrix} \begin{pmatrix} \nu_e \\ \nu_\mu \end{pmatrix}. \quad (2.5)$$

From this, it is very clear that there would be no mixing if $\delta m_{21}^2 = 0$. In order to find the time evolution, the exponential of the Hamiltonian is needed, and exploiting the properties of Pauli matrices, it is

$$\exp(-iHt) = \cos \left(\frac{\delta m_{21}^2 t}{4E} \right) I + i \sin \left(\frac{\delta m_{21}^2 t}{4E} \right) \begin{bmatrix} -\cos 2\theta_{12} & \sin 2\theta_{12} \\ \sin 2\theta_{12} & \cos 2\theta_{12} \end{bmatrix}. \quad (2.6)$$

Now, it is easy to calculate the oscillation probability from ν_e to ν_μ as function of time

$$P(\nu_e \rightarrow \nu_\mu) = |\langle \nu_\mu | \nu_e \rangle|^2 = \left| \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} e^{-iHt} \begin{pmatrix} 1 \\ 0 \end{pmatrix} \right|^2 = \sin^2 \left(\frac{\delta m_{21}^2 t}{4E} \right) \sin^2 2\theta_{12}, \quad (2.7)$$

and similarly for the other oscillation probabilities. This type of oscillations is called vacuum oscillations because they occur even when there is no background medium. They are especially important for understanding neutrinos from the atmosphere, nuclear reactors, and particle accelerators, and most of the parameters are well known.

The mixing angles are measured to be [38]

$$\begin{aligned} \sin^2(2\theta_{12}) &= 0.846 \pm 0.021, \\ \sin^2(2\theta_{23}) &= 0.999_{-0.018}^{+0.001}, \text{ }^3 \\ \sin^2(2\theta_{13}) &= 0.093 \pm 0.008. \end{aligned}$$

This means that there is an at least almost, if not exactly, maximal mixing between the muon neutrino and the tau neutrino, a large mixing between the electron neutrino and the muon neutrino, and a non-negligible mixing between the electron neutrino and the tau neutrino. Regarding the mass differences, the values known today are [38]

$$\begin{aligned} \delta m_{21}^2 &= \delta m_2^2 - \delta m_1^2 = (7.53 \pm 0.18) \times 10^{-5} \text{eV}^2, \\ |\delta m_{32}^2| &= |\delta m_3^2 - \delta m_2^2| = (2.44 \pm 0.06) \times 10^{-3} \text{eV}^2. \text{ }^3 \end{aligned}$$

Since the oscillation length is inversely proportional to the squared mass difference, we have two lengths of oscillation separated by two orders of magnitude. This means that an experiment typically only probes one of the mass differences and some combination of

³for normal mass hierarchy

the mixing angles. As indicated, the oscillations can only probe the difference between two masses, and this leaves the absolute masses of the neutrinos unknown. Furthermore, the sign of δm_{32}^2 is unknown, and this gives the possibility of $m_3 > m_2, m_1$, called the normal hierarchy, and $m_3 < m_2, m_1$, called the inverse hierarchy. Actually, the exact values of the mixing parameters depend a little on the mass hierarchy, and the values mentioned above are for the normal mass hierarchy, but the inverted hierarchy give almost identical values.

Neutrino oscillations in vacuum are quite easy to describe in terms of a couple of angles and a CP-violating phase. When any kind of matter is introduced, this picture becomes more complicated. This was realised by Lincoln Wolfenstein [9] and the theory was further developed by Stanislav Mikheev and Alexei Smirnov [10] as previously mentioned in the Introduction. Today other effects are known as well, and the originally discovered effect is called MSW-oscillations. For moderate densities of matter, the only important effect is coherent forward scattering on electrons which changes the effective mass of the neutrinos, thus altering the oscillation length. For higher densities, incoherent scattering with electrons and neutrino-neutrino self interactions can also play a role as we will see both in this part of the thesis, but also in Part II.

2.1.1 The MSW-resonance

When neutrinos travel through matter, they forward scatter on electrons as well as on nucleons, which we will come back to in Section 2.2.3. For now, it is sufficient to know that the diagonal of the Hamiltonian from Equation (2.4) is modified by an effective potential, $V_e = \sqrt{2}G_F n_e$

$$H = \begin{bmatrix} -\frac{\delta m_{21}^2}{4E} \cos 2\theta_{12} + V_e & \frac{\delta m_{21}^2}{4E} \sin 2\theta_{12} \\ \frac{\delta m_{21}^2}{4E} \sin 2\theta_{12} & \frac{\delta m_{21}^2}{4E} \cos 2\theta_{12} \end{bmatrix}. \quad (2.8)$$

The vacuum Hamiltonian can be represented as a scalar multiplication and a unitary matrix as seen in Equation (2.4), and by adding a suitably constant, which does not affect the neutrino oscillations, the MSW Hamiltonian can get the same form. The result is a matter mixing angle and matter mass difference [37]

$$\sin 2\theta_m = \sin 2\theta \frac{\delta m^2}{\delta m_m^2}, \quad \delta m_m^2 = \delta m^2 \sqrt{(2EV_e/\delta m^2 - \cos 2\theta)^2 + \sin^2 2\theta}. \quad (2.9)$$

In the limit of vanishing matter potential, we recover the vacuum result, while for very large values of V_e , the mass difference becomes large. This also means that the matter mixing angle becomes very small as it is suppressed by $\delta m^2/\delta m_m^2$. For intermediate values, we can have the situation that $2EV_e/\delta m^2$ and $\cos 2\theta$ cancels, which corresponds to maximal mixing of $\sin 2\theta_m = 1$ and $\cos 2\theta_m = 0$. This is called the MSW resonance, and except for having a very large mixing around the resonance, it is also possible to have a significant conversion of one neutrino flavour into another. As we already learnt, the higher energy neutrinos from the Sun are a good example of this. Electron neutrinos are produced in the centre of the Sun, after which they propagate outwards. As they see the changing background density of electrons, they go through the above mentioned resonance adiabatically, and they emerge as almost pure ν_2 mass eigenstates. The ν_2 is an almost equal mixture of ν_e , ν_μ and ν_τ , and as a consequence, we only measure a third of the electron neutrinos that we would expect without neutrino oscillations.

2.2 Quantum kinetic equations

An ensemble of neutrinos and antineutrinos in a thermal plasma can be described using the density matrix formalism, where the a $2n \times 2n$ density matrix describes n mixed neutrinos, and basically contains the number distribution functions on the diagonal and correlations between the different neutrinos in the off-diagonal entries. We will assume that neutrino-antineutrino correlations and the associated spin-flip can be neglected.⁴ This breaks the density matrix into two submatrices such that we can describe the neutrinos by the matrix, ρ , and the antineutrinos by the matrix, $\bar{\rho}$. The density matrix for neutrinos can be expressed as $\rho_{ij}(\mathbf{p}) = \langle a_i^\dagger(\mathbf{p})a_j(\mathbf{p}) \rangle$, where $a_i^\dagger(\mathbf{p})$ and $a_i(\mathbf{p})$ are the creation and annihilation operators for ν_i with momentum \mathbf{p} . Similarly, we can define the density matrix for antineutrinos as $\bar{\rho}_{ij}(\mathbf{p}) = \langle \bar{a}_i^\dagger(\mathbf{p})\bar{a}_j(\mathbf{p}) \rangle$ where $\bar{a}_i^\dagger(\mathbf{p})$ and $\bar{a}_i(\mathbf{p})$ are the creation and annihilation operators for $\bar{\nu}_i$. This is one of two conventions often used [45–48]. The advantage of this convention is that the definition is symmetric, while the other convention uses the antineutrino density matrix $\bar{\varrho}_{ij} = \langle \bar{a}_j^\dagger(\mathbf{p})\bar{a}_i(\mathbf{p}) \rangle$. This convention is most commonly used when treating neutrino oscillations in supernovae [49–53], but has also been applied to the early Universe [54, 55] and to more general treatments of neutrino oscillations [56]. The advantage of exchanging the indices is a more straight forward description of neutrino background effects which depend on $\rho - \bar{\varrho} = \rho - \bar{\rho}^*$. For all chapters on the early Universe, we will use the simple notation $\bar{\rho}$, but for Chapter 8, we will follow the supernova neutrino literature and use $\bar{\varrho}$. For now, the quantum kinetic equation (QKE) will be described for ρ and $\bar{\rho}$, and if it is needed for $\bar{\varrho}$, it can easily be derived.

To second order in the weak interactions, the most general form of the QKE, which describes the density matrix evolution, is

$$iL[\rho] = [H_{\text{vac}}, \rho] + [H_{\text{int}}, \rho] + C[\rho]. \quad (2.10)$$

On the left hand side, we have the Liouville operator, which describes the propagation of non-mixed, non-interacting particles if all terms on the right hand side are zero. The first term on the right hand side is the commutator between the density matrix and the vacuum Hamiltonian, H_{vac} . It describes the vacuum oscillations of mixed neutrinos. The next term similarly describes the effect of forward scattering on the background. This affects the effective mass states of the neutrinos, and lead to additional oscillation phenomena as we just saw in Section 2.1.1. Finally, the last term is a genuine second order effect and contains all non-forward collision processes between the neutrinos and the background as well as between the neutrinos themselves. The corresponding equation for antineutrinos has the same form, with the corresponding quantities \bar{H}_{vac} , \bar{H}_{int} , and $\bar{C}[\bar{\rho}]$ which we will describe later.

The task of solving Equation (2.10) in its most general form is overwhelming as it is a partial differential equation in the three dimensions of momentum space and three dimensions of real space as well as time. When we apply the QKE, we therefore need to find appropriate approximations to reduce the number of dimensions. Furthermore, we can often neglect some of the interaction terms on physical grounds because of the specific system we are considering.

⁴This is not necessarily the case in anisotropic systems [39–43], see [44] for a recent review.

In the following few sections we will review the different terms, the most common approximations and discuss the issues that these might raise.

2.2.1 Liouville operator

The Liouville operator basically gives the full time derivative. Including all dependencies on position and momentum, it has the form

$$L[\cdot] = \partial_t[\cdot] + \dot{\mathbf{x}} \cdot \nabla_{\mathbf{x}}[\cdot] + \dot{\mathbf{p}} \cdot \nabla_{\mathbf{p}}[\cdot], \quad (2.11)$$

where the dot denotes a partial time derivative. The first two terms contain derivatives of time and space and turn out to be dominating in different interesting models. The last term depends on the overall force on the neutrinos through $\dot{\mathbf{p}}$, and throughout this thesis it will be assumed to be zero. As neutrinos only feel gravity and the short ranged weak force, this is a good assumption when the gravitational potentials are not deep. For neutrinos, an appreciable effect would require a large disturbance of space-time as they are almost massless and therefore ultra-relativistic.

The two environments of interest are the early Universe and exploding supernovae. In the early Universe, we know that the background is very homogeneous and isotropic, which means that the second term on the right hand side can be neglected. The force term can similarly be neglected as we argued above because there is no net gravitational force on the neutrinos. In a supernova, the situation is quite different. Here, there is a preferred direction, and the positional derivative is important. On the other hand, neutrinos move with almost the speed of light, and the timescale at which they oscillate is much smaller than the timescales at which the supernova explodes. This means that the neutrinos see a stationary background to a very good approximation, and the time derivative term is negligible. A supernova explosion can leave behind a black hole, and in this case the space-time is most certainly deformed. The reasons we still neglect the force contribution to Equation (2.11) are twofold. Firstly, if a black hole is formed, almost all neutrino emission stops, and such a system is not that interesting from the point of neutrino oscillations as most of the interesting phenomena happen because of the large neutrino flux that usually comes from a supernova. Secondly, the densest parts of the supernova are not that interesting either, as neutrinos are trapped there, and the oscillation phenomena that we study using the QKE only happen once the neutrinos have started to stream out of the proto neutron star.

Although the above mentioned assumptions are reasonable, and there are good argument for them to hold true, it might still be interesting to study whether neutrino oscillations inside the core are as simple as we think they are, and whether inhomogeneities in the early Universe could effect neutrino oscillations. These topics are largely unexplored, but will not be covered further in this thesis, although the chaotic lepton asymmetry of Chapter 7 relates somewhat to the latter issue.

2.2.2 Vacuum oscillations

Neutrino oscillations in vacuum are due to neutrino mixing and the differences in neutrino masses. As we showed in Section 2.1, the full Hamiltonian can be well approximated by an ultra-relativistic approximation, which depends on the mass squared difference

2 General theoretical aspects




Figure 2.1: Forward scattering diagrams.

Forward scattering contributions to the effective neutrino mass at second order. f in the tadpole diagram can be any fermion coupling to the Z-boson, while l in the bubble diagram is a lepton of the same flavour as the neutrino.

between the neutrinos, $\delta m^2 = m_2^2 - m_1^2$. In the density matrix formalism, the same approximations can be done, and yield the similar result

$$H_{\text{vac}} = \frac{U^\dagger M_{\text{mass}}^2 U}{2E}, \quad (2.12)$$

where the squared mass matrix in the neutrino flavour basis is given in terms of the mass matrix in the mass basis and the mixing matrix U . In absence of any background effect, this means that the neutrinos will oscillate independently according to their energy and only depend on the mixing parameters (mixing angles and masses). For antineutrinos, we find the same expression for \bar{H}_{vac} .

2.2.3 Forward scattering on the background

When a particle travels through a medium, it acquires contributions to the effective mass from the medium as it was already discussed in Section 2.1.1. This happens through forward scatterings, where the momentum of the particle is not changed. Since the mass enters the QKE as $[H_{\text{vac}}, \rho]$, we can similarly describe the mass shift by the term $[H_{\text{int}}, \rho]$, and the different contributions to H_{int} are often referred to as background potentials. For neutrinos, the main contributors to the background potential at energies below 100MeV are electrons, neutrons, and neutrinos themselves. Protons could also contribute when present, but it turns out that the contribution is only minor due to the value of the weak mixing angle [57]. The two second order Feynman diagrams contributing to forward scattering are seen in Figure 2.1. The tadpole diagram on the left gives a contribution from all neutrinos and antineutrinos, the electron, the positron, the neutron, and a small one from the proton. The bubble diagram on the right only gives a contribution from the neutrino itself and the corresponding lepton, as lepton flavour is conserved by the weak interaction.

The contribution from neutrons to the background potential is [57]

$$H_{\text{int}, n} = -\frac{\sqrt{2}G_F}{2} (n_n - n_{\bar{n}}), \quad (2.13)$$

where G_F is the Fermi constant and $n_{n(\bar{n})}$ is the density of (anti) neutrons. Similarly the proton contribution is

$$H_{\text{int}, p} = \frac{\sqrt{2}G_F}{2} (1 - 4 \sin^2 \theta_W) (n_p - n_{\bar{p}}), \quad (2.14)$$

where θ_W is the weak mixing angle. As $\sin^2 \theta_W = 0.23864$ [38], this contribution is significantly suppressed relative to the neutron contribution. The only diagram contributing to these terms is the tadpole diagram as the baryons could not conserve lepton number in the bubble diagram. For the electrons and neutrinos, this is not true in general, and the combination of both diagrams can give rise to slightly more complicated expressions.

The contribution from electrons to the background potential for electron neutrinos is given by [57]

$$H_{\text{int}, e} = \frac{\sqrt{2}G_F}{2} (1 + 4 \sin^2 \theta_W) (n_{e^-} - n_{e^+}) - \sqrt{2}G_F \frac{8p}{3m_W^2} (u_{e^-} + u_{e^+}), \quad (2.15)$$

where n_{e^\pm} is the number density of positrons and electrons, p is the momentum of the neutrino, m_W is the W-boson mass, and u_{e^\pm} is the energy density of positrons and electrons. Notice that the first term, which is proportional to G_F , requires an asymmetry between the number of electrons and positrons just as the terms from neutrons and protons require an asymmetry, while the second term, with an additional suppression of G_F/m_W^2 , is proportional to the total energy density of electrons and positrons. For the muon- and tau neutrinos, the bubble diagram is not present, and the contribution is simply the same as it was for the protons, only with a different sign [57]

$$H_{\text{int}, e} = -\frac{\sqrt{2}G_F}{2} (1 - 4 \sin^2 \theta_W) (n_{e^-} - n_{e^+}). \quad (2.16)$$

The last contribution comes from the neutrinos themselves. The most general expression, incorporating both tadpole- and bubble diagrams and assuming that all neutrinos are described by the density matrices, is [56]

$$\begin{aligned} H_{\text{int}, \nu}(\mathbf{p}) &= \sqrt{2}G_F \int d\mathbf{q} \{G_S \text{Tr}((\rho(\mathbf{q}) - \bar{\rho}^*(\mathbf{q}))G_S) + G_S(\rho(\mathbf{q}) - \bar{\rho}^*(\mathbf{q}))G_S\} (1 - v_{\mathbf{q}} \cdot v_{\mathbf{p}}) \\ &\quad - \sqrt{2}G_F \frac{8p}{4m_Z^2} \int d\mathbf{q} q G_S(\rho(\mathbf{q}) + \bar{\rho}^*(\mathbf{q}))G_S (1 - v_{\mathbf{q}} \cdot v_{\mathbf{p}})^2. \end{aligned} \quad (2.17)$$

Here, G_S is a matrix of coupling constants, and in the standard model it is simply the identity matrix, while $v_{\mathbf{p}}$ and $v_{\mathbf{q}}$ are the velocity vectors for the given momenta. The integrals in Equation (2.17) run over the momentum vector, and the first integral contains two terms. The first term corresponds to the contribution from the tadpole diagram, and the trace means that all neutrinos in the background give a contribution to this part of the potential. The second term, on the other hand, only couples to each neutrino since it is related to the bubble diagram, and it accounts for the well known result, that neutrinos of the same species contribute with a potential twice the value of that from each of the other neutrinos [57, 58]. It is also worth noticing that, again, the first part of the expression depends on the difference between the number density of particles and antiparticles and is proportional to G_F , while the second integral depends on the sum of the energy densities, and is proportional to G_F/m_Z^2 . In an isotropic medium such as the early Universe, the $(1 - v_{\mathbf{q}} \cdot v_{\mathbf{p}})$ term integrate to 1, while $(1 - v_{\mathbf{q}} \cdot v_{\mathbf{p}})^2$ integrate to $\frac{4}{3}$, but for non-isotropic media such as a supernova, they are important.

2 General theoretical aspects

A common approximation is to ignore the oscillations between some of the neutrinos in order to significantly simplify the equations. However, the background contribution from these additional neutrinos, ν_β , can still be accommodated easily. Assuming that the background neutrino has standard model couplings, we can easily multiply out the matrices and perform the integrals. As already mentioned, it turns out that it is the trace term that accounts for the effect of other neutrinos in the background, and the result is simply

$$H_{\text{int}, \nu_\beta} = \sqrt{2}G_F \left(n_{\nu_\beta} - n_{\bar{\nu}_\beta} \right), \quad (2.18)$$

which resembles Equation (2.13) and (2.14), just with a different front factor and neutrino densities instead of neutron or proton densities.

In all the cases above, the Hamiltonian for the antineutrino, $\bar{H}_{\text{int}, x}$, can be found by reversing the sign of the first term, which is proportional to the difference between the number of particles and antiparticles, and taking the complex conjugate of all terms. This also corresponds to exchanging particles and antiparticles.

As we found in Section 2.1, oscillations only depend on the difference of squared masses between the neutrinos. This also means that the absolute potential of each neutrino is irrelevant, and only the difference between the background potentials of the different neutrinos matter. Using only this observation and some basic knowledge about neutrino interactions, we can reduce the number of interesting background particles further in each particular case.

In the early Universe, we will mostly be interested in oscillations between standard model neutrinos and sterile neutrinos, and therefore all contributions to the background should be important. However, the environment is almost perfectly symmetric between particles and antiparticles when we only consider the standard model interactions, which means that the asymmetric parts of the potentials are suppressed. We know that there is a surplus of baryons over antibaryons of 10^{-10} relative to the density of photons, but often it can be ignored in the context of neutrino oscillations. For the electrons and positrons one expects a similar asymmetry to keep the Universe uncharged, but again the contribution is very small. Finally, the neutrino asymmetries are not very well constrained as the neutrinos are uncharged, and a large asymmetry could potentially be present. However, from the standard model itself, we do not have a good reason to think that there is a large asymmetry between neutrinos and antineutrinos. Therefore, in simple models it is often useful to neglect all the asymmetric contributions to H_{int} and only include the symmetric one. For the electron neutrinos, this leaves the expression

$$H_{\text{int}, \text{sym}} = -\frac{\sqrt{2}G_F 8p}{3m_Z^2} \left((1 - \sin^2 \theta_W) (u_{e^-} + u_{e^+}) G_e + \int d\mathbf{q} q G_S(\rho(\mathbf{q}) + \bar{\rho}(\mathbf{q})) G_S \right), \quad (2.19)$$

where G_e is a matrix with a one on the diagonal entry corresponding to the electron neutrino, and we have used that $\cos \theta_W = m_W/m_Z$.

In this thesis, we will have two occasions to study the leptonic part of the asymmetric term in the early Universe. The first reason is that a large lepton asymmetry can significantly affect the production of sterile neutrinos through oscillations as we will see in Chapter 4, while the second reason is that the lepton asymmetry can give rise to very non-linear oscillations which are intrinsically interesting as we will see in Chapter 7.

Collecting all the asymmetric terms, we can express all of it as

$$H_{\text{int, asym}} = \sqrt{2}G_F L, \quad (2.20)$$

where the lepton asymmetry is

$$\begin{aligned} L = & -\frac{1}{2}(n_n - n_{\bar{n}}) + \frac{1}{2}(1 - 4\sin^2\theta_W)(n_p - n_{\bar{p}}) - \frac{1}{2}(1 - 4\sin^2\theta_W)(n_{e^-} - n_{e^+}) \\ & + G_e(n_{e^-} - n_{e^+}) + \int d\mathbf{q} \{G_S \text{Tr}((\rho(\mathbf{q}) - \bar{\rho}(\mathbf{q}))G_S) + G_S(\rho(\mathbf{q}) - \bar{\rho}(\mathbf{q}))G_S\}. \end{aligned} \quad (2.21)$$

With this definition of the lepton asymmetry, we include the contributions from the baryons as well as contribution from the leptons, but most often the baryon contribution will be insignificant in the early Universe.

In supernovae, there is a large surplus of protons, neutrons, and electrons compared to the antiparticles, but we will only consider oscillations between standard model neutrinos which means that the potential from neutrons and protons is unimportant. Similarly, can we neglect the contribution from the electrons in Equation (2.16), and combining this with the background contribution from the neutrinos themselves, we end up with

$$H_{\text{int, asym}} = \sqrt{2}G_F(n_{e^-} - n_{e^+}) + \sqrt{2}G_F \int d\mathbf{q}(\rho(\mathbf{q}) - \bar{\rho}(\mathbf{q})), \quad (2.22)$$

where we have neglected the symmetric parts of the potential as they are suppressed by m_Z^2 or m_W^2 compared to the asymmetric parts. Furthermore, we have used that G_S is the identity matrix and the trace part of the neutrino background is the same for all the neutrinos and therefore unimportant.

2.2.4 Collision term

The coherent forward scattering processes that we have just discussed are important in all of the places where neutrinos oscillate in interesting ways, and the incoherent scattering processes described by the collision term can often be ignored. However, in some environments, the density is so large that the incoherent scatterings are also important, and some of the associated complications will now be introduced.

In general, the collision term can be expected to be more complicated than the forward scatterings as it is a second order effect whereas the forward scattering is only first order in the Fermi constant. This is indeed the case if one attempts to write them out in terms of matrix products similar to Equation (2.17), but it turns out that integrated versions of the full expressions give quite good approximations as we will see in Chapter 3. Starting with the general expression, it can be written in the compact form [45, 59]

$$\begin{aligned} C[\rho_{ij}] = & \pi \mathcal{V} [k', l' | m', n'] \mathcal{V} [k, l | m, n] \delta_E(k', l' | m', n') \\ & [\rho_{n'l} \rho_{m'k} (\delta_{mk'} - \rho_{mk'}) (\delta_{il'} - \rho_{il'}) - \rho_{mk'} \rho_{il'} (\delta_{km'} - \rho_{km'}) (\delta_{n'l} - \rho_{n'l})] \delta_{nj} \\ & + \text{complex conjugate}, \end{aligned} \quad (2.23)$$

where the indices indicate the various particles and summation over all indices except for i and j is implied. In this expression, \mathcal{V} contains the matrix element for each

2 General theoretical aspects

process⁵, δ_E is a Dirac delta function in energy, δ_{xy} is the Kronecker delta function, and ρ_{xy} is the density matrix correlator between x and y . As we assume no correlation between the oscillating neutrinos and the background, many of the matrix elements are zero, and even more terms can be discarded as the processes are unphysical, and the corresponding matrix elements are zero.

We will only be concerned about the collisional term for oscillations between one active, ν_α and one sterile neutrino, ν_s , and in this case, the collision term turns out to separate nicely into different parts,

$$C[\rho] = \mathcal{R}[\rho] + \mathcal{D}[\rho]. \quad (2.24)$$

The first part describes the equilibration of the active neutrino, and the second part describes the loss of coherence when the neutrinos interact with the background. If Pauli blocking is neglected, and most of the distributions are assumed to be equilibrium distributions, we get the simple integrated expressions [46, 58, 60]

$$\mathcal{R}[\rho_{\alpha\alpha}] = C_\alpha G_F^2 p T^4 (f_{\text{eq}} - \rho_{\alpha\alpha}), \quad (2.25)$$

$$\mathcal{D}[\rho_{ij}] = -\frac{1}{2} C_\alpha G_F^2 p T^4 \rho_{ij}, \quad (2.26)$$

where $C_e \approx 1.27$, $C_{\mu/\tau} \approx 0.92$, and $j \neq i$. It is possible to make other and better approximations to the full expression in Equation (2.23), but this discussion will be postponed to Chapter 3, and the approximations in Equation (2.25) and (2.26) will be used for all other chapters where collisional effects are considered.

2.2.5 Two neutrino case

A model with n neutrinos requires a $n \times n$ density matrix to describe all the oscillations, but it is often possible to capture the essential part of the phenomena using only two oscillating neutrinos and a corresponding 2×2 density matrix. This simplifies the problem significantly, and allows an expansion of the density matrix in terms of Pauli spin matrices, $\boldsymbol{\sigma}$.

$$\rho = \frac{f_0(p)}{2} (P_0 + \mathbf{P} \cdot \boldsymbol{\sigma}), \quad (2.27)$$

where $f_0(p) = 1/(1 + e^{p/T})$, and correspondingly for the antineutrinos. With this parametrisation and neglecting the collision terms, the QKE from Equation (2.10) can be expressed as

$$\begin{aligned} L[\mathbf{P}] &= \mathbf{V} \times \mathbf{P} \\ L[P_0] &= 0. \end{aligned} \quad (2.28)$$

Here, \mathbf{V} is the potential which can be calculated from the Hamiltonians H_{vac} and H_{int} . The formulation in terms of polarisation vectors highlights the similarity between neutrino oscillations and a spherical pendulum. This is especially useful in the outer parts of a supernova, where the collision terms are negligible, and Equation (2.28) provide a good description of the oscillations [49]. In the early Universe, this analogy is less useful, but it still puts the equations on a nice form.

⁵The full expression is given in Chapter 3.

When reformulating the Hamiltonian as the vector \mathbf{V} , we find that

$$H - \frac{1}{2}\text{Tr}(H) = \frac{1}{2} \begin{pmatrix} V_z & V_x - iV_y \\ V_x + iV_y & -V_z \end{pmatrix}, \quad (2.29)$$

where the factor of $\frac{1}{2}$ accounts for the transition from the commutator formulation in Equation (2.10) to the cross product formulation in Equation (2.28). Using this new parametrisation, the vacuum term gives

$$V_x = \frac{\delta m^2}{2p} \sin 2\theta, \quad V_z = -\frac{\delta m^2}{2p} \cos 2\theta. \quad (2.30)$$

Similarly, the forward scattering Hamiltonian can be reformulated, and for the non-neutrino background, it turns out that the only contribution is to the z -component of \mathbf{V} .

$$V_z = -H_{\text{int}, \nu_a} + H_{\text{int}, \nu_b}, \quad (2.31)$$

where ν_a and ν_b are the first and second oscillating neutrino respectively. For the neutrino background, it is a bit more complicated as there can be correlations between the oscillating neutrinos, and this introduces off-diagonal terms. The matrix structure of Equation (2.17), describing the neutrino background potential, ensures that there are no off-diagonal terms if the oscillations are between an active and a sterile neutrino, and in that case, Equation (2.31) is also valid for the neutrino background. However, for active-active oscillations, the picture is more complicated, and we will consider this case when it is relevant.

In the early Universe, collisions are important, and we need to include the collision term. This is quite complicated in the most general form, but for active-sterile neutrino oscillations, it simplifies significantly as mentioned in the previous section. In this case, we can write the equations on the form

$$\begin{aligned} L[\mathbf{P}] &= \mathbf{V} \times \mathbf{P} + \frac{R}{f_0} \hat{\mathbf{z}} - D\mathbf{P}_T, \\ L[P_0] &= \frac{R}{f_0}, \end{aligned} \quad (2.32)$$

where R and D are the collision terms describing repopulation and damping, and $\mathbf{P}_T = P_x \hat{\mathbf{x}} + P_y \hat{\mathbf{y}}$ is the transverse part of the polarisation vector \mathbf{P} . R and D can be directly derived from the approximate expressions in Equation (2.25) and (2.26), and in this case they give

$$R = C_\alpha G_F^2 p T^4 \left(f_{\text{eq}} - \frac{f_0}{2} (P_0 + P_z) \right), \quad (2.33)$$

$$D = \frac{1}{2} C_\alpha G_F^2 p T^4. \quad (2.34)$$

When these terms are important, it can be an advantage to take a step back and get rid of P_0 and P_z in favour of two variables that directly relate to the active and sterile component. This can be done by defining $P_a = P_0 + P_z$ and $P_s = P_0 - P_z$, but now the resulting parametrisation resembles the density matrix quite a lot.

Fixme Note: There should be a paragraph on this in Chapter 8.

2 General theoretical aspects

Like the Pauli matrices were used as the basis for a two times two matrix here, it is also possible to use Gell-Mann matrices for three times three matrices. This means that the formalism used here can also be developed for three oscillating neutrinos [60], but the most common approach is to stay in the density matrix language when dealing with more than two oscillating neutrinos.

2.3 The expanding Universe

The Universe is homogeneous and isotropic on very large scales even today. The structures that we observe around us such as the Earth, the Sun, the Milky Way and the Local Group originate in small inhomogeneities that have grown under the influence of gravity and eventually also affected by the electromagnetic force.

In the first few seconds after Big Bang, the inhomogeneities were very small on all scales, and they can be ignored to a very good approximation when considering neutrino oscillations in the early Universe. This leaves an assumption of a perfectly homogeneous and isotropic plasma which is furthermore neutral, and therefore gravity is the only important long range force.

The gravitational force is described by general relativity where time and space are treated on equal footing. However, when we look at the Universe, we find that it is expanding, and this can be described by using the Friedmann-Robertson-Walker-Lemaitre metric

$$ds^2 = g_{\mu\nu} dx^\mu dx^\nu, \quad g_{\mu\nu} = \begin{bmatrix} -1 & 0 & 0 & 0 \\ 0 & a^2 & 0 & 0 \\ 0 & 0 & a^2 & 0 \\ 0 & 0 & 0 & a^2 \end{bmatrix}, \quad (2.35)$$

where a is the scale factor and summation over repeated indices is implied with μ and ν taking the values 0, 1, 2, and 3 representing the time coordinate and the three space coordinates.

This is not as such a statement about gravity and General Relativity, but rather a certain way to parametrise space-time. The gravity comes into play, when we wish to describe the relationship between energy density and space-time, and this requires the Einstein equation

$$G_{\mu\nu} \equiv R_{\mu\nu} - \frac{1}{2} g_{\mu\nu} \mathcal{R} = 8\pi G T_{\mu\nu}. \quad (2.36)$$

The left hand side of the equations is the Einstein tensor $G_{\mu\nu}$ which depends on the Ricci tensor, $R_{\mu\nu}$, the metric and the Ricci scalar, $\mathcal{R} \equiv g^{\mu\nu} R_{\mu\nu}$. On the right hand side, we have the gravitational constant, G , and the energy-momentum tensor, $T_{\mu\nu}$. The Ricci tensor, and hence also the Ricci scalar, can be computed from the metric. The 00 component of $T_{\mu\nu}$ is just the energy density, u , and computing the left hand side gives the Friedmann equation [61]

$$\left(\frac{\dot{a}}{a}\right)^2 = \frac{8\pi G u}{3}. \quad (2.37)$$

This gives the scale factor as a function of time, given that we know how u depends on a . That relationship can be derived from the covariant continuity equation

$$\frac{\partial u}{\partial t} + \frac{\dot{a}}{a}(3u + 3\mathcal{P}) = 0, \quad (2.38)$$

where \mathcal{P} is the pressure. Therefore, we need to know the relationship between energy density and pressure for the plasma where we wish to consider neutrino oscillations, and hence we also need the interesting temperature range.

The three normal neutrinos are in thermal equilibrium with the other particles at high temperatures in the early Universe, and therefore the main consequence of neutrino oscillations is to excite sterile neutrinos if they exist. This excitation has two main consequences. The first one is of course that a new particle is present in the Universe, but if the normal neutrinos are coupled to electrons and positrons when the oscillations happen, it will also lead to an over all increase in the energy density of radiation. For this reason, we are looking for the temperature range where the neutrinos decouple.

Neutrinos decouple from the electron-photon plasma when the scattering rate falls below the Hubble expansion rate, $H = \dot{a}/a$. This turns out to happen at a temperature of $\sim 1\text{MeV}$, which means that the QCD-transition has already happened, muons and pions have annihilated, while electrons and positrons are still relativistic. Hence, every species contributing significantly to the energy density of the Universe can be treated as radiation, although electrons and positrons start to annihilate just below 1MeV .

Having a Universe filled with radiation means that the energy density and pressure is related as $\mathcal{P} = \frac{1}{3}u$. As a result, Equation (2.38) can be solved by

$$u \propto a^{-4}, \quad (2.39)$$

which corresponds to the decrease in energy for the individual particle combined with the dilution of the number density as the Universe expands. With this relationship, it is possible to relate the scale factor and the time uniquely given the knowledge about how many particles contribute to the energy density.

2.3.1 Deriving a variable

The simple relationship between time and scale factor which can be found from Equation (2.37) and (2.39) means that we can choose either as our free variable when solving the QKE. The QKE is expressed as derivatives of time, but the simple transformation

$$\frac{\partial}{\partial t} = \frac{\partial a}{\partial t} \frac{\partial}{\partial a},$$

means that they are easy to transform to derivatives of a . However, as already mentioned, the energy of a particle is inversely proportional to the scale factor, and this relationship carries over to the temperature. Therefore, it is possible to get rid of both the scale factor and the time, and express the equations as derivatives of temperature. For doing this, we need to know the derivative of T with respect to t

$$\frac{\partial T}{\partial t} = \frac{\partial T}{\partial a} \frac{\partial a}{\partial t} = -\frac{T}{a} \dot{a} = -TH. \quad (2.40)$$

This formulation is convenient as we already need to know the temperature of the Universe in order to calculate several of the quantities in the QKE, while the scale factor and time never appears explicitly. Nonetheless, we still need to know the Hubble expansion rate which can be found from the Friedmann equation (2.37). Knowing the temperature, the energy density is given by

$$u = \int_0^\infty \frac{d^3\mathbf{p}}{(2\pi)^3} p f(p) = g_i \frac{\pi^2}{30} T^4, \quad (2.41)$$

2 General theoretical aspects

where $f(p)$ is the distribution function and g_i is the number of degrees of freedom for bosons, while it is the number of degrees of freedom times $7/8$ for fermions.

During the epoch just above $T = 1\text{MeV}$, the photons are the only bosons, and they contribute with one degree of freedom for each spin state. Similarly, the electrons and positrons both contribute with one degree of freedom for each helicity state, while the neutrinos only have one helicity state each. The total of the contributions are

$$g_* = 2 + \frac{7}{8}(2 + 2 + 3 + 3) = 10.75. \quad (2.42)$$

At higher temperatures, the muon, pion and eventually all the quarks should also be included, while at temperatures below $\sim 0.3\text{MeV}$, the electrons and positrons have annihilated, leading to lower values of g_* .

2.4 Numerical methods

The equations that we want to solve are formulated as ordinary differential equations, which in a quite general form can be written as

$$\frac{d\mathbf{y}}{dt} = \mathbf{G}(t, \mathbf{y}), \quad (2.43)$$

where \mathbf{y} is a vector of the variables that we wish to follow. For illustrating the numerical methods, we will only consider one variable, and the equation can be expressed by a simple function

$$\frac{dy}{dt} = g(t, y).$$

The most common methods to solve these equations numerically, called forward methods, evaluate $g(t, y(t))$ at some points and estimate $y(t + \Delta t)$ based on these points and on the value of $y(t)$. One example of this is the Euler method:

$$y(t + \Delta t) = y(t) + g(t, y(t))\Delta t. \quad (2.44)$$

For stiff ODEs these methods do, however, not perform well, and the step size Δt must be chosen to be very small. It is hard to give a precise definition of a stiff problem, but one characteristic is that the problems often involve very different timescales, and another characteristic is that the solution is very stable thus making the forward methods overshoot as it tries to compensate for every little deviation from the stable solution. A way to handle these problems is to use an implicit method. One of the simplest implicit methods is the implicit Euler:

$$y(t + \Delta t) = y(t) + g(t + \Delta t, y(t + \Delta t))\Delta t. \quad (2.45)$$

To take a step with this method it is necessary to solve a set of equations, and this makes each step more demanding computationally than for the Euler method. The advantage is that Δt can be chosen much larger with the same error tolerance and without spoiling the solver stability. In order to understand what is meant by solver stability, let us consider the simple equation

$$\frac{dy}{dt} = \lambda y. \quad (2.46)$$

The analytic solution to this equation is a simple exponential function $y(t) = a_1 e^{\lambda t}$, but solving the equation numerically highlights the general problem of stiff equations. For $\text{Re}(\lambda) < 0$, the solution should converge quickly towards zero for large t , but this is not always the case. If we consider a fixed step size Δt , the Euler method gives

$$y(t + \Delta t) = y(t)(1 + \lambda \Delta t).$$

Iterating this procedure, it is obvious that any value of λ with a real part smaller than $-1/\Delta t$ will lead to a diverging sequence, very different from the exponentially decaying analytic solution. On the other hand, the implicit Euler method gives the equation $y(t + \Delta t) = y(t) + \lambda y(t + \Delta t) \Delta t$ which can be solved to give

$$y(t + \Delta t) = y(t)/(1 - \lambda \Delta t).$$

This expression can be iterated indefinitely without diverging for any λ with a negative real part, and therefore the implicit Euler method is said to be a stable solver. This concept can be refined further by introducing different kinds of stability criteria, but we will not discuss it further here.

While the implicit Euler method is stable, it is also only a first order method which means that the step size needs to be small in order to keep down the error in every step. For the solution of the QKE, we instead use a numerical differentiation formulae of order 1-5 (**ndf15**) devised by Shampine and Reichelt [62], which uses higher order estimates, but is still stable in the sense of solving Equation (2.46) reliably if the order is reduced to 2. Another method which we will use in some cases and as a control is the **RADAU5** method. It is based on an implicit Runge-Kutta method, and is stable even though it is a fifth order method which should allow it to take longer steps. The disadvantage of **RADAU5** compared to **ndf15** is that the higher order gives rise to a larger system of equations which must be solved for every step, making each step approximately five times more expensive to evaluate. Hence it is not obvious which method will perform better in a given situation, and this can only be learnt by running them. Both these higher order methods were implemented by Thomas Tram, and they are described further in his thesis [63], where he also discusses their stability in more detail.

Full collision term

The QKEs have been well established for a long time, and they have been solved in numerous cases for many different models both for the early Universe and for supernova neutrinos. However, the collision term is usually treated in a relaxation time approximation based on equilibrium distributions if it is included at all. This is in stark contrast to the Boltzmann treatment of neutrino decoupling, where the first calculations including a dynamical collision term were performed two decades ago [64, 65].

One of the commonly used approximations is to neglect Pauli blocking. This approximation on its own induce errors in the order of 10% [60] when doing a pessimistic estimate, and on top of this comes the assumption of equilibrium distributions that is necessary for the evaluation of the collision integrals before the QKEs are solved. This level of errors should be compared to the latest values from the Planck satellite, which gives a relative relativistic energy density of $N_{\text{eff}} = 3.04 \pm 0.18$ [21]. Given that future observations of especially large scale structures will give even more precise values, it is timely to improve the treatment of collisions when considering models which depend crucially on neutrino oscillations.

In the paper presented in the next section, we have solved the QKEs for one active and one sterile oscillating neutrino with the full collision term calculated at every temperature. In doing this, we attempt to reduce the number of integrals that must be performed during the QKE integration by reducing them analytically and prepare tables when possible. Apart from including all the appropriate Pauli blocking terms, we also investigate the influence of the electron mass on the calculations. The results are compared to two different approximation schemes often used in the literature, and we find that they have different forces and weaknesses. The over all conclusion is that the error in general is smaller than the naive estimate from Pauli blocking would suggest, but also that there are deviations — especially for low conversion temperatures. In order to improve on this, we present a new approximation which captures much more of the physics. Finally, our calculations confirm what was already shown using the approximate approaches: A sterile neutrino with a mass of 1eV and a mixing of order 10% becomes fully thermalised in the early Universe. Compared to the published version of the paper, the introduction has been removed as the content has been covered here and in the general introduction.

3.1 \wp Active–sterile neutrino oscillations in the early Universe with full collision terms¹

Steen Hannestad, Rasmus Sloth Hansen, Thomas Tram and Yvonne Y. Y. Wong [5]

Sterile neutrinos are thermalised in the early Universe via oscillations with the active neutrinos for certain mixing parameters. The most detailed calculation of this thermalisation process involves the solution of the momentum-dependent quantum kinetic equations, which track the evolution of the neutrino phase space distributions. Until now the collision terms in the quantum kinetic equations have always been approximated using equilibrium distributions, but this approximation has never been checked numerically. In this work we revisit the sterile neutrino thermalisation calculation using the full collision term, and compare the results with various existing approximations in the literature. We find a better agreement than would naively be expected, but also identify some issues with these approximations that have not been appreciated previously. These include an unphysical production of neutrinos via scattering and the importance of redistributing momentum through scattering, as well as details of Pauli blocking. Finally, we devise a new approximation scheme, which improves upon some of the shortcomings of previous schemes.

3.1.1 Quantum kinetic equations

We consider oscillations between an active neutrino flavour ν_α , where $\alpha = e, \mu$ or τ , and a sterile flavour ν_s in an ensemble of neutrinos in the early universe. The density matrices $\rho(k)$ encode the flavour content and coherence of the ensemble, and are conveniently expressed in terms of polarisation vectors $(P_0(k), \mathbf{P}(k))$, i.e.,

$$\rho(k) = \begin{pmatrix} \rho_{\alpha\alpha}(k) & \rho_{\alpha s}(k) \\ \rho_{s\alpha}(k) & \rho_{ss}(k) \end{pmatrix} = \frac{1}{2}f_0(k)[P_0(k)\mathbb{1} + \mathbf{P}(k) \cdot \boldsymbol{\sigma}],$$

where $f_0(k)$ denotes a reference phase space density, which we take here to be the relativistic Fermi–Dirac distribution with vanishing chemical potential,² and $\boldsymbol{\sigma} = \{\sigma_x, \sigma_y, \sigma_z\}$ are the Pauli matrices. In this convention the active and sterile neutrino phase space distributions are given, respectively, by

$$\begin{aligned} f_{\nu_\alpha}(k) &= \frac{1}{2}[P_0(k) + P_z(k)], \\ f_{\nu_s}(k) &= \frac{1}{2}[P_0(k) - P_z(k)], \end{aligned} \tag{3.1}$$

and the QKEs that govern their evolution can be written as [45, 60]

$$\begin{aligned} \dot{\mathbf{P}}(k) &= \mathbf{V}(k) \times \mathbf{P}(k) + \frac{R_\alpha(k)}{f_0(k)}\hat{\mathbf{z}} - D(k)\mathbf{P}_T(k), \\ \dot{P}_0(k) &= \frac{R_\alpha(k)}{f_0(k)}. \end{aligned} \tag{3.2}$$

¹This article was first published in Journal of Cosmology and Astroparticle Physics by IOP Publishing and SISSA, which maintains the Version of Record

²Henceforth we shall reserve f_0 to denote exclusively the relativistic Fermi–Dirac distribution with zero chemical potential, and use f_{eq} to indicate an equilibrium distribution of a nonspecific form.

Here, $\mathbf{V}(k) \equiv V_x(k)\hat{\mathbf{x}} + V_z(k)\hat{\mathbf{z}}$ contains the vacuum oscillation term as well as the matter potential from forward scattering, and the components are given respectively by

$$\begin{aligned} V_x(k) &= \frac{\delta m^2}{2k} \sin 2\theta, \\ V_z(k) &= -\frac{\delta m^2}{2k} \cos 2\theta - \frac{7\pi^2 G_F}{45\sqrt{2}M_Z^2} kT^4 (n_{\nu_\alpha} + n_{\bar{\nu}_\alpha}) g_\alpha, \end{aligned} \quad (3.3)$$

where G_F is the Fermi constant, M_Z the mass of the Z boson, n_i the number density normalised to the equilibrium value, $g_{\mu/\tau} = 1$, and $g_e = 1 + 4 \sec^2 \theta_W / (n_{\nu_e} + n_{\bar{\nu}_e})$ with the Weinberg angle θ_W . Finally, $R_\alpha(k)$ and $D(k)$ are respectively the repopulation and damping terms, which we describe in more detail in Section 3.1.1.1. As we assume a vanishing lepton asymmetry, the same QKEs apply to both neutrinos and antineutrinos, and $n_{\nu_\alpha} = n_{\bar{\nu}_\alpha}$.

3.1.1.1 Repopulation and damping

The repopulation and damping terms are integrals over the matrix elements for annihilation and elastic scattering processes. Beginning with equation (24) of [45], which includes Pauli blocking and appears also in [59], we find these integrals to be

$$\begin{aligned} R_\alpha(k) &= 2\pi \int d\Pi_{k'} d\Pi_{p'} d\Pi_p \delta_E(kp|k'p') \\ &\quad \times \sum_i \mathcal{V}^2[\nu_\alpha(k), \bar{\nu}_\alpha(p)|i(k'), \bar{i}(p')] [f_i(E_{k'}) f_{\bar{i}}(E_{p'}) (1 - f_{\nu_\alpha}(k))(1 - f_{\bar{\nu}_\alpha}(p)) \\ &\quad \quad \quad - f_{\nu_\alpha}(k) f_{\bar{\nu}_\alpha}(p) (1 - f_i(E_{k'}))(1 - f_{\bar{i}}(E_{p'}))] \\ &\quad + \sum_j \mathcal{V}^2[\nu_\alpha(k), j(p)|\nu_\alpha(k'), j(p')] [f_{\nu_\alpha}(k') f_j(E_{p'}) (1 - f_{\nu_\alpha}(k))(1 - f_j(E_p)) \\ &\quad \quad \quad - f_{\nu_\alpha}(k) f_j(E_p) (1 - f_{\nu_\alpha}(k'))(1 - f_j(E_{p'}))], \end{aligned} \quad (3.4)$$

and

$$\begin{aligned} D(k) &= \pi \int d\Pi_{k'} d\Pi_{p'} d\Pi_p \delta_E(kp|k'p') \\ &\quad \times \sum_i \mathcal{V}^2[\nu_\alpha(k), \bar{\nu}_\alpha(p)|i(k'), \bar{i}(p')] [f_i(E_{k'}) f_{\bar{i}}(E_{p'}) (1 - f_{\bar{\nu}_\alpha}(p)) \\ &\quad \quad \quad + f_{\bar{\nu}_\alpha}(p) (1 - f_{\bar{i}}(E_{p'}))(1 - f_i(E_{k'}))] \\ &\quad + \sum_j \mathcal{V}^2[\nu_\alpha(k), j(p)|\nu_\alpha(k'), j(p')] [f_{\nu_\alpha}(k') f_j(E_{p'}) (1 - f_j(E_p)) \\ &\quad \quad \quad + f_j(E_p) (1 - f_j(E_{p'}))(1 - f_{\nu_\alpha}(k'))], \end{aligned} \quad (3.5)$$

where we have used the shorthand $d\Pi_p \equiv d^3\mathbf{p}/(2\pi)^3$, $\delta_E(kp|k'p') \equiv \delta^{(1)}(E_k + E_p - E_{k'} - E_{p'})$ is the 1D Dirac delta function, the summation index i runs over all spectator neutrino flavours (i.e., ν_β where $\beta \neq \alpha$) and the electron, while j runs in addition over all their antiparticles as well as the oscillating neutrino and antineutrino. The terms \mathcal{V}^2 are

$$\begin{aligned} \mathcal{V}^2[a(p), b(k)|c(p'), d(k')] &= (2\pi)^3 \delta^{(3)}(k + p, k' + p') N_a^2 N_b^2 N_c^2 N_d^2 \\ &\quad \times S|M|^2(a(p), b(k)|c(p'), d(k')), \end{aligned} \quad (3.6)$$

3 Full collision term

where $N_i = \sqrt{1/2E_i}$,³ E_i denotes the energy of particle i , and the squared matrix element for the forward process $a(p)b(k) \rightarrow c(p')d(k')$, summed (but not averaged) over initial and final spins, and symmetrised over identical particles in the initial and the final state is $S|M|^2(a(p), b(k)|c(p'), d(k'))$. If two ν_α s are present in the initial state, then $S|M|^2$ must additionally be multiplied by 2 to account for the fact that $\nu_\alpha(k)\nu_\alpha(p) \rightarrow \dots$ and $\nu_\alpha(p)\nu_\alpha(k) \rightarrow \dots$ constitute two identical processes.

The first part of both the repopulation and the damping integrals (3.4) and (3.5) pertains to annihilation processes, while the rest describes scattering processes. The repopulation integral (3.4) incorporates Pauli blocking in the form of additional multiplicative factors of the form $(1 - f_j)$ for every particle j in the final state of both the forward and reverse processes, and conforms with expectations. For the damping integral (3.5), however, Pauli blocking enters in a way that may not be entirely intuitive. Compared to the expression used by McKellar and Thomson [45],

$$D(k) = \pi \int d\Pi_{k'} d\Pi_{p'} d\Pi_p \delta_E(kp|k'p') \\ \times \sum_i \mathcal{V}^2[\nu_\alpha(k), \bar{\nu}_\alpha(p)|i(k'), \bar{i}(p')] f_{\bar{\nu}_\alpha}(p) + \sum_j \mathcal{V}^2[\nu_\alpha(k), j(p)|\nu_\alpha(k'), j(p')] f_j(E_p), \quad (3.7)$$

we find two modifications for each interaction process: one additional term (the ‘‘first term’’) and new multiplicative factors in the second term. In terms of the evolution of the density matrix ρ , Pauli blocking enters the equation of motion as a multiplicative matrix factor of the form $\delta_{ij} - \rho_{ij}$ (see equation (24) of [45]). The appearance of the first term can be traced to the off-diagonal part of this matrix factor, while the second term includes a factor from the matrix diagonal. Because the two corrections differ in sign, they cancel one another to some extent.

A naive introduction of Pauli blocking into the damping integral (3.5) might lead one to miss the first term, which would result in an underestimation of the damping. However, as it turns out, the negative correction contained in the second term dominates anyhow when using equilibrium distributions, so that the effect of Pauli blocking is similar to what would naively be expected, albeit smaller.

In appendix 3.1.A we evaluate the repopulation and damping integrals (3.4) and (3.5) using the technique described by Hahn-Woernle, Plümacher and Wong [66]. With this approach we can reduce the nine-dimensional integral to three dimensions. Of these it is possible to perform one integral analytically, but the remaining two must be calculated numerically.

3.1.1.2 Approximation schemes

The customary way to treat the collision terms is to assume most of the particles to be in equilibrium with the background photons. This simplifies the integrals so that solving the final expression is less numerically demanding. Assuming that Pauli blocking can be neglected and that all species follow equilibrium distributions except for the one being repopulated, the repopulation and damping terms evaluate in what we shall call

³The prefactor $N_i = \sqrt{1/2E_i}$ used here differs from the definition in [45] due to the normalisation of the Dirac spinor. Our choice of normalisation gives the completeness relations $u(p_i)\bar{u}(p_i) = \not{p}_i + m_i$ and $v(p_i)\bar{v}(p_i) = \not{p}_i - m_i$.

the equilibrium approximation [46, 48, 60] to

$$R_{\text{eq}}(k) = \Gamma \left(f_0 - \frac{f_0}{2} (P_0 + P_z) \right), \quad (3.8)$$

$$D_{\text{eq}}(k) = \frac{1}{2}\Gamma, \quad (3.9)$$

where $\Gamma = C_\alpha G_F^2 x T^5$, with $x = k/T$, $C_e \approx 1.27$, and $C_{\mu,\tau} \approx 0.92$ [46, 58, 67]. While this expression makes intuitive sense in terms of bringing the distribution towards equilibrium and coincides with the relaxation time approximation commonly found in the Boltzmann transport literature, it can also be derived from a firmer basis [45, 60].

An alternative approximation scheme is that introduced by Chu and Cirelli [68] (CC approximation), which is based on a second quantised approach [56]. Here, the combined collision (damping and repopulation) term is [55, 68]

$$\begin{aligned} \dot{\rho}_{\text{coll}}(k) = -\frac{G_F^2 T^5}{2} \frac{k}{\langle k \rangle} & \left(\left\{ \mathbf{G}_s^2, \boldsymbol{\rho} - \boldsymbol{\rho}^0 \right\} - 2\mathbf{G}_s(\boldsymbol{\rho} - \boldsymbol{\rho}^0)\mathbf{G}_s \right. \\ & \left. + \left\{ \mathbf{G}_a^2, \boldsymbol{\rho} - \boldsymbol{\rho}^0 \right\} + 2\mathbf{G}_a(\bar{\boldsymbol{\rho}} - \bar{\boldsymbol{\rho}}^0)\mathbf{G}_a \right), \end{aligned} \quad (3.10)$$

where $\mathbf{G}_{s,a} = \text{diag}(\gamma_{s,a}^\alpha, 0)$, with $(\gamma_s^e)^2 = 3.06$ and $(\gamma_s^{\mu,\tau})^2 = 2.22$ for scattering processes, and $(\gamma_a^e)^2 = 0.50$ and $(\gamma_a^{\mu,\tau})^2 = 0.28$ for annihilations [68], $\langle k \rangle \approx 3.15 T$ is the average momentum, and $\boldsymbol{\rho}^0 \equiv f_0 \mathbf{1}$. The CC approximation was originally applied in [68] to the momentum-averaged quantum rate equations, and to adapt it for the momentum-dependent quantum kinetic equations we have had to scale Equation (3.10) by a factor $k/\langle k \rangle$ [55] to approximate the momentum dependence. Evaluating the matrix products, we find

$$\dot{\rho}_{\text{coll}}(k) = -\frac{1}{2}G_F^2 T^5 \frac{k}{\langle k \rangle} \begin{pmatrix} 4(\gamma_a^\alpha)^2(\rho_{\alpha\alpha} - \rho_{\alpha\alpha}^0) & [(\gamma_s^\alpha)^2 + (\gamma_a^\alpha)^2]\rho_{\alpha s} \\ [(\gamma_s^{\alpha 2})^2 + (\gamma_a^\alpha)^2]\rho_{s\alpha} & 0 \end{pmatrix}, \quad (3.11)$$

which can be further recast as expressions for repopulation and damping,

$$R_{\text{CC}}(k) = 2G_F^2 T^5 \frac{x}{3.15} (\gamma_a^\alpha)^2 \left(f_0 - \frac{f_0}{2} (P_0 + P_z) \right), \quad (3.12)$$

$$D_{\text{CC}}(k) = \frac{1}{2}G_F^2 T^5 \frac{x}{3.15} \left[(\gamma_s^{\alpha 2})^2 + (\gamma_a^\alpha)^2 \right], \quad (3.13)$$

compatible with Equation (3.2).

Repopulation

Observe that in the CC approximation only the annihilation processes contribute to the repopulation term (3.12). This approximation is reasonable for the momentum-averaged quantum rate equations where the expression was first used, but is not strictly correct in the momentum-dependent quantum kinetic equations where elastic scattering processes is also expected to contribute to the equilibration of individual momentum bins. The equilibrium approximation (3.8) on the other hand does take into account scattering

⁴As we assume no asymmetry between neutrinos and antineutrinos, $\bar{\boldsymbol{\rho}} = \boldsymbol{\rho}$.

3 Full collision term

processes. However, in doing so, it also sacrifices the principle of detailed balance in that scattering processes can only redistribute momentum but not repopulate a momentum bin without reference to the population of other bins. A better solution would be to separate the two contributions, but this would require a more advanced treatment than a simple relaxation towards one equilibrium distribution.

Here, we develop a new repopulation approximation scheme that keeps the annihilation and scattering terms separate. We call this the A/S approximation. Looking first at annihilations, the full expression for $\nu_\alpha \bar{\nu}_\alpha$ annihilation into a and \bar{a} is an integral over $(f_a f_{\bar{a}} - f_{\nu_\alpha} f_{\bar{\nu}_\alpha})$, where we have ignored Pauli blocking factors. In the equilibrium approximation, it is assumed that $f_i \approx f_0$ for all particles $i = a, \bar{a}, \bar{\nu}_\alpha$ but ν_α . For a and \bar{a} this is a good assumption. For $\bar{\nu}_\alpha$, however, we know that it overestimates $f_{\bar{\nu}_\alpha}$ if $n_{\bar{\nu}_\alpha} < 1$. Thus, to compensate for the depletion of $\bar{\nu}_\alpha$, we adopt $f_{\bar{\nu}_\alpha} \approx n_{\nu_\alpha} f_0$ (remembering that $n_{\bar{\nu}_\alpha} = n_{\nu_\alpha}$) in the annihilation part of the A/S approximation, while keeping $f_i = f_0$ for $i = a, \bar{a}$.

For the scattering processes on the other hand, we must take care to preserve the neutrino number density at all times. We accomplish this in two ways, depending on whether the scattering process involves energy and momentum exchange between the ν_α population and an external bath.

1. *Nonzero energy and momentum exchange.* Scattering processes involving electrons, positrons and spectator neutrinos $\bar{\nu}_\beta^{(-)}$, all of which are assumed to be in equilibrium with the photons, serve to drive the ν_α population to a thermal distribution with temperature equal to the photon temperature T . We therefore model the process using a repopulation term of the relaxation form (3.8), but with the equilibrium distribution to which f_{ν_α} relaxes replaced by

$$f_{\text{eq, scat}} = \frac{1}{e^{x-\xi} + 1}, \quad (3.14)$$

where $\mu = \xi T$ is a pseudo-chemical potential.⁵ The ξ parameter can be determined in practice from the condition that the scattering contribution to \dot{n}_{ν_α} must be zero, or equivalently, that the third (not the second) kinetic moments of $f_{\text{eq, scat}}$ and f_{ν_α} must be equal (because the collision rate is proportional to momentum).

2. *Vanishing energy and momentum exchange.* Scattering amongst ν_α and $\bar{\nu}_\alpha$ conserves both the energy and number densities of the active oscillating neutrinos. Energy conservation implies that the ν_α population could relax to a thermal distribution with a temperature T_{ν_α} different from the photon temperature T , i.e.,

$$f_{\text{eq, } \nu_\alpha} = \frac{1}{e^{k/T_{\nu_\alpha} - \mu_{\nu_\alpha}/T_{\nu_\alpha}} + 1}. \quad (3.15)$$

Here, μ_{ν_α} and T_{ν_α} are determined from the combined condition that the scattering contributions to \dot{n}_{ν_α} and \dot{N}_{ν_α} must both be zero, where N_{ν_α} is the energy density normalised to the equilibrium value. The former constraint is equivalent to requiring equality between the third kinetic moments of $f_{\text{eq, } \nu_\alpha}$ and f_{ν_α} , while the latter constraint calls for equality between the fourth moments.

⁵A pseudo-chemical potential appears with the same sign for both particles and anti-particles, whereas the normal chemical potential has opposite signs for particles and anti-particles.

Thus, the full repopulation term in the A/S approximation can now be expressed as

$$R_{A/S}(k) = G_F^2 x T^5 \left\{ C_{\alpha,a} \left(f_0 - \frac{n_{\nu_\alpha} f_0}{2} (P_0 + P_z) \right) + C_{\alpha,s} \left(f_{\text{eq, scat}} - \frac{f_0}{2} (P_0 + P_z) \right) + C_{\alpha,\nu} n_{\nu_\alpha} \left(f_{\text{eq}, \nu_\alpha} - \frac{f_0}{2} (P_0 + P_z) \right) \right\}, \quad (3.16)$$

where $C_{e,a} = 0.180$, $C_{e,s} = 0.718$, $C_{e,\nu} = 0.407$, $C_{\mu/\tau,a} = 0.102$, $C_{\mu/\tau,s} = 0.407$, and $C_{\mu/\tau,\nu} = 0.407$. Full expressions for the coefficients can be found in appendix 3.1.B. As it turns out, the separation of the scattering processes into vanishing and non-vanishing energy exchange with an external bath is not crucial for the accuracy of the approximation; it has been included here mainly for consistency. If it were to be abandoned, one could simply truncate Equation (3.16) at the second term, and obtain a new coefficient for the scattering term by adding together the numerical values of $C_{\alpha,s}$ and $C_{\alpha,\nu}$ given above.

Finally, note that Equation (3.16) does not accommodate a large lepton asymmetry; if one were present, it would be necessary to reexamine the assumptions behind the equilibrium approximation (3.8) and all subsequent approximations that lead to (3.16).

Damping

The damping term in Equation (3.5) is affected by Pauli blocking in two ways as already discussed in Section 3.1.1.1. As the negative correction dominates when considering equilibrium distributions, $D_{\text{eq}}(k)$ in the equilibrium approximation (Equation (3.9)) tends to overestimate the amount of damping. Conversely, Chu and Cirelli [68] included only the diagonal Pauli blocking terms when calculating the numerical coefficients for $D_{\text{CC}}(k)$ in Equation (3.13), thereby underestimating the damping.⁶

Here, we propose to evaluate the damping integral in Equation (3.5) again with the approximations $f_i \approx f_0$ for $i \neq \nu_\alpha, \bar{\nu}_\alpha$ and $f_{\nu_\alpha} = f_{\bar{\nu}_\alpha} \approx n_{\nu_\alpha} f_0$. Due to Pauli blocking, these approximations lead to a damping term that is quadratic in n_{ν_α} :

$$D_{A/S} = \frac{1}{2} G_F^2 x T^5 \left(n_{\nu_\alpha}^2 C_{\alpha,2} + n_{\nu_\alpha} C_{\alpha,1} + C_{\alpha,0} \right), \quad (3.17)$$

where $C_{e,2} = -0.020$, $C_{e,1} = 0.569$, $C_{e,0} = 0.692$, $C_{\mu/\tau,2} = -0.020$, $C_{\mu/\tau,1} = 0.499$, and $C_{\mu/\tau,0} = 0.392$, the negative $C_{\alpha,2}$ values reflecting the origin of the $n_{\nu_\alpha}^2$ term in the negative part of the Pauli blocking factors (expressions for the coefficients are given in appendix 3.1.B). For convenience we shall continue to call this the A/S approximation, although, unlike in the repopulation treatment, there is no strict separation between annihilation and scattering; the coefficient $C_{\alpha,1}$ is predominantly from annihilation, while $C_{\alpha,0}$ comes mainly from scattering.

Lastly, we note that it is also possible to capture some of the k -dependence of the coefficients by omitting the k -integral in the computation of the damping coefficients

⁶The collision terms of the CC approximation have been presented in [68] in their integrated form, without details of how exactly they have been computed. However, reverse engineering suggests that they arise from the integral $\pi \int d\Pi_k d\Pi_{k'} d\Pi_{p'} d\Pi_p \delta_E(kp|k'p') \sum_i \mathcal{V}^2[\nu_\alpha(k), \bar{\nu}_\alpha(p)|i(k'), \bar{i}(p')] f_{\text{eq}}(p) f_{\text{eq}}(k) (1 - f_{\text{eq}}(p')) (1 - f_{\text{eq}}(k'))$ for the annihilations, and $\pi \int d\Pi_k d\Pi_{k'} d\Pi_{p'} d\Pi_p \delta_E(kp|k'p') \sum_j \mathcal{V}^2[\nu_\alpha(k), j(p)|\nu_\alpha(k'), j(p')] f_{\text{eq}}(p) f_{\text{eq}}(k) (1 - f_{\text{eq}}(p')) (1 - f_{\text{eq}}(k'))$ for the scatterings, both normalised by $\int d\Pi_k f_0(k)$.

3 Full collision term

(see appendix 3.1.B), and instead fit the result with a function of the form

$$C_{\alpha,\text{fit}} = a + be^{-cx} - \frac{d}{x+g}, \quad (3.18)$$

where a , b , c , d , and g are constants determined by the fit. This kind of expression improves the agreement between the approximation and the full treatment insofar as it reproduces the sterile neutrino spectrum with greater success than the constant coefficients. It does not however lead to significant changes in ΔN_{eff} and in the active spectrum which are often the most interesting quantities. For this reason we do not incorporate $C_{\alpha,\text{fit}}$ in the A/S approximation.

3.1.1.3 Numerical implementation

We employ a modified version of the public code `LASAGNA`⁷ [1] to solve the QKEs assuming first the full collision term (3.4) and (3.5), and then in the various approximation schemes discussed above. The QKEs are solved on a fixed momentum grid, with the explicit requirement that $\dot{\mathbf{P}} - \dot{\bar{\mathbf{P}}} = 0$ so as to enforce a zero lepton asymmetry.

Approximation schemes

Implementation of the approximate collision terms is straightforward in the equilibrium and the CC approximation. In the A/S approximation, however, additional root-finding routines are required to evaluate the chemical potential μ of the normal scattering term (3.14), as well as the temperature T_{ν_α} and chemical potential μ_{ν_α} of the $\nu_\alpha\nu_\alpha$ -scattering term (3.15).

The chemical potential of the normal scattering term satisfies the equation

$$\int d\Pi_k k \frac{1}{e^{k/T - \mu/T} + 1} = \frac{1}{2} \int d\Pi_k k f_0(P_0 + P_z),$$

which can be solved numerically. The parameter space of the $\nu_\alpha\nu_\alpha$ -scattering term, on the other hand, is two-dimensional and subject to the conditions

$$\int d\Pi_k k \frac{1}{e^{k/T_{\nu_\alpha} - \mu_{\nu_\alpha}/T_{\nu_\alpha} + 1} + 1} = \frac{1}{2} \int d\Pi_k k f_0(P_0 + P_z), \quad (3.19)$$

$$\int d\Pi_k k^2 \frac{1}{e^{k/T_{\nu_\alpha} - \mu_{\nu_\alpha}/T_{\nu_\alpha} + 1} + 1} = \frac{1}{2} \int d\Pi_k k^2 f_0(P_0 + P_z). \quad (3.20)$$

In order to solve for μ_{ν_α} and T_{ν_α} , we first reduce the problem to one dimension by constructing the ratio

$$\begin{aligned} \frac{(\int d\Pi_k k^2 f_0(P_0 + P_z))^{4/5}}{\int d\Pi_k k f_0(P_0 + P_z)} &= \frac{(\int d\Pi_k k^2 / (e^{k/T_{\nu_\alpha} - \mu_{\nu_\alpha}/T_{\nu_\alpha} + 1})^{4/5}}{\int d\Pi_k k / (e^{k/T_{\nu_\alpha} - \mu_{\nu_\alpha}/T_{\nu_\alpha} + 1})}, \\ &= (2\pi^2)^{1/5} \frac{[-24\text{Li}_5(-e^{\mu_{\nu_\alpha}/T_{\nu_\alpha}})]^{4/5}}{-6\text{Li}_4(-e^{\mu_{\nu_\alpha}/T_{\nu_\alpha}})}, \end{aligned} \quad (3.21)$$

⁷Available at https://github.com/ThomasTram/lasagna_public.

using Equation (3.19) and (3.20). Here $\text{Li}_s(z)$ denotes the polylogarithm. Since the RHS depends on $\mu_{\nu_\alpha}/T_{\nu_\alpha}$ but not directly on T_{ν_α} , Equation (3.21) can be solved immediately for $\mu_{\nu_\alpha}/T_{\nu_\alpha}$ using a one-dimensional root-finding algorithm.

Once we have established $\mu_{\nu_\alpha}/T_{\nu_\alpha}$, Equation (3.20) can be evaluated explicitly for the temperature T_{ν_α} . We find:

$$T_{\nu_\alpha} = \left(\frac{\int dx x^4 f_0(P_0 + P_z)/2}{-24\text{Li}_5\left(-e^{\frac{\mu_{\nu_\alpha}}{T_{\nu_\alpha}}}\right)} \right)^{1/5} T, \quad (3.22)$$

where T is the photon temperature, and again we have $x = k/T$. Finally, we take the μ_{ν_α} and T_{ν_α} values obtained from Equation (3.21) and (3.22), and further tune them by solving the *discretised* moments Equation (3.19) and (3.20) using a 2D Newton’s method initialised with these numbers. This last step ensures that the chosen μ_{ν_α} and T_{ν_α} values do indeed satisfy conservation of number and energy densities; the untuned values might not have this desired property because of the discretisation of momentum space. In practice, however, the amount of tuning is quite small.

Full collision term

For the full repopulation and damping terms (3.4) and (3.5), we use the following tricks to simplify the calculations. Consider first the repopulation integral, which we split into three parts:

$$R_\alpha = R_{\alpha,e} + R_{\alpha,\beta} + R_{\alpha,\alpha},$$

where the second subscript on the RHS labels the contributing processes. We use equilibrium distributions f_{eq} for the electrons and positrons in the processes

$$\nu_\alpha(k)e^\pm(p) \rightleftharpoons \nu_\alpha(k')e^\pm(p'), \quad (3.23)$$

$$\nu_\alpha(k)\bar{\nu}_\alpha(p) \rightleftharpoons e^-(k')e^+(p'), \quad (3.24)$$

which should be a good assumption as these particles are tightly coupled via electromagnetic interactions. This assumption enables the pre-evaluation (as opposed to real-time evaluation while solving the QKEs) of one of the two energy integrals in each of Equation (3.47), (3.48) and (3.50), so that the contribution of processes in Equation (3.23) and (3.24) to R_α can be expressed as

$$R_{\alpha,e}(k) = (1 - f(k)) \left(\int dE_{k'} R_{\alpha,e,s,1} f(k') + \int dE_p R_{\alpha,e,a,1} (1 - f(p)) \right) - f(k) \left(\int dE_{k'} R_{\alpha,e,s,2} (1 - f(k')) + \int dE_p f(p) R_{\alpha,e,a,2} \right), \quad (3.25)$$

with

$$\begin{aligned} R_{\alpha,e,s,1}(k, k') &= \int dE_p (\tilde{R}_{\alpha,s,e^-} + \tilde{R}_{\alpha,s,e^+}) f_{\text{eq}}(E_{p'}) (1 - f_{\text{eq}}(E_p)), \\ R_{\alpha,e,a,1}(k, p) &= \int dE_{k'} \tilde{R}_{\alpha,a,e} f_{\text{eq}}(E_{k'}) f_{\text{eq}}(E_{p'}), \\ R_{\alpha,e,s,2}(k, k') &= \int dE_p (\tilde{R}_{\alpha,s,e^-} + \tilde{R}_{\alpha,s,e^+}) f_{\text{eq}}(E_p) (1 - f_{\text{eq}}(E_{p'})), \\ R_{\alpha,e,a,2}(k, p) &= \int dE_{k'} \tilde{R}_{\alpha,a,e} (1 - f_{\text{eq}}(E_{k'})) (1 - f_{\text{eq}}(E_{p'})), \end{aligned} \quad (3.26)$$

3 Full collision term

where the integration kernels $\tilde{R}_{\alpha,x}(k, k', p)$ encode the kinematics of the interaction processes, and can be read off Equation (3.47), (3.48) and (3.50). We note also that p' is fixed by energy and momentum conservation once a combination of k, k', p has been specified.

In the limit of a massless electron, $R_{\alpha,e,s,i}$ and $R_{\alpha,e,a,i}$ of Equation (3.26) scale trivially with temperature as $\propto T^4$. However, as the temperature drops below ~ 1 MeV, the massless electron approximation also becomes increasingly tenuous. Reinstating a nonzero electron mass significantly complicates the temperature dependence of $R_{\alpha,e,s,i}$ and $R_{\alpha,e,a,i}$; we handle this by pre-evaluating Equation (3.26) on a grid in (k, p, T) (or (k, k', T)), which we interpolate in real time using a 3D spline when solving the QKEs.

We use the same equilibrium approximation for the spectator neutrinos ν_β and antineutrinos $\bar{\nu}_\beta$ in the processes

$$\begin{aligned}\nu_\alpha(k)\bar{\nu}_\beta(p) &\rightleftharpoons \nu_\alpha(k')\bar{\nu}_\beta(p'), \\ \nu_\alpha(k)\bar{\nu}_\alpha(p) &\rightleftharpoons \nu_\beta(k')\bar{\nu}_\beta(p'),\end{aligned}$$

where $\alpha \neq \beta$, and $\bar{\nu}_\beta^{(-)}$ are assumed to be non-oscillating. Thus $R_{\alpha,\beta}$ and the associated kernels $R_{\alpha,\beta,s,i}$ and $R_{\alpha,\beta,a,i}$ are, save for the $e \rightarrow \beta$ relabelling, formally given by Equation (3.25) and (3.26) respectively, and the corresponding expressions for $\tilde{R}_{\alpha,x}(k, k', p)$ can be read off Equation (3.43), (3.44) and (3.49). Since the massless neutrino limit always holds in our considerations, the temperature dependence of $R_{\alpha,\beta,s,i}$ and $R_{\alpha,\beta,a,i}$ is trivial, and the integrals need only be pre-evaluated on one 2D (k, p) (or (k, k')) grid.

Finally, for those processes involving only the active oscillating neutrinos and/or antineutrinos, i.e.,

$$\nu_\alpha(k)\bar{\nu}_\alpha(p) \rightleftharpoons \nu_\alpha(k')\bar{\nu}_\alpha(p'),$$

we evaluate the full double integrals in Equation (3.45) and (3.46) without approximation, since the momentum distributions of ν_α and $\bar{\nu}_\alpha$ are *a priori* unknown quantities.

For the damping term, we see from Equation (3.4) and (3.5) that it differs from the repopulation term only in the missing factors of $f(k)$ and $1 - f(k)$. This means that the same simplifications of the repopulation integral discussed above and consequently all of the pre-evaluated integrals apply also to the damping term. For example, contribution of processes (3.23) and (3.24) to D_α can be expressed as

$$\begin{aligned}D_{\alpha,e}(k) &= \left(\int dE_{k'} R_{\alpha,e,s,1} f(k') + \int dE_p R_{\alpha,e,a,1} (1 - f(p)) \right) \\ &\quad - \left(\int dE_{k'} R_{\alpha,e,s,2} (1 - f(k')) + \int dE_p f(p) R_{\alpha,e,a,2} \right),\end{aligned}\tag{3.27}$$

where $R_{\alpha,e,s,i}$ and $R_{\alpha,e,a,i}$ are exactly the pre-evaluated quantities of Equation (3.26).

3.1.2 Numerical results

The main quantity we wish to study is the change in the energy density of neutrinos due to a sterile neutrino population,

$$\Delta N_{\text{eff}} = N_{\text{eff}, \alpha} + N_{\text{eff}, s} - 1 = \frac{120}{7\pi^4 T^4} \int dk (f_{\nu_\alpha}(k) + f_{\nu_s}(k)) k^3 - 1.\tag{3.28}$$

This one quantity affects directly the Hubble expansion rate, making it possible to constrain ΔN_{eff} using various cosmological observations.

3.1 \mathfrak{P} Active–sterile oscillations in the early Universe with full collision terms




Figure 3.1: Convergence test using the full collision term. We compare the difference in the neutrino number and energy densities between using 200 momentum bins and 40 (blue long dashed), 80 (red dot-dashed), 100 (green dashed) and 150 bins (cyan dotted).

3.1.2.1 Numerical convergence

An important concern in the solution of the QKEs is detailed balance. If detailed balance is not fulfilled, at least so to a very good approximation, it will lead to unphysical excitation of the sterile neutrinos. As discussed in Section 3.1.1.2, detailed balance depends on the implementation of and approximations applied to the repopulation term; it is violated, for example, by the equilibrium approximation already at the level of the equation.

The full repopulation term, even in the presence of simplifications introduced in Section 3.1.1.3, preserves detailed balance in principle. In practice, however, numerical solution of the QKEs using the full collision term requires that we sum over a set of discretised momentum bins at every time step. This discretisation can potentially violate detailed balance, unless the number of momentum bins is sufficiently large. In Figure 3.1, we solve the QKEs for the normalised neutrino number density n_ν and effective energy density N_{eff} using the benchmark mixing parameter values $\delta m^2 = 0.1 \text{ eV}^2$ and $\sin^2 2\theta = 0.025$, employing variously 40, 80, 100 and 150 bins; the outcomes are expressed relative to the 200 bin results. Beyond 80 bins, convergence towards the 200 bin results to better than 0.002 across the whole temperature range of interest is immediately manifest.

Comparing the different bin choices, the offset at high temperatures originates simply in the improved representation of the distribution function as we increase the number of bins. The additional discrepancy at $T \lesssim 10 \text{ eV}$ is likely to have arisen from a very small deviation from detailed balance, since this is temperature regime at

3 Full collision term




Figure 3.2: Different contributions to the collision terms and their effects on the neutrino number and energy densities relative to the no-oscillations case. The red solid and dot-dash lines denote a scenario with only scattering and no annihilation. The blue dashed and dotted lines represent scattering only amongst the oscillating active neutrinos ν_α ; the thick black dashed and dotted lines denote the same scenario, but now computed using 150 momentum bins instead of the canonical 100.

which the thermalisation process is most efficient. For the remainder of the analysis we use 100 bins which gives an absolute error of ~ 0.001 for the benchmark mixing parameter values. This choice is a compromise between accuracy and speed, as the evaluation of the collision integrals scales with the number of momentum bins cubed: higher resolutions rapidly become prohibitively expensive in terms of CPU time.

The full collision term sources from a variety of scattering and annihilation processes of the oscillating active neutrinos with electrons, positrons, spectator active neutrinos, and the oscillating active neutrinos themselves. Since we have computed their individual contributions explicitly, we can now also study their individual effects on the sterile neutrino thermalisation process. This is a sanity check, but also serves to gauge the level to which our implementation of repopulation conserves energy and number densities and hence fulfils detailed balance. To this end, we consider in Figure 3.2 two scenarios without annihilation, and compute the corresponding changes in the neutrino energy and number densities relative to the no-oscillation case for the benchmark mixing parameter values.

In the first scenario, we include all scattering processes (red solid and dot-dash lines) and find that the neutrino number density is conserved to an excellent degree, while the energy density drops below that of the no-oscillation case. This decrease can be understood as follows. Energy is removed from the oscillating active neutrino ν_α population through conversion to sterile neutrinos beginning at the low end of

the momentum distribution. Some of these low-momentum active states are refilled from higher-momentum states through ν_α scattering with electrons, positrons, and the spectator neutrinos. This effectively reduces the total energy contained in the combined active and sterile neutrino population, where the surplus energy has been absorbed into the background of e^+ , e^- , and $\bar{\nu}_\beta$.

The second scenario (blue dashed and dotted lines) consists of only scattering processes amongst the active oscillating neutrinos themselves. Number density conservation is again satisfied to a good degree. Energy density is likewise approximately conserved; this is expected, as the isolated ν_α population has no contact with the background plasma with which to exchange energy. Observe that the degree of non-conservation is a larger here than in the first scenario. This is because the evaluation of the $R_{\alpha,\alpha}$ collision integrals in Equation (3.45) and (3.46) for the $\nu_\alpha\nu_\alpha$ processes are more sensitive to numerical inaccuracies owing to their nonlinear dependence on the distribution function $f_{\bar{\nu}_\alpha}$ (see Section 3.1.1.3). Increasing the number of momentum bins from the canonical 100 to 150 bins (thick black dashed and dotted lines) improves the situation, although it is clear that we still do not have perfect detailed balance. Nonetheless, the induced error is small enough that we can conclude it is under control.

3.1.2.2 Comparison of approximation schemes

The different approximation schemes have different strengths and weaknesses as already discussed in Section 3.1.1.2. On the one hand, the equilibrium approximation has the advantage that the assumptions involved have been analysed quite extensively [60]. It also allows the scattering processes to push the distributions towards the equilibrium form as they should, albeit at the cost of sacrificing detailed balance. The CC approximation on the other hand enforces detailed balance as it prevents the same scattering processes from repopulating the active neutrino states. However, in doing so, the approximation also neglects possible refilling due to redistribution between different momentum states. With these considerations in mind, we expect the equilibrium approximation to overestimate the neutrino number and energy densities, and the CC approximation to underestimate them.

These trends are reinforced and possibly enhanced by the behaviour of the damping term in the two approximation schemes. By neglecting Pauli blocking, the equilibrium approximation overestimates the damping, while the CC approximation underestimates it through a missing positive Pauli blocking term discussed in Section 3.1.1.2. The size of the damping term has direct consequences for ΔN_{eff} as the production of the sterile neutrinos in our case is non-resonant, where the effective production rate,

$$\Gamma = \frac{1}{2} \sin^2(2\theta_m) \Gamma_{\text{collision}}, \quad (3.29)$$

is directly proportional to the damping term $\Gamma_{\text{collision}} \sim D$ and the in-matter mixing angle θ_m [69]. Thus, the larger the damping, the higher the resulting ΔN_{eff} , and vice versa.

As shown in Figure 3.3, the over- and underestimation of Δn_ν and ΔN_{eff} in the equilibrium and CC approximations, respectively, are exactly what transpires in our numerical solutions of the QKEs for the benchmark mixing parameter values. In contrast, the A/S approximation introduced in this work takes the best of both worlds,

3 Full collision term

and, as is evident in Figure 3.3, comes much closer to the full collision treatment than both the equilibrium and the CC approximations. It is not obvious what to expect for the A/S approximation, and the fact that the resulting ΔN_{eff} and Δn_ν values tend towards different sides of the full solutions suggests that there is no major systematic bias.

3.1.2.3 Electron mass and Pauli blocking

Next we test the consequences of ignoring Pauli blocking, and of assuming $m_e = 0$ in the full collision term. As shown in Figure 3.3, the latter assumption makes no practical difference to Δn_ν and ΔN_{eff} , since the conversion to sterile neutrinos takes place largely before electrons become nonrelativistic. Ignoring Pauli blocking, however, induces a ~ 0.02 absolute error as T drops below ~ 5 MeV, smaller than the naive expectation of $\sim 10\%$ estimated from the Pauli blocking terms in the relevant momentum range [60].

Note that there is a small subtlety when ignoring Pauli blocking: detailed balance can be satisfied for Fermi–Dirac distributions only when the Pauli blocking is taken into account. Otherwise, ignoring Pauli blocking generally demands that we replace Fermi–Dirac with Maxwell–Boltzmann statistics in order to fulfil detailed balance. We would however like to continue using Fermi–Dirac distributions, and therefore choose to enforce detailed balance in a different way. For all processes of the form $a(p) + \nu_\alpha(k) \rightleftharpoons b(p')c(k')$, we assume that

$$f(E_{p'})f(E_{k'}) - f(E_p)f(k) = f_{\text{eq}}(E_p)(f_{\text{eq}}(k) - f(k)). \quad (3.30)$$

This is also the assumption behind the equilibrium approximation, and is exact if $f(E_{p'})$, $f(E_{k'})$ and $f(E_p)$ take on the Maxwell–Boltzmann equilibrium form. Not surprisingly we see in Figure 3.3 that the equilibrium approximation solution follows the no Pauli blocking full solution quite well; the difference between them is due mainly to rounding errors in the numerical value of C_α in the equilibrium approximation immediately below Equation (3.9).

3.1.2.4 Impact on distribution functions

It is also interesting to compare the active and sterile neutrino distribution functions in the different approximations. The distribution function of the electron neutrino affects directly the weak reaction rates in BBN [55], while the division of ΔN_{eff} between ν_α and ν_s have consequences for large-scale structure formation if δm^2 is significantly larger than the solar and atmospheric mass splittings [70, 71]. Figure 3.4 shows this comparison for the benchmark mixing parameters.

The first observation is that the full collision treatment gives rise to $f > f_0$ for the high-momentum active neutrinos at $T = 10$ MeV. This is not a numerical artefact, but originates in both the $\nu_\alpha\nu_\alpha$ -scattering and the annihilation processes in the presence of a step-like feature at low momenta, such as that produced by oscillations here. Consider first $\nu_\alpha\nu_\alpha$ -scattering. Removing neutrinos from the low momentum modes causes the energy per neutrino to increase. This in turn triggers the number and energy conserving scatterings to push the distribution towards an equilibrium with a higher effective temperature, which automatically leads to $f > f_0$. Annihilation processes on the other hand enhance f via a slightly different mechanism. For processes

3.1 \mathfrak{P} Active–sterile oscillations in the early Universe with full collision terms

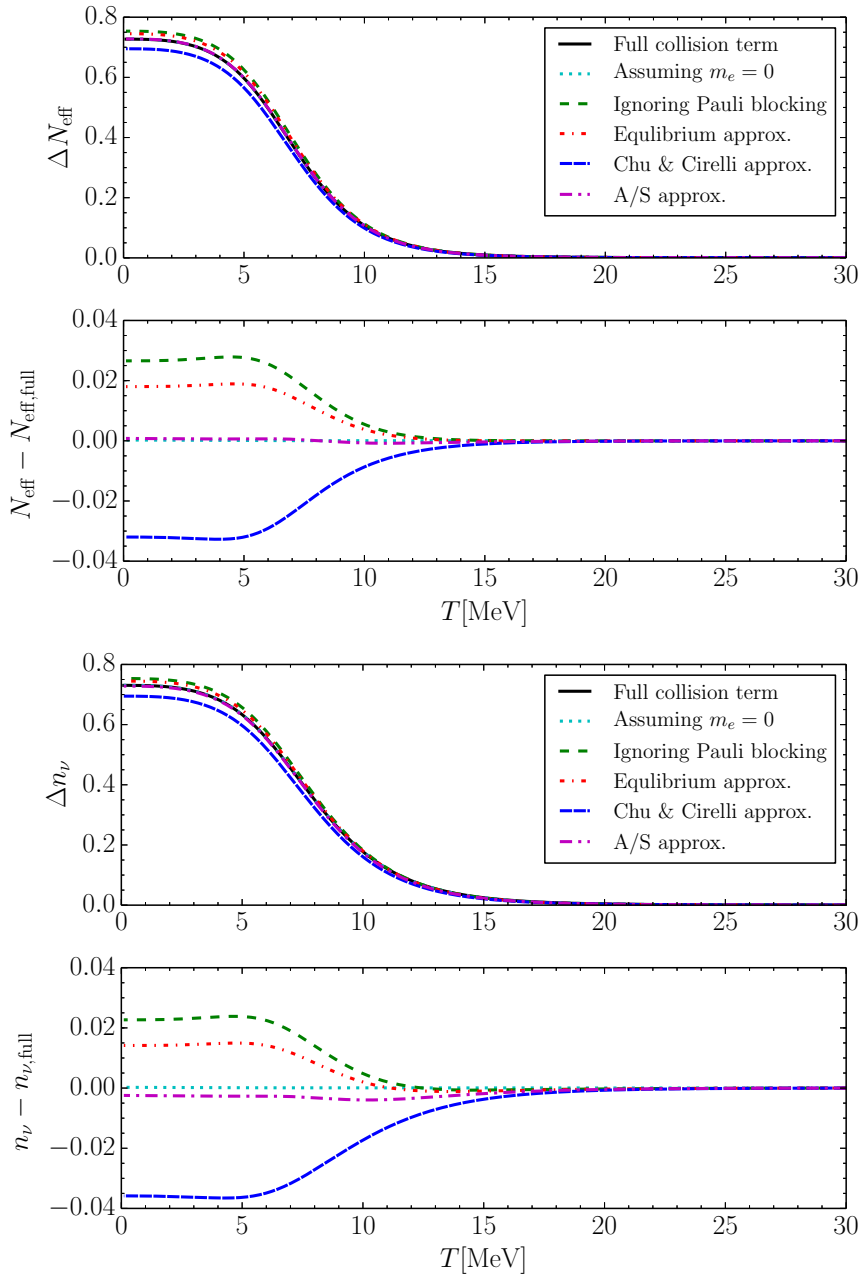


Figure 3.3: Comparison of the full treatment (black solid lines) with the equilibrium approximation (red dot-short dash), the CC approximation (blue long dashed), and the A/S approximation (purple dot-long dash) introduced in this work. Setting $m_e = 0$ (cyan dotted) has almost no effect, while ignoring Pauli blocking in the full expression (green dashed) gives almost the same result as the equilibrium approximation.

3 Full collision term

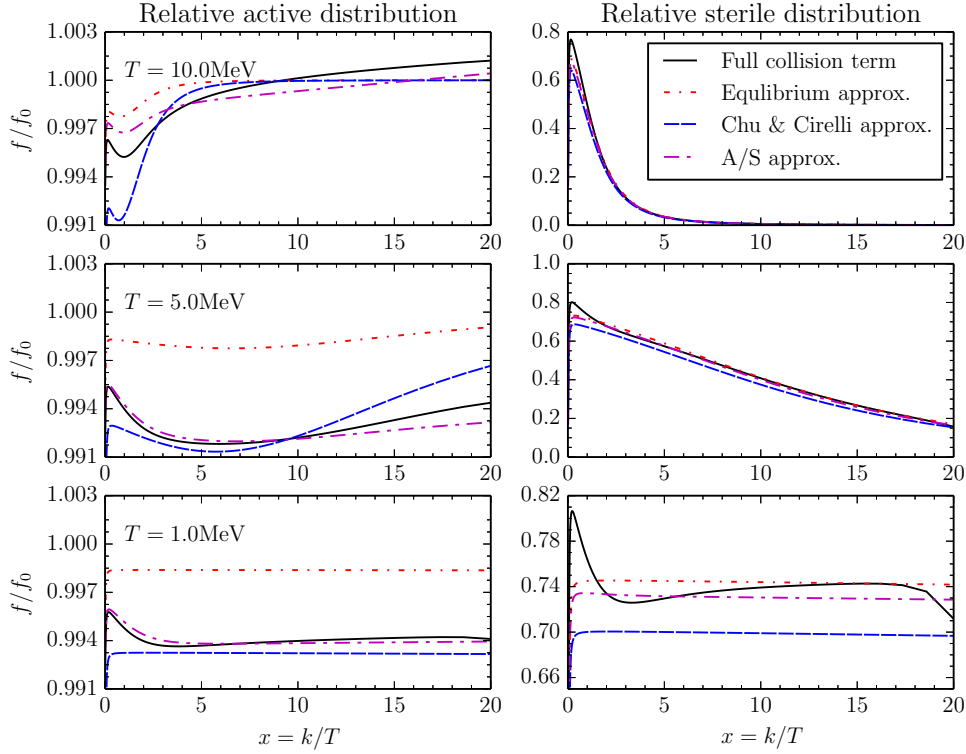


Figure 3.4: The active and sterile neutrino momentum distributions computed from the full treatment (black solid lines), the equilibrium approximation (red dot-short dash), the CC approximation (blue dashed), and the A/S approximation (purple dot-long dash) for the benchmark mixing parameters $\delta m^2 = 0.1\text{eV}^2$ and $\sin^2(2\theta) = 0.025$ at three different temperatures. The distributions have been normalised to the relativistic Fermi–Dirac distribution f_0 .

involving only states above the step, the equilibrium is maintained. For processes such as $\nu_\alpha(k_{\text{high}})\bar{\nu}_\alpha(p_{\text{low}}) \rightleftharpoons a(k')\bar{a}(p')$, however, the reaction will be pushed to the left because there is a deficit of $\bar{\nu}_\alpha(p_{\text{low}})$ relative to $a(k')$ and $\bar{a}(p')$. Thus, annihilation can also overproduce ν_α at high momenta, leading to $f > f_0$.

The A/S approximation mimics both the annihilation and the scattering effects to an extent, and is the only approximation scheme that manages to reproduce $f > f_0$ at $T = 10$ MeV, albeit only at very high momenta. The annihilation effect is captured by the n_{ν_α} factor in Equation (3.16), while the separate treatment of $\nu_\alpha^{(-)}\bar{\nu}_\alpha^{(-)} \rightleftharpoons \nu_\alpha^{(-)}\bar{\nu}_\alpha^{(-)}$ accounts for the scattering effect. The latter constitutes the main role of the $C_{\alpha,\nu}$ term in the A/S repopulation approximation (3.16); in contrast, the $C_{\alpha,s}$ term with the pseudo-chemical potential ξ as the sole degree of freedom acts in the wrong direction: it becomes negative when the active sector is depleted, giving rise to a lower value of f for any choice of momentum.

Notwithstanding its success at capturing some degree of the $f > f_0$ phenomenon, the approximate treatment of scattering remains the main source of discrepancy between

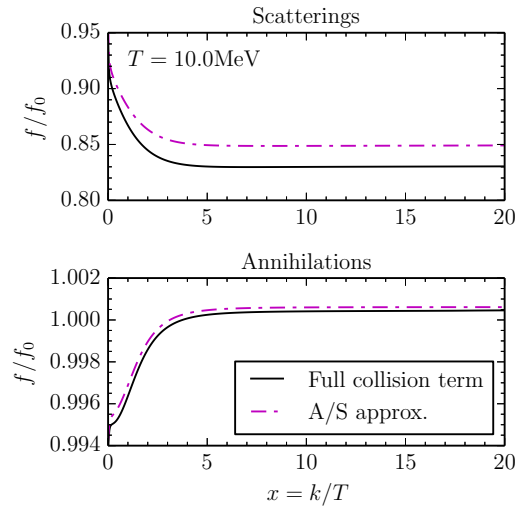


Figure 3.5: The active momentum distributions at $T = 10$ MeV calculated using a collision term incorporating only scattering processes (*top*) and only annihilations (*bottom*). The difference between the full treatment (black solid) and the A/S approximation (purple long-dot dash) is much larger for the scatterings than for the annihilations. We have used the benchmark mixing parameters $\delta m^2 = 0.1 \text{ eV}^2$ and $\sin^2(2\theta) = 0.025$, and the distributions have been normalised to the relativistic Fermi–Dirac distribution f_0 .

the full solution and the A/S approximation at $T = 10$ MeV. This is evident in figure 3.5 where we compare the distribution functions obtained assuming (i) only scattering, and (ii) only annihilation. The scattering-only result also demonstrates that these processes do not replenish the active neutrino population, yielding $f/f_0 \sim 0.85$ for most momenta, while the annihilations are very effective at bringing f/f_0 to 1.

Apart from these effects, we find that the equilibrium approximation gives too large a final distribution for both the active and the sterile neutrinos, as expected from the oversized R_{eq} and D_{eq} terms. The CC approximation, on the other hand, comes surprisingly close to the real active neutrino distribution, but is short by $\sim 5\%$ for the sterile states. As repopulation does not directly affect oscillations into sterile states, we conclude that this offset must originate in the undersized damping term D_{CC} , which as discussed in Section 3.1.2.2 affects directly the effective sterile neutrino production rate (3.29).

3.1.2.5 Dependence on mixing parameters

So far we have tested the various approximation schemes using the benchmark mixing parameter values $\delta m^2 = 0.1 \text{ eV}^2$ and $\sin^2 2\theta = 0.025$. In reality, however, the errors caused by these approximations are also sensitive to the parameter values.

Figure 3.6 shows ΔN_{eff} at $T = 0.1$ MeV as a function of δm^2 and $\sin^2 2\theta$ computed using the full collision term. Our results are generally quite similar to previous calculations using the equilibrium approximation [48, 58]. Note however that for large mixing

3 Full collision term




Figure 3.6: ΔN_{eff} at $T = 0.1$ MeV using the full collision term.

angles and small mass squared differences, the contours deviate slightly from the straight lines found in [48, 58]. To highlight this deviation, we show in Figure 3.7 the differences in ΔN_{eff} incurred respectively by the equilibrium, CC, and A/S approximations relative to the full solution in the $(\delta m^2, \sin^2 2\theta)$ -parameter space.

As expected, the deviations always occur, independently of the approximation scheme, in a diagonal band corresponding to the transition region from $\Delta N_{\text{eff}} = 0$ to $\Delta N_{\text{eff}} = 1$. Beyond this common feature, however, the different approximations incur the largest deviations at different parameter values.

The equilibrium approximation follows the result of the full calculation quite well for δm^2 values above 0.01eV^2 , but overestimates ΔN_{eff} by more than 0.1 at $\delta m^2 < 0.001 \text{eV}^2$. We can understand this deviation by looking at the conversion temperature. The temperature of maximal conversion is proportional to $(\delta m^2)^{1/6}$ [58], so that low values of δm^2 generally correspond to low conversion temperatures. If the conversion temperature is sufficiently high ($T \gg 1$ MeV), repopulation is rapid and ΔN_{eff} is limited only by how fast ν_s can be produced through oscillations and collisions [69]. The effective production rate is given by Equation (3.29), and production ceases as soon as the collision rate becomes too low. If most of the conversion occurs at low temperatures, however, collisions become too inefficient to sustain the population of the active sector and consequently for Equation (3.29) — which assumes instantaneous repopulation — to hold, thereby causing the real ΔN_{eff} contours to deviate from straight lines in relation to δm^2 and $\sin^2 2\theta$ in Figure 3.6. The equilibrium approximation errs in its overestimation of the repopulation rate, yielding almost straight ΔN_{eff} contours even in the low δm^2 , high $\sin^2 2\theta$ region.

In contrast, the top right panel of Figure 3.7 shows that the CC approximation generally underestimates ΔN_{eff} , but works somewhat better at $\delta m^2 \lesssim 0.01 \text{eV}^2$. For high δm^2 values the underestimation is due mainly to the undersized D_{CC} damping term which diminishes the sterile neutrino production rate as already discussed in

3.1 \mathfrak{P} Active–sterile oscillations in the early Universe with full collision terms




Figure 3.7: Deviations in the $\Delta N_{\text{eff}}(T = 0.1 \text{ MeV})$ values computed using various approximations from the full solutions. *Top left:* The equilibrium approximation. *Top right:* The CC approximation. *Bottom:* The A/S approximation. Note that the colour scale is linear, but the contour levels are not.

Section 3.1.2.2. For low δm^2 values, however, the agreement becomes better (a deviation of 0.02 at $\delta m^2 = 10^{-4}$, $\sin^2 2\theta = 10^{-0.5}$) because the deficiency of *sterile* neutrinos is compensated by an overproduction of active neutrinos (similar to the overpopulation of the active sector discussed above in relation to the equilibrium approximation). This overproduction comes about as follows. Although technically the CC repopulation integral incorporates only annihilation processes, numerically the constant $2(\gamma_a^e)^2/3.15 = 0.317$ is rather larger than its annihilation counterpart in the A/S approximation, $C_{e,a} = 0.177$. Thus, the CC approximation does inadvertently contain some degree of repopulation due to scattering, which drives up the active neutrino repopulation rate relative to the true rate.

Finally, the bottom panel of Figure 3.7 shows the A/S approximation. Here, the agreement is generally much better than either the equilibrium or the CC approximation, although there remains a small discrepancy of ~ 0.01 in the low δm^2 , high $\sin^2 2\theta$ corner. This is not surprising, as the region is characterised by large spectral distortions from conversion at low temperatures, and is thus most sensitive to how exactly we handle

3 Full collision term

repopulation in the active sector. It is nonetheless remarkable that for most of the parameter region, the A/S approximation is able to reproduce the full results at a precision of ~ 0.001 through a fairly simplistic description of the collision term. This deviation is in fact comparable to the numerical error expected of the full treatment due to our choice of momentum resolution (see Section 3.1.2.1). Therefore, Figure 3.7 not only validates the A/S approximation, but through comparison with a physically intuitive modelling of collisions, also confirms that the full collision term has been implemented correctly

3.1.3 Conclusions

We have calculated in this work the full collision term for active–sterile neutrino oscillations in the early universe, and for the specific case of (ν_e, ν_s) -oscillations implemented it in the computation of ΔN_{eff} from sterile neutrino thermalisation. In particular we have included for the first time a nonzero electron mass and Pauli blocking in the collision integrals. The former turns out to have a negligible impact on the final results, while ignoring the latter gives rise to noticeable discrepancies. Nonetheless, our full treatment confirms previous analyses based on approximations of the collision terms that a 1 eV sterile neutrino coupled with a mixing angle $\sin^2 2\theta \sim 0.1$ produces $\Delta N_{\text{eff}} \approx 1$; the discrepancies arising from approximations are at the level that affects only precision calculations.

We then proceeded to perform a systematic comparison of the full collision treatment with two approximation schemes found in the literature. We find that the commonly used “equilibrium approximation” [60] reproduces ΔN_{eff} at better than 0.04 for $\delta m^2 > 0.01 \text{ eV}^2$, but incurs large errors (> 0.1) for very small mass squared differences due to the unphysical repopulation of the active sector by elastic scattering. The “CC approximation” of Chu and Cirelli [68] is discrepant up to 0.04 for $\delta m^2 > 0.0001 \text{ eV}^2$, but improves for low δm^2 values because of a fortuitous cancellation between an underestimated damping term and an overestimation of repopulation from annihilation processes.

Recognising that the equilibrium and CC approximations neglect different physical effects, we have devised a new approximation scheme, the A/S approximation, in which scattering and annihilation contributions to repopulation are treated separately, and the damping term includes Pauli blocking. As the scheme is better able to capture the physics of repopulation and damping, it also pushes the error in ΔN_{eff} down to 0.001 for most of the parameter region, although minor deviations (~ 0.01) remain in the $\delta m^2 < 0.001 \text{ eV}^2$ region, where the sterile neutrino conversion temperature is lowest and hence spectral distortions from oscillations are expected to have the largest effect.

The connection between low δm^2 deviations and low temperature spectral distortions also has implications for an inverted mass spectrum, i.e., where the active state is heavier than the sterile state, as well as for the various mechanisms designed to reconcile eV-mass sterile neutrinos with cosmology. In the case of an inverted mass spectrum, sterile neutrino thermalisation proceeds via a resonance, which, depending on the adiabaticity of the resonance and hence the mixing angle, can cause more disturbance to the active neutrino spectrum than in the non-resonant case. Likewise, mechanisms that block the production of eV-mass sterile neutrinos typically work by delaying the thermalisation

to low temperatures, which by construction also makes them more sensitive to the approximations employed for the collision terms.

At the current level of observational precision, $\sigma(\Delta N_{\text{eff}}) \approx 0.2$ from Planck [21], our analysis shows that the CC and the A/S approximations can be reliably applied to most active–sterile oscillation scenarios, whereas the equilibrium approximation appears to be approaching its boundary of validity if the sterile neutrino conversion temperature is too low. In the future, large-volume galaxy surveys are expected to improve the sensitivity to ΔN_{eff} to ~ 0.03 [72]. Thus, should hints for a 1 eV sterile neutrino persist in the laboratory, a collision treatment more precise than either the equilibrium or the CC approximation can offer would become necessary. Short of evaluating the full collision terms, which is numerically costly or possibly even infeasible in some cases, the A/S approximation developed in this work appears to be a most convenient alternative.

Acknowledgments

We acknowledge use of computational resources provided by the Danish e-Infrastructure Cooperation. Y³W thanks the Institute for Nuclear Theory at the University of Washington for its hospitality and the Department of Energy for partial support during the completion of this work.

3.1.A Derivation of the full collision terms

3.1.A.1 Neutrino–electron scattering in the s -channel

Consider first the case of $\nu_\alpha + e^- \rightleftharpoons \nu_\alpha + e^-$. The repopulation term due to this process takes on the form

$$R_{\alpha,s,e} = \frac{1}{29\pi^5} E_k \int d^3\mathbf{k}' d^3\mathbf{p}' d^3\mathbf{p} \frac{1}{E_{k'} E_p E_{p'}} |M|^2 (\nu_\alpha(k), e^-(p) | \nu_\alpha(k'), e^-(p')) \\ \times \delta^{(4)}(k + p - k' - p') [f_{\nu_\alpha}(E_{k'}) f_e(E_{p'}) (1 - f_{\nu_\alpha}(E_k)) (1 - f_e(E_p)) \\ - f_{\nu_\alpha}(E_k) f_e(E_p) (1 - f_{\nu_\alpha}(E_{k'})) (1 - f_e(E_{p'}))], \quad (3.31)$$

where the matrix element is [64, 73]

$$|M|^2 = 32G_F^2 \left(A_\alpha^2(k \cdot p)(k' \cdot p') + B_\alpha^2(k \cdot p')(k' \cdot p) - m_e^2 A_\alpha B_\alpha(k \cdot k') \right), \quad (3.32)$$

with $A_e = 2 \sin^2 \theta_W + 1$, $A_{\mu,\tau} = 2 \sin^2 \theta_W - 1$, $B_e = 2 \sin^2 \theta_W$, and $B_{\mu,\tau} = 2 \sin^2 \theta_W$. Note the change of notation here in the appendix: p and \mathbf{p} now denote respectively the 4- and 3-momentum, while in the main text we use p to indicate the magnitude of the 3-momentum.

The matrix element has three different dependences on the momentum, and therefore requires three different parameterisations to simplify the integral. We label the first dependence $(k \cdot p)(k' \cdot p')$ the s -channel, the second $(k \cdot p')(k' \cdot p)$ the u -channel, and the last dependence $(k \cdot k')$ the t -channel, where for each channel it is convenient to define a 3-momentum variable, given respectively by

$$\mathbf{q} \equiv \mathbf{p} + \mathbf{k} = \mathbf{p}' + \mathbf{k}', \\ \mathbf{v} \equiv \mathbf{p} - \mathbf{k}' = \mathbf{p}' - \mathbf{k}, \\ \mathbf{w} \equiv \mathbf{k} - \mathbf{k}' = \mathbf{p}' - \mathbf{p},$$

3 Full collision term

which replaces one of the integration variables of Equation (3.31). Then, to evaluate the integral for each channel we simply follow the technique described by Hahn-Woernle, Plümacher and Wong in [66].

Taking the s -channel as a worked example, we use the 3-momentum variable \mathbf{q} as a reference and define around it a coordinate system

$$\begin{aligned}\mathbf{q} &= |\mathbf{q}|(0, 0, 1), \\ \mathbf{k} &= E_k(0, \sin \eta, \cos \eta), \\ \mathbf{k}' &= E_{k'}(\cos \varphi \sin \theta, \sin \varphi \sin \theta, \cos \theta).\end{aligned}$$

With this choice, we find the quantities

$$\begin{aligned}s &= (p + k)^2 = (p' + k')^2 = (E_p + E_k)^2 - |\mathbf{q}|^2, \\ k \cdot p &= k' \cdot p' = \frac{s - m_e^2}{2}, \\ |\mathbf{q} - \mathbf{k}'|^2 &= |\mathbf{q}|^2 + |\mathbf{k}'|^2 - 2\mathbf{q} \cdot \mathbf{k}' = |\mathbf{q}|^2 + E_{k'}^2 - 2|\mathbf{q}|E_{k'} \cos \theta, \\ |\mathbf{q} - \mathbf{k}|^2 &= |\mathbf{q}|^2 + |\mathbf{k}|^2 - 2\mathbf{q} \cdot \mathbf{k} = |\mathbf{q}|^2 + E_k^2 - 2|\mathbf{q}|E_k \cos \eta,\end{aligned}$$

and consequently

$$|M_s|^2 = 8 G_F^2 A_\alpha^2 ((E_p + E_k)^2 - |\mathbf{q}|^2 - m_e^2)^2 \quad (3.33)$$

for the first term of the matrix element in Equation (3.32).

Since the matrix element in Equation (3.33) now depends only on energies and the magnitude of \mathbf{q} , we can use the Dirac delta functions in the integral in Equation (3.31) to integrate out the directional dependencies. Evaluating first the \mathbf{p}' -integral as [66]

$$\begin{aligned}\int \frac{d^3 \mathbf{p}'}{2E_{p'}} \delta^{(4)}(k + p - k' - p') &= \int d^3 \mathbf{p}' dE_{p'} \frac{\delta(E_{p'} - \sqrt{|\mathbf{p}'|^2 + m_e^2}) \theta(E_{p'} - m_e)}{2\sqrt{|\mathbf{p}'|^2 + m_e^2}} \\ &\quad \times \delta(E_k + E_p - E_{k'} - E_{p'}) \delta^3(\mathbf{k} + \mathbf{p} - \mathbf{k}' - \mathbf{p}') \\ &= \frac{\delta(E_k + E_p - E_{k'} - \sqrt{|\mathbf{q} - \mathbf{k}'|^2 + m_e^2})}{2\sqrt{|\mathbf{q} - \mathbf{k}'|^2 + m_e^2}} \\ &\quad \times \theta(E_k + E_p - E_{k'} - m_e),\end{aligned} \quad (3.34)$$

we use the property

$$\delta(f(x)) = \sum_{f(x_i)=0} \frac{\delta(x - x_i)}{|f'(x_i)|}$$

and hence

$$\delta(E_i^2 - |\mathbf{p}_i|^2 - m_i^2) = \frac{\delta(E_i - \sqrt{|\mathbf{p}_i|^2 + m_i^2})}{2\sqrt{|\mathbf{p}_i|^2 + m_i^2}} + \frac{\delta(E_i + \sqrt{|\mathbf{p}_i|^2 + m_i^2})}{2\sqrt{|\mathbf{p}_i|^2 + m_i^2}}$$

to further simplify Equation (3.34) to [66]

$$\begin{aligned}
 & \int \frac{d^3 \mathbf{p}'}{2E_{p'}} \delta^{(4)}(k + p - k' - p') \\
 &= \delta((E_k + E_p - E_{k'})^2 - |\mathbf{q} - \mathbf{k}'|^2 - m_e^2) \theta(E_k + E_p - E_{k'} - m_e) \\
 &= \delta\left((E_k + E_p - E_{k'})^2 - |\mathbf{q}|^2 - E_{k'}^2 + 2|\mathbf{q}|E_{k'} \cos \theta - m_e^2\right) \\
 &\quad \times \theta(E_k + E_p - E_{k'} - m_e) \\
 &= \frac{1}{2|\mathbf{q}|E_{k'}} \delta\left(\cos \theta - \frac{E_{k'}^2 - (E_k + E_p - E_{k'})^2 + |\mathbf{q}|^2 + m_e^2}{2|\mathbf{q}|E_{k'}}\right) \\
 &\quad \times \theta(E_k + E_p - E_{k'} - m_e).
 \end{aligned} \tag{3.35}$$

In a similar way, we rewrite

$$\begin{aligned}
 \frac{d^3 \mathbf{p}}{2E_p} &= \int d^3 \mathbf{p} dE_p \delta(E_p^2 - |\mathbf{q} - \mathbf{k}|^2 - m_e^2) \theta(E_p - m_e) \\
 &= \int d^3 \mathbf{q} dE_p \frac{1}{2|\mathbf{q}|E_k} \delta\left(\cos \eta - \frac{E_k^2 - E_p^2 + |\mathbf{q}|^2 + m_e^2}{2|\mathbf{q}|E_k}\right) \theta(E_p - m_e),
 \end{aligned} \tag{3.36}$$

where in the second line we have also changed the integration variable from \mathbf{p} to \mathbf{q} . Then, applying Equation (3.35) and (3.36) to the integral in Equation (3.31), we find, after performing the trivial angular integrations and averaging over the direction of the incoming neutrinos ($\int d \cos \eta / 2$), the s -channel contribution

$$\begin{aligned}
 R_{\alpha,s,e,s} &= \frac{1}{27\pi^3 E_k^2} \int d \cos \eta d \cos \theta d|\mathbf{q}| dE_{k'} dE_p \delta(\cos \theta - \dots) \delta(\cos \eta - \dots) \\
 &\quad \times |M_s|^2 F \theta(E_k + E_p - E_{k'} - m_e) \theta(E_p - m_e),
 \end{aligned} \tag{3.37}$$

where F denotes all the distribution functions, and the arguments of the two Dirac delta functions can be read off Equation (3.35) and (3.36) respectively.

The integrals over $\cos \theta$ and $\cos \eta$ involve only Dirac delta functions. Taking the $\cos \eta$ -integral as an example and noting that integration limits can be equivalently expressed as step functions, we find

$$\begin{aligned}
 & \int_{-1}^1 d(\cos \eta) \delta\left(\cos \eta - \frac{E_k^2 - E_p^2 + |\mathbf{q}|^2 + m_e^2}{2|\mathbf{q}|E_k}\right) \\
 &= \theta\left(1 - \frac{E_k^2 - E_p^2 + |\mathbf{q}|^2 + m_e^2}{2|\mathbf{q}|E_k}\right) \theta\left(\frac{E_k^2 - E_p^2 + |\mathbf{q}|^2 + m_e^2}{2|\mathbf{q}|E_k} + 1\right).
 \end{aligned}$$

The two step functions can be reinterpreted as limits on $|\mathbf{q}|$, and together they confine $|\mathbf{q}|$ to the interval

$$\left| |\mathbf{k}| - |\mathbf{p}| \right| = |E_k - \sqrt{E_p^2 - m_e^2}| \leq |\mathbf{q}| \leq E_k + \sqrt{E_p^2 - m_e^2} = |\mathbf{k}| + |\mathbf{p}|.$$

Integrating over $d \cos \theta$ gives the same formal result save for the replacements $\mathbf{k} \rightarrow \mathbf{k}'$ and $\mathbf{p} \rightarrow \mathbf{p}'$. Thus the combined integration limits on $|\mathbf{q}|$ can be written equivalently as

$$\max(|\mathbf{k}| - |\mathbf{p}|, |\mathbf{k}'| - |\mathbf{p}'|) \leq |\mathbf{q}| \leq \min(|\mathbf{k}| + |\mathbf{p}|, |\mathbf{k}'| + |\mathbf{p}'|), \tag{3.38}$$

3 Full collision term

Reaction ($\alpha \neq \beta$)	$S M ^2$
$\nu_\alpha(k)\nu_\beta(p) \rightarrow \nu_\alpha(k')\nu_\beta(p')$	$32G_F^2(k \cdot p)(k' \cdot p')$
$\nu_\alpha(k)\bar{\nu}_\beta(p) \rightarrow \nu_\alpha(k')\bar{\nu}_\beta(p')$	$32G_F^2(k \cdot p')(k' \cdot p)$
$\nu_\alpha(k)\nu_\alpha(p) \rightarrow \nu_\alpha(k')\nu_\alpha(p')$	$32G_F^2 2(k \cdot p)(k' \cdot p')$
$\nu_\alpha(k)\bar{\nu}_\alpha(p) \rightarrow \nu_\alpha(k')\bar{\nu}_\alpha(p')$	$32G_F^2 4(k \cdot p')(k' \cdot p)$
$\nu_\alpha(k)e^-(p) \rightarrow \nu_\alpha(k')e^-(p')$	$32G_F^2 ((2x_W \pm 1)^2(k \cdot p)(k' \cdot p') + 4x_W^2(k \cdot p')(k' \cdot p) - (2x_W \pm 1)2x_W m_e^2(k \cdot k'))$
$\nu_\alpha(k)e^+(p) \rightarrow \nu_\alpha(k')e^+(p')$	$32G_F^2 ((2x_W \pm 1)^2(k \cdot p')(k' \cdot p) + 4x_W^2(k \cdot p)(k' \cdot p') - (2x_W \pm 1)2x_W m_e^2(k \cdot k'))$
$\nu_\alpha(k)\bar{\nu}_\alpha(p) \rightarrow \nu_\beta(k')\bar{\nu}_\beta(p')$	$32G_F^2(k \cdot p')(k' \cdot p)$
$\nu_\alpha(k)\bar{\nu}_\alpha(p) \rightarrow e^-(k')e^+(p')$	$32G_F^2 ((2x_W \pm 1)^2(k \cdot k')(p \cdot p') + 4x_W^2(k \cdot p')(k' \cdot p) - (2x_W \pm 1)2x_W m_e^2(k \cdot p))$

Table 3.1: Matrix elements for all relevant reactions involving ν_α , with $x_W = \sin^2 \theta_W = 0.23864$ [38], S is a symmetrisation factor of $1/2$ for each pair of indistinguishable particles in the initial and the final state, and $|M|^2$ has been summed but not averaged over initial and final spins. For the process with two ν_α s in the initial state, we have further multiplied the matrix element by 2 to account for the fact that $\nu_\alpha(k)\nu_\alpha(p) \rightarrow \dots$ and $\nu_\alpha(p)\nu_\alpha(k) \rightarrow \dots$ constitute two identical processes. Where there is a choice of \pm , the plus signs are for $\alpha = e$, and the minus signs for $\alpha = \mu, \tau$. The corresponding matrix elements for $\bar{\nu}_\alpha$ can be obtained by the exchange $(k \cdot p)(k' \cdot p') \leftrightarrow (k \cdot p')(k' \cdot p)$ for the elastic scattering processes.

and we note that $|\mathbf{p}'|$ is determined from $|\mathbf{p}|, |\mathbf{k}|, |\mathbf{k}'|$ by imposing energy conservation. Applying the limits in Equation (3.38) to the integral in Equation (3.37) then yields

$$R_{\alpha,s,e,s} = \frac{1}{2^7 \pi^3 E_k^2} \int_0^\infty dE_{k'} \int_{\max(m_e, k' - k + m_e)}^\infty dE_p \int d|\mathbf{q}| |M_s|^2 F \times \theta(|\mathbf{q}| - \max(|\mathbf{k}| - |\mathbf{p}|, |\mathbf{k}'| - |\mathbf{p}'|)) \theta(\min(|\mathbf{k}| + |\mathbf{p}|, |\mathbf{k}'| + |\mathbf{p}'|) - |\mathbf{q}|) \quad (3.39)$$

as our reduced repopulation integral from the s -channel. The u - and t -channel integral reduction proceeds in a similar manner, using \mathbf{v} and \mathbf{w} respectively as an integration variable.

3.1.A.2 The massive case

The reduced integral in Equation (3.39) is but one of three contributions to the repopulation of ν_α arising from neutrino scattering with electrons. The full collision term, including scattering and annihilation processes with $e^\pm, \nu_\alpha^{(-)}, \nu_\beta^{(-)}$, has in total 14 such terms to be evaluated (see Table 3.1). Fortunately, however, these 14 terms can all be recast into one of the standard s -, t -, and u -forms, and thus can be handled in ways similar to that discussed above.

As the reduction procedure concerns only kinematics and does not involve the actual matrix element besides the initial classification of the momentum dependence into s -, u - or t -forms, we shall keep the calculation as general as possible and allow for the

possibility that all initial and final states are massive. Then, the number of independent reductions to be performed is only three, which yield:

$$R_{\alpha, s\text{-channel}} = \frac{1}{2^7 \pi^3 E_k |\mathbf{k}|} \int_{m_{k'}}^{\infty} dE_{k'} \int_{\max(m_p, k' - k + m_{p'})}^{\infty} dE_p \int d|\mathbf{q}| S |M_s|^2 F \\ \times \theta(|\mathbf{q}| - \max(|\mathbf{k}| - |\mathbf{p}|, |\mathbf{k}'| - |\mathbf{p}'|)) \theta(\min(|\mathbf{k}| + |\mathbf{p}|, |\mathbf{k}'| + |\mathbf{p}'|) - |\mathbf{q}|), \quad (3.40)$$

$$R_{\alpha, t\text{-channel}} = \frac{A_k}{2^7 \pi^3 E_k |\mathbf{k}|} \int_{m_{k'}}^{\infty} dE_{k'} \int_{\max(m_p, k' - k + m_{p'})}^{\infty} dE_p \int d|\mathbf{w}| S |M_t|^2 F \\ \times \theta(|\mathbf{w}| - \max(|\mathbf{p}| - |\mathbf{p}'|, |\mathbf{k}| - |\mathbf{k}'|)) \theta(\min(|\mathbf{p}| + |\mathbf{p}'|, |\mathbf{k}| + |\mathbf{k}'|) - |\mathbf{w}|), \quad (3.41)$$

$$R_{\alpha, u\text{-channel}} = \frac{1}{2^7 \pi^3 E_k |\mathbf{k}|} \int_{m_{k'}}^{\infty} dE_{k'} \int_{\max(m_p, k' - k + m_{p'})}^{\infty} dE_p \int d|\mathbf{v}| S |M_u|^2 F \\ \times \theta(|\mathbf{v}| - \max(|\mathbf{k}| - |\mathbf{p}'|, |\mathbf{k}'| - |\mathbf{p}|)) \theta(\min(|\mathbf{k}| + |\mathbf{p}'|, |\mathbf{k}'| + |\mathbf{p}|) - |\mathbf{v}|). \quad (3.42)$$

After inserting the matrix element $S |M_x|^2$ and momentum distributions F , these reduced integrals are valid for any $2 \rightarrow 2$ process.

3.1.A.3 The full collision terms

The matrix elements for all elastic and inelastic processes involving ν_α at temperatures $T \lesssim m_\mu$ are summarised in Table 3.1. These have been computed at various times by several different groups [64, 65, 73], but can be easily obtained from first principles in the four-fermion limit. Using these matrix elements and the expressions in Equation (3.40), Equation (3.41) and (3.42), we can now determine the contribution of each process to the repopulation integral. The results are as follows.

1. $\nu_\alpha(k) \nu_\beta(p) \rightarrow \nu_\alpha(k') \nu_\beta(p')$:

$$R_{\alpha, s, \beta} = \frac{G_F^2}{2(2\pi)^3 E_k^2} \int_0^\infty dE_{k'} \int_{\max(0, E_{k'} - E_k)}^\infty dE_p \int_{\max(|E_k - E_p|, |2E_{k'} - E_p - E_k|)}^{E_k + E_p} d|\mathbf{q}| \\ \times \left[(E_p + E_k)^2 - |\mathbf{q}|^2 \right]^2 \left[f(E_{k'}) f(E_{p'}) (1 - f(E_k)) (1 - f(E_p)) \right. \\ \left. - f(E_k) f(E_p) (1 - f(E_{k'})) (1 - f(E_{p'})) \right]. \quad (3.43)$$

2. $\nu_\alpha(k) \bar{\nu}_\beta(p) \rightarrow \nu_\alpha(k') \bar{\nu}_\beta(p')$:

$$R_{\alpha, s, \bar{\beta}} = \frac{G_F^2}{2(2\pi)^3 E_k^2} \int_0^\infty dE_{k'} \int_{\max(0, E_{k'} - E_k)}^\infty dE_p \int_{|E_{k'} - E_p|}^{\min(2E_k + E_p - E_{k'}, E_{k'} + E_p)} d|\mathbf{v}| \\ \times \left[(E_p - E_{k'})^2 - |\mathbf{v}|^2 \right]^2 \left[f(E_{k'}) f(E_{p'}) (1 - f(E_k)) (1 - f(E_p)) \right. \\ \left. - f(E_k) f(E_p) (1 - f(E_{k'})) (1 - f(E_{p'})) \right]. \quad (3.44)$$

3 Full collision term

3. $\nu_\alpha(k)\nu_\alpha(p) \rightarrow \nu_\alpha(k')\nu_\alpha(p')$:

$$\begin{aligned}
R_{\alpha,s,\alpha} &= \frac{G_F^2}{(2\pi)^3 E_k^2} \int_0^\infty dE_{k'} \int_{\max(0, E_{k'} - E_k)}^\infty dE_p \int_{\max(|E_k - E_p|, |2E_{k'} - E_p - E_k|)}^{E_k + E_p} d|\mathbf{q}| \\
&\times \left[(E_p + E_k)^2 - |\mathbf{q}|^2 \right]^2 \left[f(E_{k'})f(E_{p'}) (1 - f(E_k))(1 - f(E_p)) \right. \\
&\quad \left. - f(E_k)f(E_p) (1 - f(E_{k'}))(1 - f(E_{p'})) \right].
\end{aligned} \tag{3.45}$$

4. $\nu_\alpha(k)\bar{\nu}_\alpha(p) \rightarrow \nu_\alpha(k')\bar{\nu}_\alpha(p')$:

$$\begin{aligned}
R_{\alpha,s,\bar{\alpha}} &= \frac{2G_F^2}{(2\pi)^3 E_k^2} \int_0^\infty dE_{k'} \int_{\max(0, E_{k'} - E_k)}^\infty dE_p \int_{|E_{k'} - E_p|}^{\min(2E_k + E_p - E_{k'}, E_{k'} + E_p)} d|\mathbf{v}| \\
&\times \left[(E_p - E_{k'})^2 - |\mathbf{v}|^2 \right]^2 \left[f(E_{k'})f(E_{p'}) (1 - f(E_k))(1 - f(E_p)) \right. \\
&\quad \left. - f(E_k)f(E_p) (1 - f(E_{k'}))(1 - f(E_{p'})) \right].
\end{aligned} \tag{3.46}$$

5. $\nu_\alpha(k)e^-(p) \rightarrow \nu_\alpha(k')e^-(p')$:

$$\begin{aligned}
R_{\alpha,s,e^-} &= \frac{G_F^2}{2(2\pi)^3 E_k^2} \int_0^\infty dE_{k'} \int_{\max(m_e, E_{k'} - E_k + m_e)}^\infty dE_p \\
&\times \left[\int d|\mathbf{q}| (2x_W \pm 1)^2 \theta(|\mathbf{q}| - \max(|E_k - |\mathbf{p}||, |E_{k'} - |\mathbf{p}'||)) \right. \\
&\quad \times \theta(\min(E_k + |\mathbf{p}|, E_{k'} + |\mathbf{p}'|) - |\mathbf{q}|) ((E_p + E_k)^2 - |\mathbf{q}|^2 - m_e^2)^2 \\
&\quad + \int d|\mathbf{v}| 4x_W^2 \theta(|\mathbf{v}| - \max(|E_k - |\mathbf{p}'||, E_{k'} - |\mathbf{p}||)) \\
&\quad \times \theta(\min(E_k + |\mathbf{p}'|, E_{k'} + |\mathbf{p}|) - |\mathbf{v}|) ((E_p - E_{k'})^2 - |\mathbf{v}|^2 - m_e^2)^2 \\
&\quad + \int d|\mathbf{w}| 4m_e^2 (2x_W \pm 1) x_W \theta(|\mathbf{w}| - \max(|\mathbf{p}| - |\mathbf{p}'|, |E_k - E_{k'}|)) \\
&\quad \times \theta(\min(|\mathbf{p}| + |\mathbf{p}'|, E_k + E_{k'}) - |\mathbf{w}|) ((E_k - E_{k'})^2 - |\mathbf{w}|^2) \left. \right] \\
&\times \left[f(E_{k'})f(E_{p'}) (1 - f(E_k))(1 - f(E_p)) \right. \\
&\quad \left. - f(E_k)f(E_p) (1 - f(E_{k'}))(1 - f(E_{p'})) \right].
\end{aligned} \tag{3.47}$$

6. $\nu_\alpha(k)e^+(p) \rightarrow \nu_\alpha(k')e^+(p')$:

$$\begin{aligned}
 R_{\alpha,s,e^+} &= \frac{G_F^2}{2(2\pi)^3 E_k^2} \int_0^\infty dE_{k'} \int_{\max(m_e, E_{k'} - E_k + m_e)}^\infty dE_p \\
 &\times \left[\int d|\mathbf{q}| 4x_W^2 \theta(|\mathbf{q}| - \max(|E_k - |\mathbf{p}||, |E_{k'} - |\mathbf{p}'||)) \right. \\
 &\quad \times \theta(\min(E_k + |\mathbf{p}|, E_{k'} + |\mathbf{p}'|) - |\mathbf{q}|) ((E_p + E_k)^2 - |\mathbf{q}|^2 - m_e^2)^2 \\
 &\quad + \int d|\mathbf{v}| (2x_W \pm 1)^2 \theta(|\mathbf{v}| - \max(|E_k - |\mathbf{p}'||, E_{k'} - |\mathbf{p}||)) \\
 &\quad \times \theta(\min(E_k + |\mathbf{p}'|, E_{k'} + |\mathbf{p}|) - |\mathbf{v}|) ((E_p - E_{k'})^2 - |\mathbf{v}|^2 - m_e^2)^2 \\
 &\quad + \int d|\mathbf{w}| 4m_e^2 (2x_W \pm 1) x_W \theta(|\mathbf{w}| - \max(||\mathbf{p}| - |\mathbf{p}'||, |E_k - E_{k'}|)) \\
 &\quad \times \theta(\min(|\mathbf{p}| + |\mathbf{p}'|, E_k + E_{k'}) - |\mathbf{w}|) ((E_k - E_{k'})^2 - |\mathbf{w}|^2) \left. \right] \\
 &\times \left[f(E_{k'}) f(E_{p'}) (1 - f(E_k)) (1 - f(E_p)) \right. \\
 &\quad \left. - f(E_k) f(E_p) (1 - f(E_{k'})) (1 - f(E_{p'})) \right].
 \end{aligned} \tag{3.48}$$

 7. $\nu_\alpha(k)\bar{\nu}_\alpha(p) \rightarrow \nu_\beta(k')\bar{\nu}_\beta(p')$:

$$\begin{aligned}
 R_{\alpha,a,\beta} &= \frac{G_F^2}{2(2\pi)^3 E_k^2} \int_0^\infty dE_p \int_0^{E_k + E_p} dE_{k'} \int_{|E_{k'} - E_p|}^{\min(2E_k + E_p - E_{k'}, E_{k'} + E_p)} d|\mathbf{v}| \\
 &\times \left[(E_p - E_{k'})^2 - |\mathbf{v}|^2 \right]^2 \left[f(E_{k'}) f(E_{p'}) (1 - f(E_k)) (1 - f(E_p)) \right. \\
 &\quad \left. - f(E_k) f(E_p) (1 - f(E_{k'})) (1 - f(E_{p'})) \right].
 \end{aligned} \tag{3.49}$$

3 Full collision term

8. $\nu_\alpha(k)\bar{\nu}_\alpha(p) \rightarrow e^-(k')e^+(p')$:

$$\begin{aligned}
R_{\alpha,a,e} &= \frac{G_F^2}{2(2\pi)^3 E_k^2} \int_{\min(0, 2m_e - E_k)}^\infty dE_p \int_{m_e}^{E_k + E_p - m_e} dE_{k'} \\
&\times \left[\int d|\mathbf{q}| (2x_W \pm 1) 4x_W m_e^2 \theta(|\mathbf{q}| - \max(|E_k - E_p|, ||\mathbf{k}'| - |\mathbf{p}'||)) \right. \\
&\quad \times \theta(\min(E_k + E_p, |\mathbf{k}'| + |\mathbf{p}'|) - |\mathbf{q}|) (|\mathbf{q}|^2 - (E_p + E_k)^2) \\
&\quad + \int d|\mathbf{v}| 4x_W^2 \theta(|\mathbf{v}| - \max(|E_k - |\mathbf{p}'||, ||\mathbf{k}'| - E_p|)) \\
&\quad \times \theta(\min(E_k + |\mathbf{p}'|, |\mathbf{k}'| + E_p) - |\mathbf{v}|) ((E_p - E_{k'})^2 - |\mathbf{v}|^2 - m_e^2)^2 \\
&\quad + \int d|\mathbf{w}| (2x_W \pm 1)^2 \theta(|\mathbf{w}| - \max(|E_p - |\mathbf{p}'||, |E_k - |\mathbf{k}'||)) \\
&\quad \times \theta(\min(E_p + |\mathbf{p}'|, E_k + |\mathbf{k}'|) - |\mathbf{w}|) ((E_k - E_{k'})^2 - |\mathbf{w}|^2 - m_e^2)^2 \left. \right] \\
&\times \left[f(E_{k'}) f(E_{p'}) (1 - f(E_k)) (1 - f(E_p)) \right. \\
&\quad \left. - f(E_k) f(E_p) (1 - f(E_{k'})) (1 - f(E_{p'})) \right].
\end{aligned} \tag{3.50}$$

We remind the reader again that $|\mathbf{p}'|$ is not a free parameter, but is constrained by energy conservation.

The integrals over $|\mathbf{q}|$, $|\mathbf{v}|$ and $|\mathbf{w}|$ can be evaluated analytically, and it turns out that they fall into two different functional forms:

$$\begin{aligned}
\int dx (a - x^2)^2 &= a^2 x - \frac{2}{3} a x^3 + \frac{x^5}{5} + \text{constant}, \\
\int dx (a - x^2) &= a x - \frac{x^3}{3} + \text{constant}.
\end{aligned}$$

The remaining two integrals over E_p and $E_{k'}$ must be performed numerically, although as discussed in Section 3.1.1.3 judicious assumptions about certain distribution functions in the integrand make it possible to pre-evaluate some of the integrals only once, rather than evaluating them in real time simultaneously with the numerical solution of the QKEs.

3.1.B Repopulation and damping coefficients in the A/S approximation

We summarise here the full expressions for the dimensionless repopulation and damping coefficients in the A/S approximation discussed in Section 3.1.1.2. The quantity

3.1 \mathfrak{P} Active–sterile oscillations in the early Universe with full collision terms

$A \equiv 2\pi/\int d\Pi_k k f_0(k)$ is a normalisation factor.

$$C_{\alpha,a} = A \int d\Pi_k d\Pi_{k'} d\Pi_{p'} d\Pi_p \delta_E(kp|k'p') \sum_i \mathcal{V}^2[\nu_\alpha(k), \bar{\nu}_\alpha(p)|i(k'), \bar{i}(p')] f_0(p) f_0(k),$$

$$C_{\alpha,s} = A \int d\Pi_k d\Pi_{k'} d\Pi_{p'} d\Pi_p \delta_E(kp|k'p') \sum_{j \neq \nu_\alpha, \bar{\nu}_\alpha} \mathcal{V}^2[\nu_\alpha(k), j(p)|\nu_\alpha(k'), j(p')] f_0(p) f_0(k),$$

$$C_{\alpha,\nu} = A \int d\Pi_k d\Pi_{k'} d\Pi_{p'} d\Pi_p \delta_E(kp|k'p') \sum_{j = \nu_\alpha, \bar{\nu}_\alpha} \mathcal{V}^2[\nu_\alpha(k), j(p)|\nu_\alpha(k'), j(p')] f_0(p) f_0(k),$$

$$C_{\alpha,0} = A \int d\Pi_k d\Pi_{k'} d\Pi_{p'} d\Pi_p \delta_E(kp|k'p') f_0(k) \times \left\{ \sum_i \mathcal{V}^2[\nu_\alpha(k), \bar{\nu}_\alpha(p)|i(k'), \bar{i}(p')] f_0(k') f_0(p') + \sum_{j \neq \nu_\alpha, \bar{\nu}_\alpha} \mathcal{V}^2[\nu_\alpha(k), j(p)|\nu_\alpha(k'), j(p')] f_0(p) (1 - f_0(p')) \right\},$$

$$C_{\alpha,1} = A \int d\Pi_k d\Pi_{k'} d\Pi_{p'} d\Pi_p \delta_E(kp|k'p') f_0(k) \times \left\{ \sum_i \mathcal{V}^2[\nu_\alpha(k), \bar{\nu}_\alpha(p)|i(k'), \bar{i}(p')] f_0(p) (1 - f_0(p') - f_0(k')) + \sum_{j \neq \nu_\alpha, \bar{\nu}_\alpha} \mathcal{V}^2[\nu_\alpha(k), j(p)|\nu_\alpha(k'), j(p')] f_0(k') (f_0(p') - f_0(p)) + \sum_{j = \nu_\alpha, \bar{\nu}_\alpha} \mathcal{V}^2[\nu_\alpha(k), j(p)|\nu_\alpha(k'), j(p')] f_0(p) \right\},$$

$$C_{\alpha,2} = A \int d\Pi_k d\Pi_{k'} d\Pi_{p'} d\Pi_p \delta_E(kp|k'p') f_0(k) \times \sum_{j = \nu_\alpha, \bar{\nu}_\alpha} \mathcal{V}^2[\nu_\alpha(k), j(p)|\nu_\alpha(k'), j(p')] (f_0(k') f_0(p') - f_0(k') f_0(p) - f_0(p) f_0(p')).$$

See Equation (3.16) and (3.17) for the implementation of these coefficients in the A/S repopulation and damping terms.



3.2 Lepton asymmetries and active-active oscillations

In the previous sections, we have treated oscillations between one active and one sterile neutrino in the early Universe assuming that there is no asymmetry between particles and antiparticles. This is most probably wrong, as we know that there is a small baryon asymmetry of order $\eta = (n_b - n_{\bar{b}})/n_\gamma \sim 10^{-9}$ which has given rise to stars, planets, humans and PhD theses. However, such a small asymmetry would not affect the neutrino oscillations much, but while we know the baryon asymmetry quite well, we do not have the same handles on the lepton asymmetry, which could be as large as ~ 0.1 . Therefore, the asymmetric case would be one interesting expansion of the improved collision treatment.

The full calculation of the collision terms is in principle quite easy to extend to the asymmetric case, as the integrals are exactly the same. The main issue is that, as we will see in the next chapter, especially models with a large lepton asymmetry tend to present much more complicated solutions compared to what we have seen in this chapter. This added complexity requires the integrator to take much smaller and hence many more steps. For the full collision treatment this gives two challenges. The first one is that of numerical evaluation cost which will make high resolution calculations unfeasible. The second challenge comes from the many evaluation of the collision integrals which could give tiny deviations from detailed balance in each step. As the calculation is much more complex than any of the approximations, it will be necessary to have a thorough control of the numerical error induced by the combination of many ODE steps and the collision integrals.

For the approximations, one need to be more explicit on the adoption to the asymmetric case. The usual approach for both the equilibrium approximation and for the CC approximation is to modify the equilibrium distributions using a chemical potential defined by the condition

$$L = \frac{1}{N_\gamma} \int \frac{d^3\mathbf{k}}{(2\pi)^3} \frac{1}{e^{k/T-\mu/T}} - \frac{1}{e^{k/T+\mu/T}}. \quad (3.51)$$

Similar to what we noticed just after Equation (3.15), this definition of μ will not in itself lead to conservation of the total lepton number. For the CC approximation this is not a problem, as the scattering term has a vanishing diagonal, and the $\bar{\rho}$ in the last term of Equation (3.10) ensures cancellation between the neutrino and antineutrino annihilation terms. For the equilibrium approximation the problem is usually avoided by including an additional equation for \bar{L} when solving the QKEs, thereby explicitly avoiding any undesired effect from the collision term on L .

An alternative to this could be to modify the definition of the chemical potential mentioned above. If the integral was done over $\int d^3\mathbf{k}k/(2\pi)^3$ as in Equation (3.19), L would actually be conserved as long as Equation (3.51) with the modified integrant is fulfilled numerically.

The A/S approximation similarly requires some modifications. For the two scattering terms, it would be necessary to calculate the equilibrium distributions for neutrinos and antineutrino separately as they could have quite different distribution functions. The annihilation term need no calculation of equilibrium distributions, but requires another modification. Using Equation (3.16) and remembering to replace n_{ν_α} appropriately, the

3.2 Lepton asymmetries and active-active oscillations

collisional contribution to the lepton asymmetry time derivative would have the form

$$\begin{aligned}
 \dot{L}_{\text{coll}} &\propto \int dk k^2 \left(R_{A/S}(k) - \bar{R}_{A/S}(k) \right) \\
 &\propto \int dk k^3 (f_0 - n_{\bar{\nu}_\alpha} \rho_{\alpha\alpha} - f_0 + n_{\nu_\alpha} \bar{\rho}_{\alpha\alpha}) \\
 &\propto n_{\nu_\alpha} u_{\bar{\nu}_\alpha} - n_{\bar{\nu}_\alpha} u_{\nu_\alpha},
 \end{aligned} \tag{3.52}$$

given that the scattering terms are already handled, and where u_x is the energy density in particle x . Hence, the solution is to redefine the annihilation term using a normalised energy density rather than the normalised number density, which will ensure that annihilations do not change L .

Actually, it is in general a better solution to use energy rather than number density. We know that R scales approximately linearly with the momentum k , but all the underlying annihilation reactions are symmetric under $\nu_\alpha \leftrightarrow \bar{\nu}_\alpha$. Therefore, the integral that is done to get rid of the antineutrino distribution in R and the neutrino distribution in \bar{R} should also scale approximately linearly in the integrated momentum p , which makes the energy density a better approximation than the number density. While the results in Section 3.1 show that the difference is negligible for the neutrino-antineutrino symmetric case, the violation of lepton asymmetry makes it paramount for the asymmetric case.

While the adaption of the full collision term to the asymmetric case takes little effort, a more extensive but also interesting change would be to expand it to three oscillating neutrinos. This would allow the treatment of all the three neutrinos in the Standard Model, and such a calculation could improve the neutrino decoupling calculation, where the current state of the art only uses the full integral for the repopulation part of the collision term [74]. However, apart from expanding *LASAGNA* to three neutrinos, it would also be necessary to include new integrals. McKellar and Thomson [45] describe two new types of terms that only arise when more than one active neutrino is involved. The first type is related to the annihilation of a neutrino of one flavour with an antineutrino of another flavour due to the mixing which is obviously absent when only one active neutrino is considered. The other type is related to \mathbf{P}_T^* , which induce additional quantum damping terms in addition to the usual damping that is described by D . While McKellar and Thomson did not include Pauli blocking terms, and their result therefore cannot be applied directly, there is no reason to think that these effects should not be included in a complete treatment of active-active oscillations. One option would now be to take the equations from McKellar and Thomson and rederive all the collision terms when Pauli blocking is included. However, Sigl and Raffelt have already done this derivation [56], albeit with a quite different notation, and their result could also be used as the starting point of the complete calculation.

Large lepton asymmetry

The idea of investigating how a large lepton asymmetry and neutrino oscillations affect each other has been around for several decades [75–77]. The first mention by Khlopov and Petcov [75] contained only a very schematic description of the effects, while the later work by Savage, Malaney and Fuller [76] discussed the phenomenology in much more detail. However, the first detailed numerical study was performed by Enqvist, Kainulainen and Maalampi [77], where they solved the momentum averaged quantum rate equations giving the first detailed description of combining a large lepton asymmetry with neutrino oscillations.

Our interest in a large lepton asymmetry is motivated by the tension between the observational determination of N_{eff} from cosmological observations and the hints from laboratory experiments which can be interpreted as an 1eV sterile neutrino. As we saw in Chapter 3, such a sterile neutrino would be fully thermalised, and therefore it would contribute with a full neutrino degree of freedom to the relativistic energy density. This is, as already mentioned, excluded by observations [21], and a mechanism to suppress the sterile neutrino thermalisation would be welcome.

The first study suggesting that a large lepton asymmetry could prevent the production of sterile neutrinos in the early Universe came from Foot and Volkas [78]. They included the asymmetric contribution to the background potential and estimated the effects analytically comparing the results to the most stringent bounds from BBN that were available at the time. The first treatment using a momentum dependent treatment was done by Kirilova and Chizhov [79], but by neglecting the collision terms, they prevented themselves from making any predictions about N_{eff} . A series of other studies using various levels of refinement in their analytical estimates followed [80–82], concluding that the suppression mechanism is quite complex, and depends in a intricate way on the different momentum states once the bulk of the lepton asymmetry is depleted by converting only neutrinos and no antineutrinos. Despite this, the following numerical attempts at treating the problem used the momentum averaged quantum rate equations [54, 68]. For Chu and Cirelli [68], the numerical methods and corresponding computational power were not sufficiently mature to solve the full momentum dependent QKEs, while Mirizzi, Saviano et al. [54] used the momentum averaged case as an stepping stone towards the solution of QKEs. The first full solution of the QKEs including the collision terms was done by Hannestad, Tamborra and Tram [48] using an early version of LASAGNA which will also be used for the calculations in this chapter. The paper presented precise values of ΔN_{eff} both in the case of no lepton asymmetry, and with $L = 10^{-2}$. The model used was a two neutrino approach which will also be used here, but the first QKE calculation using a more realistic three neutrino model has also been presented by Saviano, Mirizzi et al. [55]. However, the momentum resolution used in that calculation is quite low, and the produced neutrino spectra look quite different to

4 Large lepton asymmetry

what one would expect. The authors claim that more precise calculations are premature due to the very large uncertainty on the sterile mixing parameters. However, it would be interesting to have an independent test of their results, especially considering the just mentioned issues.

4.1 Quantum kinetic equations

We will now briefly summarise the equations that describe neutrino oscillations in the presence of a large lepton asymmetry.

As mentioned before, we assume that the system is well described by oscillations between only two neutrinos – an active neutrino, ν_α , and a sterile neutrino, ν_s . This assumption is reasonable as the overall agreement between the results obtained by Hannestad et al. and Saviano et al. demonstrates, although a three or four neutrino treatment obviously would be better. The limitation to two neutrinos means that we can use the polarisation vector formalism from Section 2.2.5, just as we did in the preceding section treating the full collision term. The main difference is that we now wish to distinguish between the evolution of P_i describing the neutrinos and \bar{P}_i describing the antineutrinos. As the potential depends on the difference between these two quantities, it is useful from a numerical point of view to consider the sums and differences of P_i and \bar{P}_i .

$$P_i^\pm = P_i \pm \bar{P}_i, \quad i = 0, x, y, z. \quad (4.1)$$

Finally, we can simplify the equations a bit by changing back from P_i , $i = 0, x, y, z$ to the more physical variables as it was also suggested in the end of Section 2.2.5.

$$P_a^\pm = P_0^\pm + P_z^\pm = 2\frac{\rho_{aa}^\pm}{f_0}, \quad P_s^\pm = P_0^\pm - P_z^\pm = 2\frac{\rho_{ss}^\pm}{f_0}. \quad (4.2)$$

Now, we can expand Equation (2.32) as

$$\dot{P}_a^\pm = V_x P_y^\pm + \Gamma \left[2f_{\text{eq}}^\pm / f_0 - P_a^\pm \right], \quad (4.3)$$

$$\dot{P}_s^\pm = -V_x P_y^\pm, \quad (4.4)$$

$$\dot{P}_x^\pm = -(V_0 + V_1) P_y^\pm - V_L P_y^\mp - D P_x^\pm, \quad (4.5)$$

$$\dot{P}_y^\pm = (V_0 + V_1) P_x^\pm + V_L P_x^\mp - \frac{1}{2} V_x (P_a^\pm - P_s^\pm) - D P_y^\pm. \quad (4.6)$$

Here, the equilibrium distribution functions are given by

$$f_{\text{eq}}^\pm = \frac{1}{1 + e^{x-\xi}} \pm \frac{1}{1 + e^{x+\xi}}, \quad (4.7)$$

where $x = p/T$ is the comoving momentum and the degeneracy parameter $\chi = \mu/T$ is defined in terms of the chemical potential, μ . The potentials have been split up emphasising that only the vacuum term contributes to V_x ¹, while the z -component gets contributions from both the vacuum term, V_0 , the symmetric matter term, V_1^a , and the

¹Notice that the x here is the coordinate in polarisation space and has nothing to do with p/T .

asymmetric matter term, V_L . The expressions for the various contributions are

$$V_x = \frac{\delta m_s^2}{2xT} \sin 2\theta_s, \quad (4.8)$$

$$V_0 = -\frac{\delta m_s^2}{2xT} \cos 2\theta_s, \quad (4.9)$$

$$V_1^\alpha = -\frac{7\pi^2}{45\sqrt{2}} \frac{G_F}{M_Z^2} xT^5 [n_{\nu_\alpha} + n_{\bar{\nu}_\alpha}] g_\alpha, \quad (4.10)$$

$$V_L^\alpha = \frac{2\sqrt{2}\zeta(3)}{\pi^2} G_F T^3 L_{(\alpha)}. \quad (4.11)$$

The two matter terms depend on which ν_α is considered through g_α which is given by $g_e = 1 + 4 \sec^2 \theta_W / (n_{\nu_e} + n_{\bar{\nu}_e})$ and $g_{\mu,\tau} = 1$, and also through $L_{(\alpha)}$ expressed as

$$L_{(e)} = \left(\frac{1}{2} + 2 \sin^2 \theta_W\right) L_e + \left(\frac{1}{2} - 2 \sin^2 \theta_W\right) L_p - \frac{1}{2} L_n + 2L_{\nu_e} + L_{\nu_\mu} + L_{\nu_\tau}, \quad (4.12)$$

$$L_{(\mu)} = L_{(e)} - L_e - L_{\nu_e} + L_{\nu_\mu}, \quad (4.13)$$

$$L_{(\tau)} = L_{(e)} - L_e - L_{\nu_e} + L_{\nu_\tau}, \quad (4.14)$$

when we take all contributions to Equation (2.21) into account, and define $L_f \equiv (n_f - n_{\bar{f}})N_f/N_\gamma$ with N_f and N_γ being the fermion and photon number densities in equilibrium respectively, and n_f being the number density normalised to 1 in thermal equilibrium.

Finally, for the collision term we use the equilibrium approximation from Equation (2.25) and (2.26). This gives us the damping

$$D = \frac{1}{2}\Gamma, \quad \Gamma = C_\alpha G_F^2 xT^5, \quad (4.15)$$

where the collision rate, Γ , also appears explicitly in repopulation term of Equation (4.3).

The approximate form of the repopulation term pose a problem for the conservation of lepton number, a problem that was also discussed in Section 3.2, where several solutions were suggested. For now we will follow the previously used approach. While the active neutrino lepton number is obviously not conserved, the global lepton number including the sterile component should be. For this to be true, the integral

$$\int dx x^2 f_0 (\dot{P}_a^- + \dot{P}_s^-) = \int dx x^3 C_\alpha G_F^2 T^5 (2f_{\text{eq}}^- - f_0 P_a^-)$$

must be zero. This is only true if the third moments of $2f_{\text{eq}}^-$ and $f_0 P_a^-$ are identical. However, when determining the degeneracy parameter, we use the approximation [48]

$$\xi = \frac{-2\pi}{\sqrt{3}} \sinh \left(\frac{1}{3} \operatorname{arcsinh} \left[-\frac{18\sqrt{3}\zeta(3)}{\pi^3} L_{(\alpha)} \right] \right), \quad (4.16)$$

which gives no such guarantee. Therefore, we need to evolve the lepton asymmetry using the equation

$$\dot{L}_{(a)} = \frac{1}{8\zeta(3)} \int_0^\infty dx x^2 f_0 V_x P_y^-. \quad (4.17)$$

4 Large lepton asymmetry

With this expression we have explicitly removed the globally lepton number violating term and only included the conversion between active and sterile neutrinos.

The equations are solved using `LASAGNA` for the mixing parameters $\delta m^2 = 0.89\text{eV}^2$ and $\sin^2 2\theta = 0.025$ for mixing between an electron neutrino and a sterile neutrino [55, 83]. The evolution was followed all the way to $T = 0.1\text{MeV}$ which is necessary to calculate the effect on BBN.

4.2 Convergence problems

Before the effects of the large lepton asymmetry and the sterile neutrino on BBN can be found, a stable result of the QKEs is needed. So far this has failed in at least three important aspects. First of all, N_{eff} does not converge towards any value, when the momentum resolution and precision of the solver is improved. Secondly, the evolution of L neither converges, and it fails to show a persistent evolution pattern for the various choices of resolution and precision. Finally, the momentum distribution function attains unphysical negative values in a few cases. Due to these issues, it has not yet been possible to produce any interesting results, but in the interest of documenting the problems, they will now be described in some details along with some of the attempted solutions.

Before we dive into the unpleasant details of convergence problems, it is interesting to consider briefly what others have done. The well informed reader will know that a very similar computation was already done with the same code by Hannestad, Tamborra and Tram [48], where the convergence problems were overcome. However, the method employed there was to restart the code when $L \approx 0$ using the adiabatic approximation as new initial conditions in order to avoid very fast oscillations that happens once L is very small. While this works well for producing a stable N_{eff} , it would change the momentum spectra which we are interesting in for the BBN calculation and the method might not capture all the relevant contributions to N_{eff} . Another similar calculation is due to Saviano et al. [55]. As already mentioned in the introduction to this chapter, they only use quite few momentum bins in their calculation. Furthermore, they use Gaussian quadratures to integrate the momentum distributions, and while this is a questionable approach from a physical point of view because Gaussian quadratures require quite smooth functions, it could also have a stabilising effect on the evolution of both L and N_{eff} .

In the present case it would be preferable to avoid both restarting the code and the use of Gaussian quadratures, but so far the convergence is lacking. Although all three problems might of course be linked, the non-negativity of the distribution functions stands a bit out compared to the two genuine convergence problems, and we will begin by considering the latter.

As it was demonstrated in the preceding section, it is possible to obtain results with less than one percent deviations in N_{eff} even when using less than 100 momentum bins. This was for the full collision term, but without a lepton asymmetry. However, we should expect to obtain convergence at the percent level for the lepton asymmetric case as well. In Figure 4.1, we see that the various runs of `LASAGNA` have produced quite different values for both ΔN_{eff} and L depending on both the number of momentum bins, v_{res} and the absolute error tolerance, ε . For the $v_{\text{res}} = 150$ cases (dashed lines), it could




Figure 4.1: Lack of convergence in N_{eff} and L .

Top: Change in relative relativistic energy density. *Bottom:* Lepton asymmetry. Both as a function of decreasing temperature. The three dashed lines correspond to 150 momentum bins with absolute error tolerance 10^{-4} (cyan short), 10^{-6} (green medium) and 10^{-7} (blue long). The remaining two have a tolerance of 10^{-6} along with 100 (magenta dotted) and 200 (red dot-dashed) bins.

4 Large lepton asymmetry




Figure 4.2: Negative values for $\rho_{\alpha\alpha}$ and ρ_{ss} .

Left: The active neutrino distributions for three different temperatures around 2.4MeV and for three different solutions of the QKEs using 150 bins and a tolerance of 10^{-6} (green short-dashed) and 10^{-7} (blue long-dashed) and using 200 bins and a tolerance of 10^{-6} . *Right:* Same for the sterile neutrino distribution functions. All values above 1 and below 0 are unphysical for the parameters shown.

seem like there is convergence when the tolerance is decreased from 10^{-4} (cyan short) to 10^{-6} (green medium) and 10^{-7} (blue long). However, when other numbers of bins are considered, this appears like a mere accident. There is some difference between the 100 (magenta dotted) and the 150 bin cases with the same tolerance, but the difference between 150 and 200 (red dot-dashed) is larger. The same picture presents itself for the lepton asymmetry. There is a hint of convergence for the 150 bin cases, but the general picture is not convincing. Furthermore, there is the complication of L crossing zero. The complicated structure around $T = 2.4$ MeV comes from repeated zero crossings, and as long as these are not reproduced to a reasonable extend by the different parameter choices, the computation cannot be expected to be stable.

The third problem is related to unphysical negative values of ρ . An example is seen

in Figure 4.2, where ρ_{ss} attains large negative values in the lower right panel. At the same time $\rho_{\alpha\alpha}$ shows similarly large positive values in the lower left panel. The upper two panels illustrate the situation at temperatures where V_L still dominates V_z , while the panels in the middle show the situation immediately after. Although it is possible to get values of ρ above 1 due to the collision term, this cannot possibly happen on such a short timescale as it is seen here, and the large $\rho_{\alpha\alpha}$ values arise from the same problem as the negative ρ_{ss} values.

The unphysical values arise within a very short time period, and must originate in numerical inaccuracy of the solver. As the differential equation describing the oscillations can be written using a cross product as seen in Equation (2.28), the length of \mathbf{P} should be conserved. This works well down to $T = 2.445\text{MeV}$, where the lepton asymmetry is decreasing, and active neutrinos up to $x \sim 0.3$ are converted in an adiabatic resonance. The polarisation vectors give a useful picture of this which we will also use in Chapter 8.2. When V_L dominates, the polarisation vector rotates around V_z pointing in the $+z$ -direction in polarisation space. As V_0 begins to dominate V_z , its direction is gradually changed to the $-z$ -direction, and the polarisation vector is dragged along. In this picture, it is also easy to understand why the complications above $x \sim 0.3$ arise in the lower panels of Figure 4.2. At $T = 2.439\text{MeV}$, the resonance has started converting the bulk of the neutrino distribution, and the resonance moves very quickly through all the remaining momenta. When doing this, it pushes all the polarisation vectors out of the rotation around $+z$ and leaves them with a V_z pointing in the $-z$ -direction. The result is that polarisation vectors describing different momentum states move around more or less independently and oscillates very fast on very non-trivial trajectories.

Although the instability just described is the origin of the numerical issues, it does not present an immediate solution. The main problems identified is the failing convergence for N_{eff} and L and the accumulation of error that leads to unphysical values of the polarisation vector. In order to solve these problems we have attempted two different solutions. The first and simplest is to use another integrator, RADAU5, which use a different method for solving stiff differential equations, and the second is to introduce a minimal stepsize in the hope that this will limit the accumulated error.

4.2.1 Solving the QKEs using RADAU5

So far the method used for solving the QKEs was `ndf15` with the maximal order restricted to 2, but we will now attempt to use RADAU5 instead. As it was mentioned in Section 2.4, the advantage of RADAU5 is the larger step size, which could possibly limit the accumulated error. However, the higher order of RADAU5 also means that every step is much more expensive than in `ndf15`, and the result is that the calculations with most bins and lowest tolerance cannot be done in a reasonable time with RADAU5.

The convergence of ΔN_{eff} and L is shown in Figure 4.3. The convergence of ΔN_{eff} looks slightly better than in Figure 4.1. This time it is for increasing number of momentum bins, where the three cases with $\varepsilon = 10^{-4}$ and 100 (cyan short dashed), 150 (magenta dotted) and 200 bins (red dot-dashed) agree reasonably well. However, the lower tolerance solution deviates significantly, again questioning the apparent convergence. When looking at the lepton asymmetry, it looks slightly better. The only real outlier is the low resolution, low tolerance calculation. However, the persistent variation for $\varepsilon = 10^{-6}$ is not very promising as the naive expectation would be that

4 Large lepton asymmetry




Figure 4.3: RADAU5 does not solve the problems for ΔN_{eff} and L .
Top: Change in relative relativistic energy density. *Bottom:* Lepton asymmetry. Both as a function of decreasing temperature. The two dashed lines correspond to 100 momentum bins with absolute error tolerance 10^{-4} (cyan short) and 10^{-6} (green medium). The remaining two have a tolerance of 10^{-4} along with 150 (dotted, red) and 200 (red dot-dashed) bins.




Figure 4.4: RADAU5 can also give too large values for $\rho_{\alpha\alpha}$.

The active neutrino distributions for three different solutions of the QKEs using 100 bins and a tolerance of 10^{-6} (green dashed), 150 bins and a tolerance of 10^{-4} (magenta dotted) and 200 bins and a tolerance of 10^{-4} (red dot-dashed). All values above 1 are unphysical for the parameters shown.

an earlier stabilisation of L is better for the numerical stability. A final point that is also worth noticing is the agreement between the high resolution case using `ndf15` and the high precision case using `RADAU5`. It could be that both are wrong and it is just a coincidence, but it could also point to the correct behaviour of ΔN_{eff} .

Regarding the unphysical values of $\rho_{\alpha\alpha}$ and ρ_{ss} , it looks somewhat promising. In Figure 4.4, we see that the high resolution, high tolerance calculations still give rise to too high values for $\rho_{\alpha\alpha}$, and the corresponding values of ρ_{ss} are again negative. However, the low resolution, low tolerance calculation show no unphysical values, and this is true for all temperatures. Nonetheless, this might not be true for higher resolution calculations with the same tolerance as we saw in Figure 4.2 that the highest resolutions also leads to the largest violations of ρ_{ss} non-negativity. An additional problem is that the `RADAU5` algorithm is not as fast as `ndf15`, and it requires hundreds of computation hours to make the highest resolution and lowest tolerance calculations.

4.2.2 Fixed minimal step size

This attempt to improve the convergence is based on a modification of the `ndf15` method, where a minimal possible step size is introduced in the otherwise adaptive algorithm. This obviously violates the algorithms attempt at limiting the local error to some set tolerance, but the hope is that the violation is not too large and that the accumulated error will be smaller. Given that the adaptive algorithm is there for a reason, this might be a naive hope, but the results will tell.

Apart from setting the minimal step size, h , it is also necessary to modify the algorithm such that it does not abort when it reaches a step size that would normally make it give up. In Figure 4.5, the convergence of ΔN_{eff} and L is shown. The spread for ΔN_{eff} is again smaller than it is for the unmodified `ndf15`, but it is hard to pick an objectively better choice for the minimal step size. From ΔN_{eff} , it looks like $h = 10^{-12.5}$ gives the most significant outliers (magenta dotted). Similarly if we consider L , we

4 Large lepton asymmetry




Figure 4.5: No convergence in ΔN_{eff} and L for a fixed step size.

Top: Change in relative relativistic energy density. *Bottom:* Lepton asymmetry. Both as a function of decreasing temperature. The three dashed lines correspond to 100 momentum bins with a minimal step size $10^{-12.5}$ (cyan short), 10^{-13} (green medium) and $10^{-13.5}$ (blue long). The remaining three have 150 momentum bins and again the minimal step sizes $10^{-12.5}$ (magenta dotted), 10^{-13} (red dot-dashed) and $10^{-13.5}$ (yellow dot-dot-dashed). The absolute tolerance for all runs was $\varepsilon = 10^{-6}$.




Figure 4.6: A fixed step size can also give $\rho_{\alpha\alpha} > 1$.

The active neutrino distributions for three different solutions of the QKEs using 150 bins and a minimal step size of $10^{-12.5}$ (magenta dotted), 10^{-13} (red dot-dashed) and $10^{-13.5}$ (yellow dot-dot-dashed). All values above 1 are unphysical for the parameters shown.

see that it is one of the runs with $h = 10^{-12.5}$ (cyan short-dashed) that deviates the most from the general trend. Although the final value looks reasonably consistent, the structure at intermediate temperatures still varies quite a lot. Despite the method converging reasonably well, the lack of error control is worrisome and it is remarkably that the runs converge towards a value significantly below the value for ΔN_{eff} that the high resolution `ndf15` run and the high precision `RADAU5` run both pointed to. Unfortunately, it is hard to tell if this is an indication that the minimal step size method removes problems inherent in the other runs, or if it is a problem in the minimal step size method itself.

Turning our attention to the value of $\rho_{\alpha\alpha}$ in Figure 4.6, we see that the momentum distribution again attains unphysical values. The exception is the run with $h = 10^{-12.5}$ which showed a lack of convergence for ΔN_{eff} . Hence, we continually lack a method that can generate convergent solutions and produce a physical momentum distribution.

4.2.3 Future improvements

For future improvements, there are a few interesting options. The first one is to improve the numerical methods. A good approach would be to focus on the unphysical values of ρ and find a differential equation solver that is especially well suited for conserving the length of \mathbf{P} . The major problem here is that the equations are stiff apart from exhibiting this instability, which means that most normal forward differential equation solvers will be very slow at solving the QKEs and will most likely fail.

The second option for improving the treatment of ρ is to consider alternative parametrisations of the equations. This could be as simple as going back to using P_0 and P_z instead of P_a and P_s , but there might also be other clever ways to express the equations, which give rise to a natural conservation of $|\mathbf{P}|$ when appropriate.

A third approach could be to use an adaptive momentum grid in order to better resolve the problematic values of the momentum distribution. This was attempted

4 Large lepton asymmetry

briefly, but `ndf15` reached too small step sizes as soon as the problematic region was encountered, and a preliminary attempt of combining an adaptive grid with the fixed minimal step size was unsuccessful. Furthermore, the negative values of ρ_{ss} seem to arise easier if anything when the momentum resolution is higher, so the adaptive momentum grid has no advantage from that point of view.

Finally, the results from the full collision term calculation in Chapter 3 should be implemented. Using the full collision term itself is probably unfeasible because of the large numerical cost of evaluating the full integrals in every integration step, but the distinction between scattering and annihilation contributions should be included as it was discussed in Section 3.2.

Sterile interactions

The concept of sterile interactions may at first seem self-contradictory as the idea of a sterile or inert neutrino is to have no interactions. The reason this topic still has drawn some attention the last few years is that a more fruitful definition of the sterile neutrino is to demand that it is a Standard Model singlet as we did in the beginning of Part I. This allows the sterile neutrino to be coupled to itself and to other particles in a hidden sector, while it does not feel any of the usual forces.

Sterile neutrinos in cosmology are currently interesting mainly for two reasons and let us review them again. The first is the eV experimental neutrino anomalies, and the second is that a sterile neutrino could contribute to the dark matter of the Universe. However, the simple non-resonant Dodelson-Widrow production [84] cannot account for the production in any of the two cases, so other effects must be at play if these sterile neutrinos exist. The problem with non-resonant production is actually opposite in the two cases. For the eV sterile neutrino, the non-resonant production lead to a larger abundance than observations permit, while the sterile neutrino dark matter abundance from non-resonant production would be very insignificant due to current limits on mass and mixing angle.

The idea behind introducing a new interaction between neutrinos in the early Universe is to obtain an effect similar to what is accomplished by a large lepton asymmetry which seems capable of solving both problems [55, 85]. The additional background potential from the new neutrino interaction can suppress production of eV-scale sterile neutrinos just as the background potential of a large lepton asymmetry would, and for other parameters, it could possibly induce a resonance giving rise to a reasonable abundance of dark matter.

Before introducing an interaction between sterile neutrinos, the possibility of using an interaction for the active neutrinos should be explored. The topic of non-standard interactions (NSI) for active neutrinos with matter is a well established field, and the limits are routinely updated by neutrino oscillation experiments. The limits on neutrino-neutrino interactions are harder to improve, and they mainly come from indirect measurements such as the LEP measurement of the Z-boson, arguments on supernova cooling and cosmology [86–88]. However, both these options are very well constrained, and it is not possible to get a significant effect on sterile neutrino production within the limits. Therefore, the sterile interactions is the next option, although we should keep in mind that the introduction of new interactions could also lead to effects for active neutrinos which might lead to constraints.

When introducing the new interaction, there are two different choices for the boson mediator. The first option is a vector boson which gives rise to a force very similar to the weak force mediated by the Z-boson, while the second option is the the pseudoscalar boson which could be e.g. a majoron. Majoron models have previously been explored

5 Sterile interactions

as an explanation for now-defunct 17keV sterile neutrino [89, 90], but vector models have not been explored previously.

In the following, we will first consider models with a vector boson that inhibits the production of an 1eV sterile neutrino. The various constraints from large scale structures, the cosmological background radiation and BBN are considered as well as the connection between such a sterile interaction and interacting dark matter. After discussing the vector boson interaction, we turn to the pseudoscalar case, where late-time phenomenology as well as dark matter connections are also considered. Finally, we consider the resonant production of warm dark matter in the form of sterile neutrinos.

5.1 Vector boson mediator

From the Standard Model of particle physics, we know of several fundamental vector bosons, both a massless one in the form of the photon and massive ones in the form of the Z- and W-bosons. From a theoretical point of view, the Z-boson and photon are the simpler ones, as they have no charge, and the Z-boson furthermore has a mass which means that it does not give long-range forces. Therefore, copies of the Z-boson are some of the most popular exotic particles, and in the next section we will investigate the consequences of the coupling between such a vector boson and sterile neutrinos for the production of sterile neutrinos.

The following paper in Section 5.2 was originally inspired by new constraints on neutrino-neutrino interactions from the CMB [91]. However, such an interaction is heavily constrained from other sources [86] as already mentioned, and we turned to sterile interactions instead. By introducing the X vector boson, we obtain an effective 4-point interaction for the sterile neutrinos similar to the weak interaction in the Standard Model. This type of interaction does in general lead to a potential similar to what we found in Equation (2.17). If we consider equal amounts of neutrinos and antineutrinos, we can discard the asymmetric first term, and this leaves us with the last symmetric term

$$H_{\text{int},\nu_s} \propto -\frac{G_X p}{M_X^2} \int d\mathbf{p} p G_S (\rho + \bar{\rho}^*) G_S, \quad (5.1)$$

where G_X is the coupling constant of the new interaction and M_X is the X -boson mass. In a two-neutrino model with one active and one sterile neutrino, the matrix of coupling constants, G_S , is given by $\text{diag}(0, 1)$ since the interaction is only seen by the sterile neutrinos. Now using the relationship between the Hamiltonian and the potential from Equation (2.29), the potential is

$$V_{z,\text{int},\nu_s} \propto \frac{G_X p}{M_X^2} u_{\nu_s}. \quad (5.2)$$

Here, u_{ν_s} is the sterile neutrino energy density, and the proportionality constant depends on the definition of G_X . Notice that the contribution to V_z is positive as opposed to the effect from the Z-boson which is always negative. This is not a reflection of differences between the nature of the two vector bosons, but simply a consequence of our definition of \mathbf{V} in terms of Pauli matrices.

In the paper, we introduce this new potential to solve the QKEs for the interacting sterile neutrinos. We show that it is possible to suppress the growth of N_{eff} due to

sterile neutrinos, and we discuss the impact of neutrino oscillation on the active neutrino distributions. We also demonstrate that the mechanism is not crucially dependent on the neutrino oscillation parameters, and we explain an interesting observation of the sterile neutrino contributing more to N_{eff} than one active neutrino does. The introduction section has been removed compared to the published version, and the link to **LASAGNA** has been updated.

5.2 \wp How self-interactions can reconcile sterile neutrinos with cosmology

Steen Hannestad, Rasmus Sloth Hansen and Thomas Tram [3]

Short baseline neutrino oscillation experiments have shown hints of the existence of additional sterile neutrinos in the eV mass range. However, such neutrinos seem incompatible with cosmology because they have too large an impact on cosmic structure formation. Here we show that new interactions in the sterile neutrino sector can prevent their production in the early Universe and reconcile short baseline oscillation experiments with cosmology.

Scenarios

We are considering a hidden gauge boson with mass M_X , and we take the mass to be $\gtrsim 100\text{MeV}$ such that we can use an effective 4-point interaction for all temperatures of interest. The interaction strength is then written as

$$G_X \equiv \frac{g_X^2}{M_X^2}. \quad (5.3)$$

We will assume a 1+1 scenario, specifically a muon neutrino (or tau neutrino) and one sterile neutrino species, a simplification which does not qualitatively alter any of our findings. The system can then be fully characterized by a momentum dependent, 2×2 Hermitian density matrix $\rho(p)$. Since we are not assuming any lepton asymmetry, the evolution of the anti-particle density matrix is trivial, since $\rho(p) = \bar{\rho}(p)$. We expand the density matrix in terms of Pauli matrices:

$$\rho = \frac{1}{2}f_0(P_0 + \mathbf{P} \cdot \boldsymbol{\sigma}), \quad (5.4)$$

where $f_0 = (e^{p/T} + 1)^{-1}$ is the Fermi-Dirac distribution and $\boldsymbol{\sigma}$ is a vector consisting of the three Pauli matrices. The evolution equations for P_0 and \mathbf{P} are called the quantum kinetic equations (QKE), and they were first derived in [45, 56, 92, 93] (for a presentation closer to the present one, see [46, 48]). It is convenient to form the linear combinations:

$$P_a \equiv P_0 + P_z = 2\frac{\rho_{aa}}{f_0}, \quad P_s \equiv P_0 - P_z = 2\frac{\rho_{ss}}{f_0}, \quad (5.5)$$

which separates the sterile and the active sector. The equations of motions are then given by

$$\dot{P}_a = V_x P_y + \Gamma_a \left[2\frac{f_0}{f_0} - P_a \right], \quad (5.6a)$$

$$\dot{P}_s = -V_x P_y + \Gamma_s \left[2\frac{f_{\text{eq},s}(T_{\nu_s}, \mu_{\nu_s})}{f_0} - P_s \right], \quad (5.6b)$$

$$\dot{P}_x = -V_z P_y - DP_x, \quad (5.6c)$$

$$\dot{P}_y = V_z P_x - \frac{1}{2}V_x(P_a - P_s) - DP_y. \quad (5.6d)$$

5.2 \mathfrak{P} How self-interactions can reconcile sterile neutrinos with cosmology

The Γ_s -term is an approximation to the full scattering kernel which is valid in the limit of strong coupling. The sterile equilibrium distribution: $f_{\text{eq},s}(T_{\nu_s}, \mu_{\nu_s}) = (e^{(p-\mu_{\nu_s})/T_{\nu_s}} + 1)^{-1}$, where T_{ν_s} and μ_{ν_s} are the sterile neutrino temperature and pseudo-chemical potential respectively, is uniquely determined from the requirement that the interaction must respect energy conservation and number conservation. Γ_a and Γ_s are related to the 4-point interaction constants as

$$\Gamma_a = C_\mu G_F^2 p T^4, \quad \Gamma_s = G_X^2 p T_{\nu_s}^4 n_{\nu_s}, \quad (5.7)$$

where $C_\mu \simeq 0.92$, while n_{ν_s} is the normalized number density of sterile neutrinos, $n_{\nu_s} = \frac{2}{3\zeta(3)T^3} \int p^2 \rho_{ss}(p) dp$. D quantifies the damping of quantum coherence in the system and is approximately half of the scattering rates, $D \simeq \frac{1}{2}(\Gamma_a + \Gamma_s)$. We have chosen to define Γ_s in analogy with Γ_a , and this means that we do not have exact conservation of ΔN_{eff} for the scattering term in Eq. (5.6b) since Γ_s depends on p . However, none of the results change significantly when we let $p = 3.15T$ in the expression for Γ_s .

In order to include the sterile neutrino self-interaction, we repeat the derivation in [56] for the self-interaction due to the Z -boson in the active sector, but now for an X -boson in the sterile sector. This gives an addition to the matter-potential V_z . The potentials are now

$$V_x = \frac{\delta m_s^2}{2p} \sin 2\theta, \quad (5.8a)$$

$$V_z = V_0 + V_a + V_s, \quad (5.8b)$$

$$V_0 = -\frac{\delta m_s^2}{2p} \cos 2\theta, \quad (5.8c)$$

$$V_a = -\frac{14\pi^2}{45\sqrt{2}} p \left[\frac{G_F}{M_Z^2} T_\gamma^4 n_{\nu_a} \right], \quad (5.8d)$$

$$V_s = +\frac{16G_X}{3\sqrt{2}M_X^2} p u_{\nu_s}. \quad (5.8e)$$

Here δm_s^2 is the mass difference, θ is the vacuum mixing angle, M_Z is the mass of the Z -boson, M_X is the mass of the boson mediating the new force, and u_{ν_s} is the physical energy density of the sterile neutrino. We solve the system of equations using a modified version of the public code LASAGNA [1] available at <https://github.com/ThomasTram/LASAGNA>.

Results

In Figure 5.1 we show the degree of thermalization of the sterile neutrino, quantified in terms of the total energy density in the active plus sterile sector,

$$N_{\text{eff}} \equiv \frac{u_{\nu_a} + u_{\nu_s}}{u_{\nu_0}}, \quad u_{\nu_0} \equiv \frac{7}{8} \left(\frac{4}{11} \right)^{4/3} u_\gamma. \quad (5.9)$$

We have chosen $g_X = 0.1$ and a sample of values for G_X , and we show how ΔN_{eff} develops with the decreasing temperature. We can see that the thermalization of the

5 Sterile interactions




Figure 5.1: The evolution of ΔN_{eff} as the temperature drops for $g_X = 0.1$ and different values of the coupling constant G_X .




Figure 5.2: Contours of equal thermalization. ΔN_{eff} is given by the colors. The solid, dashed, and dot-dashed lines correspond to hidden bosons with masses $M_X = 300$ MeV, 200 MeV, and 100 MeV respectively.

sterile neutrino moves to lower temperatures when the interaction becomes stronger, and this is what we would expect since a strong interaction means that even a small background of sterile neutrinos can prevent further thermalization.

The amount of thermalization depends on both g_X and G_X , and in Figure 5.10 we show ΔN_{eff} as a function of both. It shows that thermalization can be almost completely blocked by the presence of the new interaction for high values of G_X and low values of g_X .

Another interesting observation is that the degree of thermalization depends almost entirely on the mass of the new boson, M_X , not on the dimensionless coupling g_X . This can be understood qualitatively from the following simple argument: At high temperature the production of sterile neutrinos is suppressed by rapid scattering (the quantum Zeno effect), but as soon as production commences the thermalization rate of

5.2 \mathfrak{P} How self-interactions can reconcile sterile neutrinos with cosmology




Figure 5.3: Dependence of ΔN_{eff} on the mixing parameters. $g_X = 0.01$ has been used for all the models while G_X has been changed to give the variation in mass. The standard parameters are $\sin^2(2\theta) = 0.05$ and $\delta m^2 = 1\text{eV}$.

a sterile neutrino can be approximated by $\Gamma_t \sim \Gamma \sin^2(2\theta_m)$, where Γ is the rate with which “flavor content” (in this context meaning active vs. sterile) is measured by the system, and θ_m is the in-medium mixing angle (see e.g. [94, 95] for a discussion of this in the context of active neutrinos). Γ is entirely dominated by the interaction via X so that $\Gamma \propto G_X^2$ and the in-medium mixing angle is likewise dominated by the potential generated by the new interaction so that $\sin^2(2\theta_m) \propto 1/V_s^2 \propto M_X^4/G_X^2$ leading to the sterile thermalization rate being proportional to M_X^4 , i.e. Γ_t does not depend on g_X , only on M_X .

The determination of mixing parameters from accelerator experiments is quite uncertain, and it is therefore interesting to know how our results would be affected if we changed the vacuum mixing angle or the mass difference. The results of such a variation are seen in Figure 5.3. Regarding the ability to inhibit thermalization, the results do not change much. A somewhat higher or lower mass will be needed for the hidden boson, but $\Delta N_{\text{eff}} = 0.6$ can for example be reached by using $M_X = 100$ MeV even at $\delta m^2 = 10$ eV². There are, however, two other interesting observations. First, note that when the hidden boson mass is high, ΔN_{eff} decreases with decreasing $\sin^2(2\theta)$ or δm^2 - the well known limit for non-interacting sterile neutrinos (see e.g. [48, 58]). As the boson mass is lowered, the new interaction first permits full thermalization of the sterile neutrino before we reach the mass range where the new interaction inhibits the thermalization.

The other interesting observation is that $\Delta N_{\text{eff}} > 1$ for some values of M_X . At first this seems very puzzling and counterintuitive. In a model with only oscillations and no new interactions this would be impossible since the number density and energy density of the sterile neutrinos could never exceed the densities of the active neutrinos, the net production of steriles would simply shut off as soon as $\rho_{ss} \sim f_0$. However, in the model presented here there are two effects at play simultaneously: The production of steriles due to oscillations and the redistribution of sterile states due to the new interaction. If the redistribution of energy is sufficiently fast it can keep $\rho_{ss} < f_0$, allowing for




Figure 5.4: The sterile energy distribution relative to f_0 at $T = 4.3$ MeV, where ΔN_{eff} crosses 1 for $\delta m^2 = 1$ eV², $\sin^2(2\theta) = 0.05$, $G_X = G_F$, and $g_X = 0.01$ which corresponds to $M_X = 2.9$ GeV.

more production of steriles. Figure 5.4 provides an illustration of the effect by showing a snapshot of the distributions at the point where ΔN_{eff} crosses 1 for a model with $M_X = 2.3$ GeV. Sterile neutrinos are still being produced in the region close to the resonance at $p/T \approx 5$ since $\rho_{ss} < f_0$ and oscillations therefore populate sterile neutrinos from the active sector. At the same time ρ_{ss} continues to grow at lower p/T due to the redistribution of states. In total this means that ΔN_{eff} is still growing and will do so until the resonance has moved to very high p/T where f_0 becomes very small or the active neutrinos decouple from the electrons. Naively we would expect ΔN_{eff} to be highest for low values of M_X because the energy redistribution becomes more efficient. However, when M_X is decreased the suppression of oscillations due to the effect of M_X on the matter potential quickly wins and ΔN_{eff} decreases rapidly with decreasing M_X . Therefore $\Delta N_{\text{eff}} > 1$ can only occur in a limited transition region of M_X if it occurs at all (which depends on the mixing parameters, δm^2 and $\sin^2(2\theta)$).

Finally, we again stress that our treatment is only consistent if $M_X \gg T$ for any temperature relevant to our calculation. For the typical mass differences favoured by SBL measurements the production of sterile neutrinos takes place at temperatures well below 100 MeV and we have taken this as a representative minimum mass for the new boson. Note that such a low mass might be excluded for a boson coupling to the active sector [86]¹. However, provided that the coupling is diagonal in “flavor” such that X couples only to the sterile state, such bounds are irrelevant. We also note that in the case where G_X becomes extremely high free-streaming of sterile neutrinos will be inhibited and structure formation bounds changed. However, such effects require G_X to be extremely large, $G_X \gtrsim 10^7 G_F$ (see e.g. [88, 91]).

¹It should be noted that the bounds quoted in [86] are based on the assumption that the interaction can be treated as a 4-point interaction at $E \sim m_Z$. Since we are looking at very low mass bosons this assumption does not hold and the bound is therefore expected to be much looser.




Figure 5.5: The active neutrino distribution for different temperatures. The parameters used are $G_X = 3 \cdot 10^2 G_F$ and $g_X = 0.025$. This corresponds to a hidden boson with the mass $M_X = 424$ MeV.

Big Bang Nucleosynthesis (BBN)

Apart from the additional energy density in the sterile sector the oscillations can have another important effect, namely a distortion of the active neutrino distribution. This can happen even after neutrino decoupling because energy can still be transferred between the active and sterile sectors after the active neutrino decouples from the plasma. In models where the active-sterile conversion is delayed, such as the one presented here or models with a non-zero lepton asymmetry [55] this can in certain cases be the dominant cosmological effect. The reason is that the electron neutrino takes part in the nuclear reaction network relevant for Big Bang Nucleosynthesis (see e.g. [55]). Even if the sterile neutrino mixes primarily with ν_μ or ν_τ , active-active oscillations will transfer part of the distortion to the electron sector. However, a detailed investigation of this effect is beyond the scope of the present paper and here we simply point out that interesting effects on BBN might occur. For illustration we show in Figure 5.5 how the active distribution can vary as a function of temperature compared to its unperturbed state, f_0 .

Discussion

We have demonstrated that additional self-interactions of a sterile neutrino can prevent its thermalization in the early Universe and in turn make sterile neutrinos compatible with precision cosmological observations of structure formation. Arguably the model discussed here is more natural than invoking a non-zero lepton asymmetry, relying only on the sterile sector possessing interactions similar to those in the standard model. In order for the model to work the new gauge boson mediating the interaction must be significantly lighter than M_Z , but can easily be heavy enough that no significant background of such particles can exist at late times. We finally note that if this scenario is indeed realized in nature, future precise measurements of N_{eff} will effectively pinpoint the mass of the hidden gauge boson. In summary, the framework presented here

5 Sterile interactions

presents a natural way of reconciling short baseline neutrino experiments with precision cosmology.

Acknowledgments

We thank Georg Raffelt for valuable comments on the manuscript. Note added: Shortly after the submission of this manuscript, another paper [96] on the same topic appeared.






Figure 5.6: BBN exclusions for the vector boson mediator.

The lines indicate the regions allowed at each number of σ for $D/H = (2.53 \pm 0.04) \times 10^{-5}$ and the best fit baryon density from Planck, $\Omega_B h^2 = 0.02207$. From Figure 6 in [97].

5.3 BBN bounds and coupling to dark matter

Since our paper was published along with one by Dasgupta and Kopp [96], several other groups have continued to work on similar models. One of the issues that we left unfinished was the question of the influence on BBN, which has since been addressed by Saviano et al. [97]. In their paper, they consider a three neutrino model, with two active and one sterile neutrino. They verify our result from Figure 5.10, and also confirm the claim that BBN gives rise to additional constraints. These constraints both come from the abundance of ${}^4\text{He}$, but also from deuterium measurements.

The abundance of ${}^4\text{He}$ is mainly determined by the n/p ratio after freeze out as almost all neutrons end up in ${}^4\text{He}$ during BBN. The ratio is determined by chemical equilibrium resulting in a ratio $n/p \sim \exp(-(m_p - m_n)/T)$ until freeze out, which means that a higher freeze out temperature will result in a higher n/p and vice versa. The effect from sterile neutrino oscillations on ${}^4\text{He}$ is twofold. An increase in N_{eff} corresponds to a higher energy density of the Universe which leads to faster expansion and therefore a higher freeze out temperature. Similarly, a deficit of electron neutrinos would lead to lower reaction rates which would also result in a higher freeze out temperature. In conclusion we expect to find a higher helium abundance when including the sterile neutrinos. However, if the oscillations are inefficient until the n/p ratio has settled at $T \sim 0.8\text{MeV}$, the effect on ${}^4\text{He}$ abundance is very small. When comparing the calculations to observations, it turns out that a large observational error means that the constraints are very weak. On the other hand, the limits from D are much stronger, and




Figure 5.7: Resolving cold dark matter problems.

Various constraints on the dark matter mass and the vector boson mass. The grey region to the right indicates a too large production of the sterile neutrino, while the two bands in the middle represent solutions to the small scale problems. Finally, the region in the lower left corner is excluded by cluster observations. From Figure 2 in [96].

in Figure 5.6, we see that only models with very light vector bosons corresponding to values of g_X below what is represented in the plot might be allowed. Unfortunately, the effects from sterile neutrino oscillations are much harder to understand for D, although an increase in N_{eff} in general results in a larger D/H ratio [97]. Although the limits from deuterium are quite strong, they can be alleviated somewhat by using a higher value of the baryon density, and even the variance within the limits from Planck [21] can move the limits by more than 1σ .

Another interesting direction of research is to couple the vector boson to dark matter. This idea was already presented by Dasgupta and Kopp [96], but has also been investigated by Bringmann et al. [98] as well as by Ko and Yang [99]. The general idea is that self-interacting dark matter can solve some of the discrepancies between simulations and observations of galaxies. The three relevant problems are the lack of small satellite galaxies of the Milky Way (*missing satellites*) [22], the existence of large dark matter subhaloes in simulations of Milky Way-like galaxies which should have formed luminous dwarf galaxies that are not seen in observations (*too big to fail*) [23], and the final issue regards the discrepancies between the simulated core profile of dwarf galaxies and the observations of rotation curves (*cusp vs. core*) [24].

It turns out that the interaction obtained by introducing a new vector boson has the right form to conceivably solve all these problems. First of all, the cross section is velocity dependent which allows the interaction to solve problems on galactic scale while avoiding bounds from galaxy clusters. The cusp vs. core and too big to fail problem can both be resolved by having self-interacting dark matter. The cuspy cores from simulations are smoothed out by dark matter scattering with other dark matter

particles, and apart from explaining the cored centers of observed dwarf galaxies, it also explains why the simulated subhaloes from the too big to fail problem never reach a core density high enough to support dwarf galaxies. The missing satellites problem requires a somewhat different solution, where the sterile neutrinos become important. When the dark matter is coupled to light sterile neutrinos, this efficiently delays the time of dark matter decoupling and thereby washes out structures at the smallest scales. This turns out to be possible with the parameters that can also explain both the eV sterile neutrino and the other small scale problems as it is seen in Figure 5.7.

Recently there have been new developments in the understanding of the small scale problems which we will discuss further in Section 5.7 after the pseudoscalar mediator has also been considered.

5.4 Late time production

The introduction of a sterile interaction can push the production of sterile neutrinos to the epoch after neutrino decoupling and thereby avoid large ΔN_{eff} as we have seen. It is even possible to choose such a low mediator mass that BBN is unaffected, but eventually the temperature will become so low that the T^5 -scaling of V_s will be dominated by the vacuum term which scales as T^{-1} . Once this happens, we get the vacuum mixing angle $\sin^2(2\theta) \sim 0.1$ which will lead to production of sterile neutrinos if they also have a reasonable scattering rate [69]. This effect was pointed out by Mirizzi et al. [70] who also noted that such a population of sterile neutrinos with a mass of $\sim 1\text{eV}$ could be in conflict with mass bounds from cosmology, where the latest Planck results give $\sum m_\nu < 0.195$ at 95% confidence [21].

The production rate of sterile neutrinos can be estimated as $\frac{1}{2} \sin^2(2\theta_m) \Gamma_{\text{collision}}$ which was also pointed out in Equation (3.29). The matter mixing angle θ_m can be calculated from the total matter potential, and the final production rate should be compared to the Hubble expansion rate. This is done in Figure 5.8, where the rate has been calculated using the largest of the sterile matter potential from Equation (5.8e) and the vacuum potential. For the collision rate, the expression from Equation (5.7) was used, and both the relative number and energy density of the sterile neutrinos were assumed to be 0.1.

The three cases in Figure 5.8 correspond to; the lower right corner of Figure 5.10 (blue solid); the region which solve the small scale problems that Dasgupta and Kopp identified in Figure 5.7 (green dashed); and some intermediate parameters (red dot-dashed). The two first cases give rise to eras where the production rate is above the Hubble expansion rate which means that a sizeable sterile neutrino population should be expected in both cases just as pointed out by Mirizzi et al. [70]. However, it is also worth noting that it is quite easy to come up with parameters where this production mechanism is not efficient as the red dot-dashed line shows, although these parameters might not solve the small scale problems for cold dark matter. An interesting property of the production rates is, that the curve on the right side of the maximum only depends on the vector boson mass, M_X , as it was also pointed out just below Equation (5.9), while the left part of the curve only depends on G_X since the collision rate in Equation (5.7) is independent of M_X . An important ‘‘aber dabei’’ is that this does not take resonant production into account which will happen when $V_s = V_0$ at the top of each curve in




Figure 5.8: Relative sterile neutrino production rate.

Production rate divided by the Hubble expansion rate at different temperatures. The blue solid line corresponds to the lower right corner of Figure 5.10, the green dashed line represents the region where all small scale cold dark matter problems could be resolved in Figure 5.7, while the red dot-dashed line demonstrates that parameters exist which give no non-resonant production.

Figure 5.8. The resonant production will give rise to an additional contribution to the sterile neutrino population if the resonance is wide enough to be efficient, and while it has not been attempted to include this effect for the late time production for eV sterile neutrinos, we will consider the effect for sterile neutrino dark matter in Section 5.8. One might now think that the late time production of sterile neutrinos means that the coupling between eV sterile neutrinos and dark matter cannot resolve all the small scale cold dark matter problems, but it turns out that an effect has been omitted.

This was recently pointed out by Chu, Dasgupta and Kopp [71]. The problem is that when discussing the mass bound on self-interacting sterile neutrinos, it is not appropriate to use the mass bounds derived for weakly interacting active neutrinos. When the vector boson mass becomes small enough to suppress the sterile neutrino production well after BBN, it might also become strong enough to limit the free-streaming of the sterile neutrinos at later times, and hence give a significantly smaller effect on structure formation. An estimate of the effects shows that the free streaming scale can be decreased by a factor of ~ 5 for parameters representative of the region that resolve small scale dark matter problems [71]. This is enough to give a much better agreement between observations and the self-interacting model compared to non-interacting models.

However, at least one of the mainly active neutrino mass states will also be affected due to the mixing if the self-interaction becomes too large. This will give rise to a lack of free-streaming compared to the standard Λ CDM model, and although most analyses of this effect have only been considering interactions for all three active neutrinos [88, 91], it is questionable if the observations would allow the absence of free-streaming for one of the active neutrinos. The possibly most relevant analysis for this question is the one presented in Section 5.6, and even from that, the answer is not clear.

When considering the sweet-spot where free-streaming is suppressed for the sterile neutrino while it is still present for all the active neutrinos, it turns out that it coincides with the parameters that can solve the small scale dark matter problems. Due to the lucky coincidence, this model is the most elegant way to avoid the large scale structure constraints, although various other models involving an even higher number of sterile neutrinos have also been suggested [71, 100].

5.5 Pseudoscalar boson mediator

The vector bosons that were discussed in the previous sections have a very solid motivation in the Standard Model of particle physics, were we find several fundamental examples of similar particles in form of the photon as well as the weak bosons. However, the Standard Model also presents examples of pseudoscalar force carriers, although at a slightly less fundamental level. While the strong nuclear force is mediated by gluons when discussing the inner workings of hadrons, it is mainly pions that mediate the force between protons and neutrons in nuclei, and these are prime examples of pseudoscalars. The pions can be interpreted as almost massless pseudo-Nambu-Goldstone bosons from the approximate chiral symmetry of the strong interaction, and various extensions to the standard model can create pseudoscalars by either similar or different mechanisms. The most interesting pseudoscalar in neutrino physics is the majoron which arises from the spontaneous breakdown of lepton number [101] and is often used in neutrino masses models and in models for neutrino-less double beta decay. As these symmetry breaking models are the motivation, we will only consider very light pseudoscalars with a mass say well below the eV scale.

One could also wonder why a parity even scalar particle is not considered, and the effort is concentrated on the parity odd pseudoscalar particle. In principle, a scalar particle could have an effect similar to the pseudoscalar mediator, but scalar particles couple to energy whereas pseudoscalars couple to spin, and while the energy of a macroscopic object is significant, the overall spin is in general very small. Therefore a scalar mediator would give rise to long range forces and consequently, the limits on scalar mediators are much stronger than the limits on pseudoscalars.

Compared to the vector boson mediator, the interaction mediated by the pseudoscalar has the interesting property, that it recouples rather than decouples at late times. This means that the sterile neutrinos become very strongly self-interacting at late times, and in the next section we will see that this has interesting consequences not only for the evolution of N_{eff} in presence of an eV sterile neutrino, but also for bounds from observations of CMB anisotropies.

The aim of the paper in the next section (Section 5.6) was to investigate the consequences of an interaction for the sterile neutrinos mediated by a pseudoscalar.

5 Sterile interactions

Originally, the expectation was that the overall phenomenology would be very similar to the vector boson case, and that the main effect of the interaction would be to suppress the production of the sterile neutrinos. However, as the work progressed, it became clear that the late time phenomenology of the pseudoscalar case was very different from the vector model that we had investigated, and that these late time effects were playing an pivotal role with some analogies to the later findings of Chu, Dasgupta and Kopp in the vector boson model [71].

While the potential does indeed suppress sterile neutrino production if it is large enough, the effects of an interacting component around the time of recombination are more interesting. It turns out that the combination of an increased N_{eff} and the lack of free-streaming for the sterile neutrino gives rise to a quite good fit to CMB observations, and when combined with the locally measured Hubble parameter, we get an even better fit than Λ CDM. This is an interesting observation, although it does by no means indicate that the pseudoscalar model is more plausible than Λ CDM. One could e.g. have included observations of baryon acoustic oscillations and the Lyman α forest which would have favoured the Λ CDM over the pseudoscalar model, but such an extensive comparison was beyond the scope of our paper. Apart from the over all good fit, we find that a value of N_{eff} above 3 is preferred. This means that the suppression from the potential should not be too large, and combined with other constraints, it puts the coupling constant around $10^{-5} - 10^{-6}$. When considering a similar interaction for dark matter which must have some effect on galactic dynamics, we find that the coupling constant should have a similar or lower value, although that part of our analysis is very crude.

Compared to the published paper [4], a few pages have been omitted. These are the introduction and parts of the numerical details which has been shortened to avoid a repetition of Section 5.2. Also, the section on Sommerfeld enhanced scattering is removed as it is quite technical and mainly served to prove that this effect can be neglected for our type of model.

5.6 \wp Cosmology with self-interacting sterile neutrinos and dark matter - A pseudoscalar model

Maria Archidiacono, Steen Hannestad, Rasmus Sloth Hansen, Thomas Tram [4]

Short baseline neutrino oscillation experiments have shown hints of the existence of additional sterile neutrinos in the eV mass range. Such sterile neutrinos are incompatible with cosmology because they suppress structure formation unless they can be prevented from thermalising in the early Universe or removed by subsequent decay or annihilation. Here we present a novel scenario in which both sterile neutrinos and dark matter are coupled to a new, light pseudoscalar. This can prevent thermalisation of sterile neutrinos and make dark matter sufficiently self-interacting to have an impact on galactic dynamics and possibly resolve some of the known problems with the standard cold dark matter scenario. Even more importantly it leads to a strongly self-interacting plasma of sterile neutrinos and pseudoscalars at late times and provides an excellent fit to CMB data. The usual cosmological neutrino mass problem is avoided by sterile neutrino annihilation to pseudoscalars. The preferred value of H_0 is substantially higher than in standard Λ CDM and in much better agreement with local measurements.

Model framework

Instead of constructing an explicit model we base our discussion on a simplified setup which, however, does contain all the relevant physics. The sterile neutrino is coupled to a new light pseudoscalar with mass $m_\varphi \ll 1\text{eV}$ via

$$\mathcal{L} \sim g_s \varphi \bar{\nu} \gamma_5 \nu. \quad (5.10)$$

Later we will look at dark matter with a similar coupling to φ

$$\mathcal{L} \sim g_d \varphi \bar{\chi} \gamma_5 \chi. \quad (5.11)$$

One important note is in order at this point: We assume the coupling to be diagonal in mass basis, such that the 3 mainly active mass states are completely uncoupled. This is the most natural assumption given that φ is associated with new physics and not related to standard model flavour. The new interaction is also felt partly by the active Standard Model neutrinos, although suppressed by the mixing angle. Limits from cosmology [88] are not relevant, as the active neutrino mass states do not feel the new coupling, but constraints from supernovae [87, 102] and laboratory measurements [103] do apply. The supernova bounds are derived by requiring that the pseudoscalars do not carry away a significant amount of the energy released by the supernova which results in a bound on the coupling of electron neutrinos to the pseudoscalar [87], $g_e \lesssim 4 \times 10^{-7}$. If the coupling becomes much larger, the pseudoscalars will be caught in the supernova, and the bound disappears again. However, almost all of these values are excluded by laboratory experiments [103], and we will only consider the supernova limit here. For the sterile neutrinos, the bound on g_e comes from the process $\nu_e \nu_e \rightarrow \varphi$, and it translates into the bound $g_s \lesssim g_e / \sin^2 \theta_s = 3 \times 10^{-5}$, using $\sin^2 2\theta_s \sim 0.05$ from the short baseline experiments [104, 105], where θ_s is a mixing angle representative for

5 Sterile interactions

(ν_e, ν_s) mixing or (ν_μ, ν_s) mixing. Although supernovae give the strongest bounds on the coupling strength, they are quite dependent on details in the assumptions about the supernova, and it might be more appropriate to quote the bound as $g_s \lesssim 10^{-4}$.

Let us now go through the implications of this new interaction, first for the sterile neutrinos and subsequently for the dark matter.

Sterile neutrinos

The new interaction introduces a matter potential for sterile neutrinos of the form [89, 90]

$$V_s(p_{\nu_s}) = \frac{g_s^2}{8\pi^2 p_{\nu_s}} \int p dp (f_\varphi + f_{\nu_s}), \quad (5.12)$$

where f_φ is the Bose-Einstein distribution for the pseudoscalar and f_{ν_s} is the distribution for the sterile neutrinos (see e.g. [45, 56, 58, 93, 94] for a discussion of matter potentials in the standard model). Note that the potential in Eq. (5.12) arises from bubble diagrams and is non-zero even in a CP-symmetric medium.

Before proceeding with a quantitative calculation we can estimate how large g_s needs to be in order to block thermalisation. Consider a scenario with thermal φ and ν_s distributions characterised by a common temperature T . The potential is then

$$V_s \sim 10^{-1} g_s^2 T. \quad (5.13)$$

In the absence of non-standard effects, the sterile neutrinos would be thermalised through oscillations at $T \sim 10(\delta m^2/\text{eV}^2)^{1/6} \text{MeV} \sim 10 \text{MeV}$ [58]. To prevent this effect, we need to suppress the mixing angle in matter, θ_m as the production rate is proportional to $\sin^2 2\theta_m$. This is achieved if the matter potential dominates the energy difference associated with vacuum oscillation, i.e.

$$V \gtrsim \frac{\delta m_{\nu_s}^2}{2E} \sim \frac{\delta m_{\nu_s}^2}{T}, \quad (5.14)$$

prior to neutrino decoupling at $T \sim 1 \text{MeV}$ so that

$$g_s^2 \gtrsim 10 \frac{\delta m_{\nu_s}^2}{T^2} \sim 10^{-11}. \quad (5.15)$$

So a priori we expect that a value of $g_s \sim 3 \cdot 10^{-6}$ is sufficient to block thermalisation. It should be noted here that since the pseudoscalar coupling is diagonal in mass basis the active state feels an additional matter potential associated with the φ background. The magnitude of the potential felt by the active state is approximately $V \sim \sin^2(\theta_s) V_s \sim 0.01 V_s$. The only effect is a minute shift in the effective mass difference, corresponding to a shift of less than one percent in g_s .²

²For active-active oscillations this additional potential is important, but effects from active-active oscillations are expected to be small as all active neutrinos have almost identical spectra in the early Universe.

Thermal history of the sterile neutrino

The sterile neutrino can in principle be thermalised via incoherent processes such as $\varphi\varphi \leftrightarrow \bar{\nu}_s\nu_s$, assuming that there is a pre-existing background of φ . The thermally averaged cross section in the highly relativistic limit can be calculated to be [106]

$$\langle\sigma|v|\rangle = \frac{g_s^4}{8\pi T^2}. \quad (5.16)$$

Conservatively assuming that $g_s \sim 10^{-4}$ we find that ν_s and φ come into equilibrium at a temperature of $T \sim 1$ GeV, i.e. significantly before the oscillation process becomes important [54]. However, since the dark sector is decoupled it does not share the entropy transfer to the standard model particles, and the end result is that when oscillations become important at $T \sim 10$ MeV, a low-temperature background of φ and ν_s exists. However if g_s is significantly lower no thermalisation occurs before the oscillation period.

Results and numerical implementation

We compute the thermalisation process by solving the Quantum Kinetic Equations (QKEs) for a simplified two-neutrino framework with oscillations between ν_μ and ν_s using a modified version of our public code LASAGNA [1]. The formulation of the QKEs [45, 46, 48, 56, 58, 92, 93] is based on an expansion of the density matrices, ρ , in terms of P_a , P_s , P_x , and P_y just as in Section 5.2. The only difference between the treatment there is that we use the sterile potential from Equation (5.12) and the sterile scattering rate

$$\Gamma_s = \frac{g_s^4}{4\pi T_{\nu_s}^2} n_{\nu_s}. \quad (5.17)$$

We compute the sterile neutrino contribution to the potential in Eq. (5.12) from the actual numerical distribution. The contribution from the φ -background is computed analytically assuming that the φ -particles were produced thermally above a TeV. They will then follow a Bose-Einstein distribution with a reduced temperature of

$$T_\varphi = \left(\frac{g_\star(T_\gamma)}{g_\star(1\text{TeV})}\right)^{\frac{1}{3}} T_\gamma \simeq \left(\frac{10.75}{106.7}\right)^{\frac{1}{3}} T_\gamma \simeq 0.47T_\gamma, \quad (5.18)$$

where the approximation is valid in the temperature range of interest. We are ignoring momentum transfer between the sterile neutrinos and the pseudoscalars for simplicity, but we suspect that including it would have a negligible effect on our results. When sterile neutrinos are produced, they will create non-thermal distortions in the sterile neutrino distribution, and the sterile neutrino spectrum might end up being somewhat non-thermal. In Fig. 5.9 we show the final contribution to the energy density N_{eff} from a sterile neutrino with mixing parameter $\sin^2 2\theta_s = 0.05$ and $m_{\nu_s} = 1$ eV, close to the best fit value from neutrino oscillation data [104, 105]. The transition from full thermalisation to zero thermalisation happens in the region $10^{-6} < g_s < 10^{-5}$, confirming the simple estimate in Eq. (5.15)³.

³Note that in the absence of a pre-existing population of φ and ν_s , sterile neutrino production would still be suppressed for the same values of g_s as soon as a small amount of ν_s has been produced through oscillations. The assumption is thus not crucial to the scenario.




Figure 5.9: The contribution of the sterile neutrino to the relativistic energy density $\Delta N_{\text{eff}} = N_{\text{eff}} - 3$ as a function of the coupling parameter g_s .

Late time phenomenology

In a recent paper by Mirizzi et al. [70] it was pointed out that even if strong self-interactions prevent thermalisation of the sterile neutrino before active neutrino decoupling it will eventually be almost equilibrated by oscillations at late times. This leads to a scenario in which active and sterile neutrino distributions have similar temperatures and both contribute to the combined N_{eff} . Even if early thermalisation is prevented this still leads to a sterile neutrino population with a temperature only slightly below that of standard model active neutrinos and therefore the usual cosmological neutrino mass bound still applies to this model.

However, unlike the previously studied Fermi-like interaction, sterile neutrinos and pseudoscalars interact via a variety of $2 \leftrightarrow 2$ processes which in general have a scattering rate of order $\Gamma \sim g_s^4 T$ because there is no mass scale involved. This is true for example for the pair annihilation process $\nu_s \bar{\nu}_s \rightarrow \varphi \varphi$ where we already found the thermally averaged cross section to be $\langle \sigma |v| \rangle = g_s^4 / (8\pi T^2)$ in the relativistic limit, implying a reaction rate $\Gamma = \langle \sigma |v| \rangle n_{\nu_s} \approx 3.6 \times 10^{-3} g_s^4 T$. This should be compared to the Hubble expansion rate $H \sim 10T^2/m_{\text{Pl}}$. As long as $g_s \gtrsim 10^{-6}$ the $\nu_s - \varphi$ plasma becomes strongly self-interacting before the sterile neutrinos become non-relativistic around recombination. Therefore, the rest mass constraint does not apply to this model: As soon as sterile neutrinos become non-relativistic, they annihilate into φ . This annihilation has two immediate effects. It leads to an overall increase in the energy density of the $\nu_s - \varphi$ fluid, and it leads to a temporary decrease in the equation of state parameter for the fluid. Both of these effects were discussed in detail in [107].

The strong self-interactions of the combined fluid also leads to a complete absence of free-streaming and in turn an absence of anisotropic stress in the $\nu_s - \varphi$ plasma. The scenario where all neutrinos are strongly interacting is strongly disfavoured by current data (see e.g. [88, 91, 107–113] for discussions of self-interacting neutrinos and cosmic structure formation). However, this is not necessarily true for models in which standard model neutrinos are free-streaming, and the interaction is confined to the sterile sector.

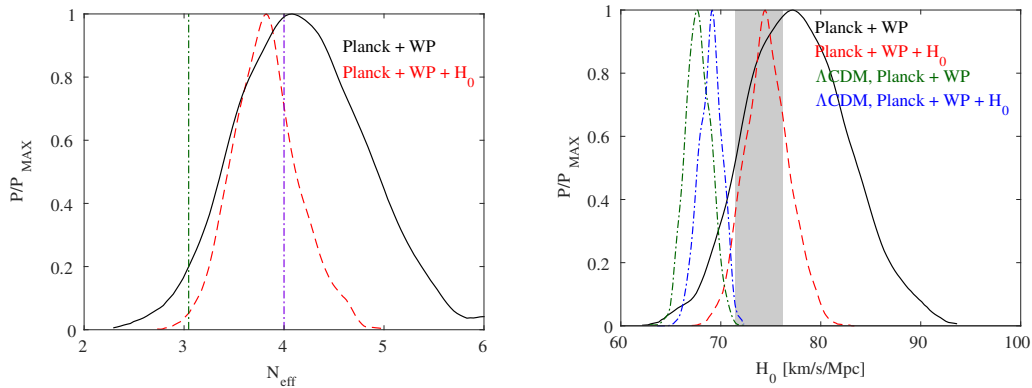


Figure 5.10: 1D marginalised posteriors for N_{eff} (Top panel) and H_0 (Bottom panel) obtained by assuming the pseudoscalar scenario and using only CMB data (black/solid line) and CMB data plus the H_0 prior (red/dashed line). (Top panel) The green dash-dot line refers to the Λ CDM model ($N_{\text{eff}} = 3.046$) and the purple line is the complete thermalization case ($N_{\text{eff}} \simeq 4$). (Bottom panel) The green and the blue dash-dot lines show the posteriors obtained in the Λ CDM model using Planck and Planck+ H_0 , respectively. The H_0 prior is marked by the grey shaded region [116].

We note here that since the pseudoscalar coupling is diagonal in mass basis it does not induce self-interactions in the three active mass states.

We have performed a study of how this model is constrained by current CMB data through an MCMC sampling of the cosmological parameter space performed with CosmoMC [114] and using CMB data from the Planck mission as well as CMB polarisation data from the WMAP satellite [115] (we refer to this data combination as “Planck+WP”). We describe the neutrino sector by the overall energy density after thermalisation, N_{eff} and assume a sterile mass of 1 eV. We assume that all neutrino species and the pseudoscalar equilibrate at some temperature between the thermalisation scale at a few MeV and the CMB scale ($T \sim 1$ eV), so that the energy density in the active sector is $21/32N_{\text{eff}}$ with the remaining $11/32N_{\text{eff}}$ is in the $\nu_s - \varphi$ fluid.

In the top panel of Fig. 5.10 we show the 1D marginalised posterior for N_{eff} for the Planck+WP data, as well as for the same data, but with the direct measurement of H_0 from [116] included. The data shows a clear preference for high values of N_{eff} and the most extreme case with complete thermalisation of the sterile neutrino, corresponding to $N_{\text{eff}} \simeq 4$, is well within the 1σ allowed region. It is also of interest to compare the difference in χ^2 between this model and the standard Λ CDM cosmology. We find that $\Delta\chi^2$ of the pseudoscalar model compared to the reference Λ CDM model is $\Delta\chi^2 = \chi^2_{\text{pseudoscalar}} - \chi^2_{\Lambda\text{CDM}} = 1.3$, while if in our model we assume $N_{\text{eff}} \simeq 4$ then $\Delta\chi^2 = 0.6$.

Interestingly for this model with a subdominant, strongly interacting neutrino sector we also find a preference for a higher value of H_0 . This effect was seen already in [107] but with a much more dramatic increase in H_0 because all neutrinos were assumed to be strongly interacting. In the bottom panel of Fig. 5.10 we show the 1D marginalised

posterior for H_0 for this model as well as for Λ CDM. The increase in H_0 alleviates the tension between the locally measured value of H_0 and the much lower value inferred from Planck data when the standard model is assumed. We see this effect very directly when comparing χ^2 values: $\Delta\chi^2 = \chi^2_{\text{pseudoscalar}} - \chi^2_{\Lambda\text{CDM}} = -2.5$, while if in our model we assume $N_{\text{eff}} \simeq 4$ then $\Delta\chi^2 = -3.9$. We thus find that in this case the model with a strongly interacting $\nu_s - \varphi$ sector is a better fit to current data than Λ CDM (and of course a vastly better fit than Λ CDM with an additional 1 eV sterile neutrino).

Dark matter

We will now investigate the possibility that dark matter also couples to the new pseudo-scalar with a dimensionless coupling strength, g_d . We assume that the dark matter is produced at a very high temperature by e.g. inflaton decay. Once dark matter is coupled to the new interaction, there is the potential worry that it will pair annihilate via the process $\chi\bar{\chi} \rightarrow \varphi\varphi$ with the same cross section as in Eq. (5.16). If the annihilation process is in equilibrium where χ goes non-relativistic, it will dilute the density of χ while transferring an unacceptable amount of entropy to φ . Due to the nature of the interaction, it is decoupled at high temperatures, and the cross section likewise drops when the dark matter becomes non-relativistic. Therefore, we only need to ensure that the dark matter annihilation rate is low enough at $T_{\text{max}} \sim m_\chi$. We assume that the cross section is given by the highly relativistic expression for $\langle\sigma|v|\rangle$ in Eq. (5.16), and use the condition $\Gamma(T_{\text{max}}) = \langle\sigma|v|\rangle n_\chi < H(T_{\text{max}})$ to derive the condition,

$$g_d \lesssim 2 \times 10^{-5} \left(\frac{m_\chi}{\text{MeV}} \right)^{1/4}, \quad (5.19)$$

for the new interaction not to overly dilute the density of χ .

Additionally, the new coupling also induces a Yukawa type potential between the dark matter particles. This in turn leads to dark matter self-interactions which might have observable consequences for galactic dynamics. Rather than going through a detailed calculation we will simply estimate the mean time between dark matter scatterings in order to estimate whether self-interactions are important. In order to do so we will follow the prescription given in [117]. First, following Ref. [118] we write

$$V(r) = -\frac{g_d^2}{m_\chi^2} \frac{e^{-m_\varphi r}}{4\pi r^3} h(m_\varphi r) \mathcal{S}, \quad (5.20)$$

where $h(m_\varphi, r) = 1 + m_\varphi r + \frac{1}{3}(m_\varphi r)^2$ and \mathcal{S} is a spin-dependent factor which we assume to be one.

The interaction potential in Eq. (5.20) causes elastic scattering of dark matter, and following the prescription in [117] we can estimate the value of g_d needed in order to have a significant impact on galactic dynamics. The calculation in [117] was performed for a massless $U(1)$ vector so the potential is Coulomb-like. This in turn leads to both ‘‘soft’’ and ‘‘hard’’ scattering of roughly equal importance. Here we can safely neglect the contribution from soft scatterings because of the steepness of the potential.

The ratio of the scattering time scale $\tau_{\text{scat.}}$ to the dynamical time scale in the galaxy $\tau_{\text{dyn.}}$ is given by Eq. 17 in [117],

$$\frac{\tau_{\text{scat.}}}{\tau_{\text{dyn.}}} = \frac{2R^2}{3N\sigma}, \quad (5.21)$$

where R is the radius of the galaxy, N is the number of DM particles in the galaxy and σ is the scattering cross section. For a hard scatter we have $\sigma \simeq b^2$ where the impact parameter b is the radial distance such that the sum of kinetic and potential energy is zero,

$$\frac{\alpha_d}{m_\chi^2 b^3} = \frac{1}{2} m_\chi v^2, \quad (5.22)$$

where we have used that $m_\varphi b \sim m_\varphi/m_\chi \ll 1$ which leads to the approximation $V(r) \approx -\alpha_d/(m_\chi^2 r^3)$ where $\alpha_d = g_d^2/4\pi$. We then find that

$$\left(\frac{\tau_{\text{scat.}}}{\tau_{\text{dyn.}}} \right)^3 = \frac{2R^4 m_\chi^8 G^2}{27N \alpha_d^2}, \quad (5.23)$$

where G is Newton's constant. The condition for the time scale of scattering to be less than the age of the Universe is⁴ $\tau_{\text{scat.}}/\tau_{\text{dyn.}} \lesssim 50$. Plugging in numbers for a Milky Way size halo and using $\alpha_d = g_d^2/4\pi$, we find

$$g_d \gtrsim 6 \times 10^{-8} \left(\frac{m_\chi}{\text{MeV}} \right)^{\frac{9}{4}}. \quad (5.24)$$

The value of g_d in Eq. (5.24) can be seen as a lower bound on the value required to have a significant effect. The actual value required might be somewhat larger.

In order for elastic scattering to be important in itself the mass of the dark matter particle is therefore required to be quite small. For example, $g_d \sim 10^{-5}$ leads to the requirement that $m_\chi \lesssim 10$ MeV. So depending on the unknown mass of the Dark Matter particle, hard scattering on this potential could have a *direct* impact on galactic dynamics.

Dark acoustic oscillations?

Since our model couples dark matter to a background of dark radiation we might worry that the $\chi - \varphi$ system can undergo acoustic oscillations close to the epoch of recombination and thus distort the observed CMB spectrum (see e.g. [119] for a recent discussion). The interaction around the epoch of CMB formation is primarily Compton scattering, $\chi\varphi \rightarrow \chi\varphi$, and we can directly compare it to the normal Compton scattering rate of photons and electrons. The Compton cross section scales as $\sigma \propto \alpha^2/m^2$ where m is the fermion mass. As long as $g_d^2 \ll \alpha$ and $m_\chi \gg m_e$, the dark sector acoustic oscillations will be completely negligible and therefore cosmologically safe. This of course also means that late-time Compton scatterings can be safely ignored since they have no impact on the ability of χ to cluster gravitationally. Scaling relative to the electron-photon process we can formulate the bound as

$$g_d^2 \ll 1.6 \times 10^{-2} \left(\frac{m_\chi}{\text{MeV}} \right). \quad (5.25)$$

⁴We take $\tau_{\text{dyn.}}$ to be the dynamical time scale of a Milky Way size halo.

Discussion

We have studied a model with secret sterile neutrino interactions mediated by a massless or very light pseudoscalar. The model has some of the same features as the previously studied models based on Fermi-like interactions mediated by heavy vector bosons in the sense that it provides a background potential which can block the production of sterile neutrinos and resolve the apparent inconsistency between cosmology and short baseline neutrino oscillation data.

However, the model has very different late-time phenomenology. The very low mass of the pseudoscalar makes the sterile neutrino strongly self-interacting at late times, an effect which is perfectly consistent with current cosmological data, but might be used to uniquely identify the model once more precise measurements become available. In order to accommodate the mass bound from cosmological large scale structure [70], we need $g_s \gtrsim 10^{-6}$ to allow the sterile neutrinos to annihilate when they become non-relativistic. Our analysis of the CMB suggests $N_{\text{eff}} \approx 4$, and this suggestion is amplified if we also consider the direct measurements of H_0 . At 95% confidence we can rule out $N_{\text{eff}} = 3.046$ when we include the H_0 measurement, and this formally corresponds to an upper limit on g_s of $g_s \lesssim 10^{-5}$ according to Fig. 5.9. However, this bound is very dependent on the set of data we have used, and might both be strengthened and weakened by including more data. We finally arrive at a combined bound on g_s of

$$10^{-6} \lesssim g_s \lesssim 10^{-5} (\text{CMB} + H_0). \quad (5.26)$$

A more robust determination of N_{eff} would allow the possible values for g_s to be further confined, and a precise value of $N_{\text{eff}} > 3.046$ would allow us to pinpoint a corresponding coupling strength. We also note that since the fundamental coupling strength is very low and restricted to the sterile sector in this model it is unlikely to produce observable effects on neutrino physics in general (see e.g. [103] for laboratory constraints). Considering non-standard energy loss from the proto-neutron star in SN1987a also leads to an upper bound on g_s in the $\sim \text{few} \times 10^{-5}$ range (see e.g. [120] for a discussion).

In addition to the coupling to sterile neutrinos we hypothesise that the pseudoscalar also couples to the dark matter particle. Provided that the dark matter particle is sufficiently light this can lead to significant effects on dark matter clustering in galaxies and clusters and possibly resolve some of the apparent discrepancies between the standard Λ CDM model and observations [121]. These discrepancies include the ‘‘Too big to fail’’ problem [23] and the ‘‘cusp vs. core’’ problem (see [24] and references herein), but not the ‘‘missing satellites’’ problem [22] which would require a stronger coupling between neutrinos and DM.

In order for the model to be viable, the dark matter coupling must be sufficiently low that the pair annihilations do not transfer excess entropy to the plasma of sterile neutrinos and pseudoscalars. Conversely, the dark matter coupling must be strong enough to produce an observable effect on galactic dynamics. In Fig. 5.11 we show these two constraints simultaneously and include the bound from warm dark matter [25]. We are left with a viable DM candidate with a mass between few keV and ~ 10 MeV and couplings from 10^{-13} to 10^{-5} . A more detailed treatment of the ‘‘cusp vs. core’’ and ‘‘Too big to fail’’ problems could probably constrain the dark matter further, but that is beyond the scope of this article. The type of dark matter, that we have described, is very different from the normal WIMP cold dark matter. However, it is




Figure 5.11: Constraints in $m_\chi - g_d$ space. The green region is ruled out from Eq. (5.19) due to overproduction of φ -particles from χ -annihilations, while the purple region will have no effect on galactic dynamics, cf. Eq. (5.24).

entirely possible that dark matter consists of an additional sterile neutrino species with extremely suppressed mixing to the active sector. If this is the case it cannot be produced via the usual scattering and oscillation mechanism. However, unlike an MeV sterile neutrino produced via the normal oscillation and scattering mechanism it also remains stable on cosmological timescales. The actual production mechanism for the dark matter particle might be via direct inflaton decay at reheating or from the thermal background at very high temperature.

In summary, sterile neutrino and dark matter interactions via a light pseudoscalar seems an extremely interesting possibility for explaining a variety of different problems in cosmology and certainly merits further study.

Acknowledgments

We thank the referees for comments and suggestions, which significantly contributed to improving the quality of the manuscript. MA acknowledges partial support from the European Union FP7 ITN INVISIBLES (Marie Curie Actions, PITN- GA-2011-289442)



5.7 Current state of sterile interactions and self-interacting dark matter

Our knowledge about self-interacting sterile neutrinos has been described in the past few sections with a few detailed highlights. The problem that these models were brought into the world to solve is still present as the experimental anomalies for both reactor-, accelerator- and gallium neutrinos persist. However, many experiments that will investigate the sterile neutrino solution to the anomalies are being prepared, and some are already running. As a result, this particular chapter of neutrino physics will soon be either closed or kicked wide open.

If the 1eV sterile neutrino is found, and we wish to reconcile it with cosmology by introducing sterile interactions, we now know of a few viable options:

1. The interaction is mediated by a vector boson, X , with mass $m_X \lesssim 10\text{MeV}$.
 - a) $g \lesssim 10^{-6}$. For a low enough g , the sterile neutrinos are never thermalised, but the interaction cannot solve all problems for cold dark matter (CDM) on galactic scales.
 - b) Depending on the mass, $10^{-4} \lesssim g < 1$, and the sterile neutrino is strongly coupled until its free-streaming only has a small effect on large scale structures. The CDM problems on galactic scales might be resolved.
 - c) More light sterile neutrinos are introduced either to redistribute the energy density from the 1eV neutrino or to solve the problems separately.
2. The interaction is mediated by a very light pseudoscalar with a coupling constant $g \sim 10^{-5} - 10^{-6}$. This gives an improved fit for CMB observations combined with measurements of the Hubble constants using supernovae, and the consequences for large scale structures are small as the sterile neutrinos decay to the pseudoscalar once it becomes non-relativistic. Some of the CDM problems on galactic scale might be solved as well.

Due to the sterile nature of the new interaction in these models, they will be hard to test experimentally, and the best option might be to constrain them using cosmological observations. For model 1a and 2, there should be an observable effect of the relativistic component without free-streaming and anisotropic stress, while 1a only distinguishes itself from ΛCDM by a very small population of sterile neutrinos, which might be very hard to measure. Finally, model 1c is not well defined as it is phrased here, and it could have observable consequences depending on the details of the additional sterile neutrinos.

While some effort has been put into understanding if these types of models can have parameters that will solve the too big to fail-, the cusp vs. core- and the missing satellites problems, there have also been a significant progress in the understanding of these problems in the standard ΛCDM cosmology, and recent results suggest that baryon physics might resolve at least some of them.

It is the event of complex hydrodynamical models that have facilitated this improved understanding [122, 123]. By including more realistic gas dynamics and baryonic feedback mechanisms, it has become possible to reproduce a representative population

of various types of galaxies; elliptical, spiral and even some of the more irregular. Also star formation rates, stellar luminosity of galaxy clusters, column density of intergalactic gas, hydrogen large scale statistics, and metal content of galaxies can be reproduced by these new simulations. When taking one of these simulation suites and attempting to reproduce our Local Group of galaxies, it is seen that both the missing satellites and the too big to fail problems are solved [124]. The missing satellites problem is resolved due to the effect of reionization that prevents star production in most of the small dark matter halos. For the too big to fail problem, the effect is more complex, but the inclusion of hydrodynamics gives rise to less massive halos of the dwarf galaxies compared to dark matter only simulations. This is enough to push the number of predicted large dwarfs down to 3 – 4 which is in agreement with what we see in the Milky Way.

The last problem, which was discussed in relation to self-interacting dark matter, is the cusp vs. core problem, which has also received some recent attention. The usual simple statement of the cusp vs. core problem is that simulations produce cores with a cusp in the center, while observations of rotation curves in galaxies indicate a core instead. However, this picture might be too simplistic as some galaxies are observed to have rotation curves which are in good agreement with simulations while others do indeed have rotation curves that deviate significantly from the predictions [125]. As a result, this problem does not call for a solution that will lead to uniformly more cored halo profiles as this would give rise to a tension with the galaxies that are currently well described. Although it would be impossible to say for sure without doing a detailed study, this suggests that an introduction of self-interactions for dark matter would be a too general modification of the Λ CDM cosmology to resolve the cusp vs. core problem.

Although these new results indicate that the necessity of self-interacting dark matter as an explanation of galactic dynamics is fading, there are many other unanswered questions that remain to be answered regarding the elusive substance that represents about 80% of the matter in our Universe, and dark matter physics continue to be an exciting field.

5.8 Dark matter production

In the recent years there have been a large interest in sterile neutrino dark matter, and the observation of an unexpected line in X-ray spectra of the Perseus cluster, the Andromeda galaxy [27], stacked cluster data [26] and the Milky Way center [28] which can be interpreted as a 7keV sterile neutrino has only intensified the attention. Although the signal has been seen in so many different objects, it has also been questioned by several other groups who have been unsuccessful in reproducing a signal that could come from dark matter if they see a signal at all [29–35]. Regardless if the final verdict is that the signal is present and from dark matter or not, sterile neutrinos with a mass above a few keVs are good candidates for dark matter. However, if all the dark matter is sterile neutrinos, it could not be produced by simple non-resonant oscillations [126]. Instead models often use the resonant Shi-Fuller production, where a large lepton asymmetry induces a resonance that can produce the correct abundance of dark matter [85]. In a similar way it is conceivable that the self-interaction of the sterile neutrinos themselves could lead to a significant resonant production even when the non-resonant production

5 Sterile interactions

is inefficient. One of the constraints that can be imposed on sterile neutrino dark matter comes from the large scale structures probed by the Lyman- α forest. As the sterile neutrino is often very light compared to e.g. WIMPs (Weakly Interacting Massive Particles), it will suppress small structures to a larger extent because it is faster for a given a kinetic energy. The current constraints are around $m_{\nu_s, \text{thermal}} \gtrsim 3.3\text{keV}$ [25] although the constraints from different groups sometimes disagree up to a factor of 2. However, if the sterile neutrino is produced with a very non-thermal momentum distribution, the particles will have a comparatively lower velocity, and it can mimic a much heavier dark matter candidate. Therefore, it is a desirable feature of the production mechanism to produce a cold distribution, especially if the considered sterile neutrino mass is below 3.3keV.

In the following we will begin with the possibility of having a vector boson mediator and solve the QKEs for a few representative cases. However, we will also investigate the adiabaticity of the resonance and try to develop a simpler model that can capture some of the relevant physics. Afterwards, we will discuss the pseudoscalar mediator and what differences and similarities we can expect for that case.

5.8.1 Vector boson mediator

To simplify the analysis, we will only consider the case of a 4-point interaction as we did in Section 5.2 with the coupling constant $G_X = g_X^2/M_X^2$, and we will again use the two expressions for the collision rate and the background potential

$$\Gamma_s = G_X^2 p T_{\nu_s}^4 n_{\nu_s}, \quad (5.7) \quad V_s = \frac{16G_X}{3\sqrt{2}M_X^2} p u_{\nu_s}. \quad (5.8e)$$

This means that our model is only accurate for temperatures well below the mass scale of the vector boson, but in the interest of testing our approximate model, we will ignore this problem in some of our cases. Similarly to the collision rate and the potential, the QKEs describing active-sterile neutrino oscillations are given in Equation (5.6) in Section 5.2. These can be solved using LASAGNA, but the calculations turn out to be CPU heavy due to the many steps used for solving the differential equations with these parameters. As an alternative, we will attempt to develop an approximate solution based on integrated quantities.

5.8.1.1 Averaged adiabatic approximation

For the thermalisation of the sterile neutrino, we are mainly interested in the number density, but assuming that the temperature of the sterile neutrino is equal to that of the photons, we can calculate the number density directly from the energy density, which is easier to find as it appears directly in V_s from Equation (5.8e). This typically introduces an error of $\sim 20\%$ around neutrino decoupling. The energy density is given by

$$u_{\nu_s} = \int \frac{d^3p}{(2\pi)^3} p f_{\nu_s}(p) = \frac{4\pi T^4}{(2\pi)^3} \int dx x^3 \frac{P_s f_0}{2}.$$

Using Equation (5.6b), we find that

$$\dot{u}_{\nu_s} = \frac{4\pi T^4}{(2\pi)^3} \int dx x^3 \frac{-V_x P_y + \Gamma_s \left[2 \frac{f_{\text{eq},s}}{f_0} - P_s \right]}{2} f_0 = \frac{4\pi T^4}{(2\pi)^3} \int dx x^3 \frac{-V_x P_y f_0}{2}. \quad (5.27)$$

The equilibrium distribution function $f_{\text{eq},s}$ is chosen just such that it will conserve number and energy density, so the Γ_s -term integrates to zero, which leaves us with an equation depending on only V_x and P_y . The potential term can be easily calculated, while the y -component of the polarization vector can be adiabatically approximated by [60]

$$P_y = \frac{V_x D}{D^2 + V_z^2} \frac{1}{2} (P_a - P_s) \approx \frac{V_x D}{D^2 + V_z^2} \frac{1}{2} P_a \quad (5.28)$$

assuming that $P_a \gg P_s$ which must be true for most momentum modes to avoid over closing the Universe. For the adiabatic approximation to be accurate, there are a number of assumptions that must be true. The first one is that $V_x \ll \sqrt{D^2 + V_z^2}$, which is in general true away from resonance, as $\sin 2\theta \ll 1$. On resonance, $V_z \approx 0$, and the condition simplifies to $V_x \ll D$. Assuming that $n_{\nu_s} \ll 1$, we can approximate $D \approx \frac{1}{2}\Gamma_a$, and the condition is

$$T \gg \left(\frac{\delta m^2 \sin 2\theta}{x^2 C_\alpha G_F^2} \right)^{1/6} \sim 50 \text{MeV} \left(\frac{3.15}{x} \right)^{1/3} \left(\frac{\delta m^2}{\text{keV}^2} \right)^{1/6} \left(\frac{\sin^2 2\theta}{10^{-9}} \right)^{1/12}.$$

As a consequence, the approximation could break down before the interesting temperature range has been covered. The second assumption that must hold is a slow change of the effective mass basis. As the resonance is largely controlled by the sterile neutrino density, this assumption might also be stretched when the feedback of a larger sterile density moves the resonance into the bulk of the phase space distribution. Despite these severe limitations, we are now left with a simple differential equation in Equation (5.27) which can be solved by using the expressions for V_x , V_z and Equation (5.28).

This approximate method gives results very similar to the full QKE treatment at high temperatures, but there are significant deviations at lower temperatures as expected, where the difference can be several orders of magnitude.

5.8.1.2 Adiabaticity

In an attempt to estimate the adiabaticity, we follow Shi and Fuller [85], where the adiabaticity parameter is given by

$$a = V_x^2 \frac{dx_{\text{res}}}{dV_z} \left(\frac{dx_{\text{res}}}{dt} \right)^{-1}.$$

The resonant comoving momentum x_{res} is found by putting $V_z = 0$. Assuming that the ordinary matter effect is much smaller than the one from sterile neutrinos, we find

$$x_{\text{res}} = \sqrt{\frac{45\delta m^2}{7\pi^2 G_X u_{\nu_s}}} \frac{M_X}{\sqrt[4]{2} T^3}.$$

When evaluating dx_{res}/dt , we need to determine du_{ν_s}/dt . For this purpose, we assume that there is no redistribution of the momentum states and that each momentum state is filled when the resonance passes, $du_{\nu_s}/dt = f(x_{\text{res}})dx_{\text{res}}/dt$. Both the first and the second assumption might prove to be inaccurate. In this simplified setup, we can determine dx_{res}/dt and calculate the adiabaticity parameter

$$a = \frac{\delta m^2 \sin^2(2\theta) x_{\text{res}} (2u_{\nu_s} (e^{x_{\text{res}}} + 1) + x_{\text{res}})}{24HT u_{\nu_s} (e^{x_{\text{res}}} + 1) x_{\text{res}}^2}.$$

5 Sterile interactions




Figure 5.12: QKEs vs. averaged adiabatic approximation.

All results were found using the parameters $m_{\nu_s} = 1000\text{eV}$, $\sin^2 2\theta = 10^{-9}$ and $G_X = G_F$. The three line styles indicate models with $g_s = 0.0003$, $M_X = 88\text{MeV}$ (dot-dashed), $g_s = 0.00027$, $M_X = 79\text{MeV}$ (dashed), $g_s = 0.0003$, $M_X = 878\text{MeV}$ (solid). Thick color lines are solutions to the QKEs while thin black lines were found using the averaged adiabatic approximation. *Top*: Normalised energy density of the sterile neutrino. *Middle*: Position of the resonance in comoving momentum. *Bottom*: The sterile neutrino temperature relative to the photon temperature.

If the adiabaticity parameter is larger than one, the resonance is adiabatic, and a significant conversion of neutrinos can be expected. If it is smaller than one, the conversion rate ($\sim V_x$) is too low compared to the width of the resonance ($V_x dx_{\text{res}}/dV_z$) and the speed at which it moves through the momentum distribution (dx_{res}/dt).

5.8.1.3 Simulations

The general thermalisation mechanism in this model follows a certain pattern. First, non-resonant oscillations produce a small amount of sterile neutrinos at high energy as

it is seen just below $T = 200\text{MeV}$ in the upper panel of Figure 5.12. If this production is efficient enough to give $V_s = -V_0 - V_a$, the resonant production begins as it is seen in the second panel for the solid and the dashed line just above $T = 100\text{MeV}$. The resonance starts at high momentum and rapidly moves to low momentum, where it eventually stops and is typically stuck for a large temperature range. This gives rise to a quite cold population of steriles, but as the temperature of the Universe decreases, the resonance eventually starts to move towards higher momenta again. In this process, a much hotter component is added to the sterile population resulting in a temperature $T_{\nu_s} \sim (0.8 - 1.2)T_{\nu_a}$. This behaviour is especially clear for the dot-dashed line in the second and third panels. When looking at all the solutions for the parameters that we have considered, the case of $g_s = 0.0003$ (dot-dashed line) is most representative of a random choice of parameters with a non-negligible production. As the large difference between $g_s = 0.003$ and $g_s = 0.0027$ illustrates, it is not possible to find parameters which produce a partly thermalised, cold sterile neutrino. Either there is only a very limited production, or there is a large production. For certain parameters such as $g_s = 0.003$, it was possible to trap the resonance before it reached low temperatures which however resulted in a warm momentum distribution, but the only slightly lower value of $g_s = 0.0027$ produced a large number density.

We have also compared the energy density that can be calculated using Equation (5.27) to the solutions of the full QKEs in the first panel of Figure 5.12. For large temperatures we find a reasonable agreement for most parameters, but for certain parameters, where the resonance is on the threshold of getting important, the results deviate significantly as it is seen for the solid lines.

Finally, we have also calculated the adiabaticity parameter for these models and compared it to the results from the simulation. An example of this is seen in Figure 5.13. Here we can see how the resonance becomes adiabatic almost from the beginning, and stays so throughout the evolution. The reason for the reduced production at low temperatures is that the resonance has moved to very high momentum at this point (see also the second panel of Figure 5.12) where the Fermi-Dirac distribution is almost zero. While the adiabaticity calculation works well for this particular example, it turns out to be less successful in predicting the behaviour when we find $a < 1$, so although it can give some understanding of the dynamics of the resonance, it is not accurate for determining the degree of thermalisation at this point.

It is also worth noticing again that the vector boson mass is below 100MeV in this case, which means that a wrong potential was used at high temperatures. Therefore, this is not a realistic representation of any physical model, but only a tool to develop a better understanding of the resonant behaviour when self-interacting sterile neutrinos are present.

5.8.2 Improvements and a pseudoscalar mediator

To sum up the results, it is possible that a narrow range of parameters exist, where a number density that could be relevant for sterile neutrinos as dark matter is produced, but we expect it to require a high degree of fine tuning of the model parameters in case it is possible. Furthermore, our simulations and analytic analysis suggest that such a population of sterile neutrinos would be hot rather than cold if it is produced. All of this is of course only true under our assumption of $T \gg M_X$.




Figure 5.13: Resonance adiabaticity.

The adiabaticity condition as a function of temperature and energy density. The example here is for the parameters $m_{\nu_s} = 1\text{keV}$, $\sin^2 2\theta = 10^{-9}$, $G_X = G_F$, $g_s = 0.0003$, and $M_X = 88\text{MeV}$. The contours give the logarithmic value of the adiabaticity condition, while the solid line comes from the solution of the QKEs.




Figure 5.14: Potentials for sterile neutrino dark matter.

The various components of the potential as a function of temperature are indicated by the colored lines. In the region around $T = M_X$ indicated by a dotted blue line, the potential has been interpolated for visualisation. The dashed black line gives the total potential. The parameters used here are $m_s = 1\text{keV}$ and $\sin^2 2\theta = 10^{-9}$.

A major improvement of the vector boson model would be to include the correct potential at high temperatures. Before embarking on this project, it would be good to have a schematic understanding of where the various contributions to V_z dominates. However, such an insight is complicated by the fact that V_s depends on the sterile energy density which can change by many orders of magnitude. With this caveat in mind, we proceed to Figure 5.14, where it was assumed that $T_{\nu_s} = p_{\nu_s} = T$. Apart from this, a number of other simplifying assumptions were made. The sterile potential was calculated using the asymptotic expressions from [96], but around $T = M_X$, a cubic spline interpolation was used to visualise the resonance. Furthermore, the decoupling of various particles in the background has not been taken into account, as it would not alter the general patterns. At very high temperatures, the normal matter effect must eventually dominate because it scales as T^5 unless we reach the weak scale, where this scaling stops. If the sterile density is large enough, its potential will dominate V_a even for some temperatures above M_X , but the resonance is not encountered until V_s changes sign just below $T = M_X$. Finally, as the temperature drops further, we encounter the resonance that has been considered in the preceding section which arises because the vacuum contribution to V_z begins to dominate over V_s .

To get the potential as a simple function of temperature, it is necessary to assume some representative momentum, but in reality the potential also depends on momentum, and different parts of the momentum distribution could be dominated by the different contributions to V_z . Furthermore, the expressions used for V_s in Figure 5.14 are only valid if both p_{ν_s} and T_{ν_s} are respectively below or above M_X . Hence for a detailed treatment it would be necessary to find the sterile potential through numerical integration for every momentum and temperature [96].

Despite the introduction of the resonance due to the vector boson mass, it seems unlikely that this type of model would give rise to a small cold population of sterile neutrinos suitable to explain dark matter. However, the new resonance will probably have some effect. As the Universe expands and the temperature decreases, the lowest momenta will be the first to pass through the vector boson mass resonance, which could result in a quite cold population of sterile neutrinos. However, such a cold population would accelerate the transition to the low temperature regime, and the transition itself would move the resonance to higher momenta and possibly heat the sterile neutrinos. Even if the transition to the low temperature regime preserves the cold sterile neutrinos, our results from the resonance between V_s and V_0 suggest that a large population of sterile neutrinos is unavoidable as soon as the sterile neutrino population is non-negligible and the resonance is important. Although these very rough and preliminary considerations do not bode well for sterile neutrino dark matter produced by resonant oscillations due to self-interactions, only a more detailed treatment using the QKEs can give a final verdict.

With a slightly pessimistic assessment of the vector boson mediator, we will now briefly consider the pseudoscalar case. For this model, there are no complications of a boson mass as we assume it to be negligible compared to all interesting temperatures and energies. As we found in Equation (5.12), the sterile potential is proportional to T^2/p_{ν_s} . Therefore both V_s and V_0 are inversely proportional to p meaning that all momentum modes will be resonant at the same time when the equation $V_0 + V_s = 0$ is fulfilled. As a consequence, it is hard to imagine that a cold population of sterile neutrinos could be produced even if the resonance is adiabatic at the relevant temperature. Instead, it

5 Sterile interactions

might be possible to produce the dark matter at higher temperature, where $V_a > V_0$. This transition happens $\sim 100\text{MeV}$ which means that the temperature approaches the QCD transition where the background plasma changes significantly. Although this complicates matters further, the resonances will probably emerge from large momenta as the sterile density grows and end up at some rather low momenta. If this produces a cold population, and it is possible to suppress the production at the V_s, V_0 resonance, this might be a viable way of producing dark matter as a sterile neutrino.

In conclusion, the first attempts to produce a sterile neutrino population that can account for the dark matter in the Universe using the resonance induced by a self-interaction were not successful. Furthermore, the prospects of finding a working mechanism using a vector mediator are not very good, while it looks better for a pseudoscalar mediator. Despite this, it would be interesting to investigate both possibilities further as the present considerations are very qualitative and incomplete.

Part II

Chaotic neutrino oscillations

Does the flap of a butterfly's wings in Brazil set off a tornado in Texas?

— *Edward Lorenz*

In Part I we have discussed the thermalisation of sterile neutrinos through oscillations for various models. If one consider the quantum kinetic equations that were introduced in Chapter 2, they are actually linear if the contribution from the neutrino background is ignored and the simplest approximation is used for the collision term. However, for all of the cases that we have considered, we have included some term that introduces a non-linearity in the QKEs. For the full collision term in Chapter 3, the non-linearity was not important in itself, and the numerically expensive part was simply the evaluation of the collision term itself. The large lepton asymmetry in Chapter 4 and the sterile interactions in Chapter 5 were another story. Here the non-linearity gave an important physical feedback, but it was also an annoyance because it complicated the solution of the QKEs significantly. In the case of the large lepton asymmetry, the complications rose to a point where we still have not obtained a numerically stable solution.

While the non-linearity of the equations is often a impediment to solving the problem numerically, it also give rise to unique dynamics that would never be found in the solutions of linear equations. The most extreme case of this is the concept of chaos, where tiny modifications of the initial conditions can lead to widely different final solutions, and the second part of this thesis will be dedicated to that topic.

The introduction of new sterile interactions can certainly introduce non-linearities in the equations, but here we will stay with the usual weak interaction which presents a lot of unsolved problems.

The non-linearity can be introduced in different ways. In the early Universe, the non-linear term is often parametrised by the lepton asymmetry, which depends on the neutrino content of the model. Using this approach we tackle a long-standing problem of alleged chaoticity for certain mixing parameters in the presence of a small lepton asymmetry in Chapter 7. Behaviour that seem chaotic is well documented in the momentum averaged approach, but the results of solving the momentum dependent equations are more ambiguous, and our solution of the problem requires the use of advanced numerical methods.

Another realm of non-linearity is found for supernova neutrinos that experience what is called collective oscillations, and although it is hard to find a butterfly to flap its wings in a supernova, the non-linearities are surely there. Where the non-linear effect in the early Universe is independent of propagation direction, this is a pivotal concept for supernova neutrinos. A number of different instabilities are related to the various directional symmetries that can be broken by neutrino flavour oscillations. In order to widen our knowledge of these, one can attempt to study them one at a time and thereby obtain an improved understanding of their nature. One of these grossly simplified models is studied in Chapter 8, where we assess the degree of chaoticity that it exhibits and characterise the stable and unstable directions in phase space.

The equations that we will solve were already introduced in Chapter 2, but we will need some tools for analysing the chaotic behaviour. These are introduced in Chapter 6 where we discuss stability analysis. We start from simple linear stability analysis and go all the way to Lyapunov exponents and various kinds of Lyapunov vectors that can be used for the characterisation of a chaotic system.

Stability analysis

When a system is described mathematically, this is often done using differential equations or other similar constructions, which one can then attempt to solve analytically or numerically to get a detailed understanding of the system. However, while one such solution can tell everything about that particular case, it is often desirable to know if the solution is also stable towards any small perturbations that might occur. This sort of questions can be answered using stability analysis.

One of the places, where the stability is important is in the numerical solution of differential equations themselves, which we discussed briefly in Section 2.4. Here it is not the stability of the differential equation that is the concern, but rather that of the methods used to solve it — you do not want a differential equation solver to return a result that has no resemblance to the analytical solution. In this chapter, our attention is not going to be on numerical methods, but instead on the general case of a differential equation.

6.1 Linear stability analysis

The simplest and most common approach to stability analysis is to use a linear approximation of the equations of motion. For a vector of variables \mathbf{y} , one of the most general formulations of an ordinary differential equation is

$$\dot{\mathbf{y}} = \mathbf{G}(t, \mathbf{y}). \quad (6.1)$$

For such an equation, we can introduce a small perturbation, \mathbf{w} , and we can approximate the differential equation by

$$\dot{\mathbf{y}} + \dot{\mathbf{w}} = \mathbf{G}(t, \mathbf{y} + \mathbf{w}) \approx \mathbf{G}(t, \mathbf{y}) + \mathbf{w} \frac{\partial \mathbf{G}(t, \mathbf{y})}{\partial \mathbf{y}}.$$

Subtracting Equation (6.1) and defining the Jacobian or stability matrix $\mathbf{J} = \partial \mathbf{G}(t, \mathbf{y}) / \partial \mathbf{y}$, we find a linear equation for \mathbf{w}

$$\dot{\mathbf{w}} = \mathbf{J}\mathbf{w}. \quad (6.2)$$

Assuming that \mathbf{J} can be diagonalised by going to the basis $\mathbf{w}' = \mathbf{U}^\dagger \mathbf{w}$, where \mathbf{U} is a unitary matrix, a complete local solution to the equations has the form

$$\mathbf{w} = \sum_j a_j \mathbf{u}_j e^{\theta_j t}, \quad (6.3)$$

where \mathbf{u}_j is the j 'th column of \mathbf{U} , θ_j is the corresponding eigenvalue of \mathbf{J} and a_j is a constant determined from the initial condition.

6 Stability analysis

The eigenvalues are not constrained to be in the real numbers, but can also attain complex values. From the solution, we can see that a purely imaginary eigenvalue will not give rise to any significant change in the long run, but only an oscillating behaviour, while a negative real part of the eigenvalue will mean that the corresponding solution decays exponentially. Finally, the most interesting is eigenvalues with positive real parts which will give rise to exponentially diverging solutions. Since it is assumed that \mathbf{w} is a small perturbation to \mathbf{y} , this situation will be unstable, and the assumption will soon break down as \mathbf{w} grows large. Hence it is possible to predict whether a system is stable by calculating the eigenvalues of \mathbf{J} .

This method was first introduced to neutrino oscillations in dense media by Banerjee, Dighe and Raffelt [127], where they analysed a multi-angle model for supernova neutrinos. Subsequently, the method has been applied with success to many systems [128–131] with the most prominent result being the discovery of an until then unknown instability for supernova neutrinos emitted in different azimuthal angles [52].

While this approach of linear stability analysis has been very successful, it can only give stationary information on each point in the parameter and phase space as the eigenvalues contain no information about how the solution moves around in phase space. In reality, however, the dynamics of the solution can often play an important role for the actual instability of the system.

6.2 Lyapunov analysis

The goal of Lyapunov analysis is to incorporate the dynamical information of the solutions when doing a stability analysis of a given system. This does however require more complicated tools than the linear stability analysis, and the added complexity foster several different techniques that can be used.

Much of the development in Lyapunov analysis comes from meteorology where the weather systems are deeply chaotic and forecasting is one of the most important tasks. For the present description, we will however try to stay general and only dwell at the tools that we have used for neutrino oscillations in the early Universe and in supernovae.

6.2.1 Lyapunov exponents

The most important quantity in Lyapunov analysis is an analogue to the eigenvalue of the Jacobian matrix which is called the Lyapunov exponent, λ_i . Where the eigenvalue gives rise to local exponential growth with the rate θ_i as we saw in Equation (6.3), the definition of the Lyapunov exponent is such that the distance between two infinitesimally nearby orbits on average grow by $e^{\lambda_i t}$ during the time t . If we define the difference between such two solutions to be \mathbf{v}_i , we can find the Lyapunov exponent as [132–134]

$$\lambda_i = \lim_{t \rightarrow \infty} \frac{1}{t} \ln \frac{\|\mathbf{v}_i(t)\|}{\|\mathbf{v}_i(0)\|}, \quad (6.4)$$

where the limit means that we take the average over time.

As for the eigenvalues, a positive Lyapunov exponent indicates that trajectories diverges, while a negative exponent leads to convergence. In contrast to the eigenvalues, the Lyapunov exponents cannot take imaginary values by definition, and must therefore

be vanishing for situations corresponding to the oscillating solutions we found in the linear stability analysis. Here it is important to notice that there of course is no one to one correspondence between the \mathbf{J} eigenvalues and the Lyapunov exponents of a system, as the point of defining the Lyapunov exponents is to investigate how the solutions actually behave and which instabilities they pick up.

All of this is only really interesting when the dimension of the system, n , is above one, and in that case the different directions in parameter space might not diverge at the same rate. Actually, it will often be the case that some directions converge while others are marginally stable with vanishing Lyapunov exponents and others again do diverge. Hence we can define a full spectrum of Lyapunov exponents: $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_n$. The most interesting Lyapunov exponent is λ_1 as it is the largest and hence tends to dominate. If you pick a random perturbation, it will most likely contain some component that grows with the rate λ_1 , and eventually this component will be larger than any other. This is also the reason that our definition of chaos will rely on λ_1 . If the value of λ_1 is positive and the trajectory is bounded — that is, it has to return to the same part of parameter space over and over again — then we say that the system is chaotic. This said, a larger Lyapunov exponent of course leads to faster growth of perturbations than a smaller exponent, and in that sense the Lyapunov exponents quantify the degree of chaoticity.

Apart from the definition of chaos, the Lyapunov exponents have a few other interesting properties. If the same exponent appears m times in the spectrum, we say that its multiplicity is m . This means that there are also m Lyapunov vectors associated with said Lyapunov exponent, and these vectors span a m -dimensional vector space, where the m Lyapunov vectors can be picked at will as long as they are linearly independent.

For each of the Lyapunov exponents, it is possible to define several different types of Lyapunov vectors that characterise the direction in which the divergence happens. Unfortunately this choice is not unique as one could have hoped, and different desirable properties of the Lyapunov vectors turn out to be mutually exclusive.

A good visualisation of the full Lyapunov spectrum can be obtained by imagining a generalised box with the first k Lyapunov vectors as the sides. Then the generalised volume of this box will grow at the rate $\sum_{i=1}^k \lambda_i$. In this way, the Lyapunov exponents describe how the local vicinity of a point is transformed under the systems dynamics (see also [133] and [135]).

This picture also relates somewhat to properties of the system. If the system is conservative and hence also invertible, the sum of all Lyapunov exponents must be zero [132], and if the sum is negative, it means that the system is dissipative [133]. Furthermore, for Hamiltonian systems, it can be proven that the spectrum must be symmetric such that $\lambda_1 = -\lambda_n, \lambda_2 = -\lambda_{n-1}, \dots, \lambda_{n/2} = -\lambda_{n/2+1}$ [134], and that every conserved quantity in the Hamiltonian system give rise to two vanishing Lyapunov exponents [132].

In order to calculate the Lyapunov exponents, we need to follow the evolution of perturbations \mathbf{v} , and as in the case of linear stability analysis, these are determined by the equation $\dot{\mathbf{v}}(t) = \mathbf{J}(t)\mathbf{v}(t)$. Using these perturbations, it is also possible to define a propagator from t_1 to t_2

$$\mathbf{v}(t_2) = \mathbf{M}(t_1, t_2)\mathbf{v}(t_1). \quad (6.5)$$

6 Stability analysis

With this definition, the propagator must obviously also be a solution to the differential equation

$$\dot{\mathbf{M}}(t_1, t) = \mathbf{J}(t)\mathbf{M}(t_1, t). \quad (6.6)$$

On the basis of \mathbf{M} , it is possible to give a slightly different definition of the Lyapunov exponents. We construct the matrix [135, 136]

$$\mathbf{W}_+(t) = \lim_{t' \rightarrow \infty} \left(\mathbf{M}(t, t')^T \mathbf{M}(t, t') \right)^{1/(2(t'-t))}, \quad (6.7)$$

which according to Oseledets multiplicative theorem exist for almost every t using some quite weak assumptions [137]. Now we can define the Lyapunov exponents as the logarithms of the eigenvalues of \mathbf{W}_+ . The relationship to Equation (6.4) can be seen by writing $\mathbf{v}(t) = \mathbf{M}(0, t)\mathbf{v}(0)$.

Apart from an alternative definition of the Lyapunov exponents, the propagator will also form the basis for our exploration of various different Lyapunov vectors.

6.2.2 Lyapunov vectors

The simplest approach to associate a vector with every Lyapunov exponent, is to pick a set of arbitrary linearly independent perturbations and evolve them according to Equation (6.2). When taking this approach, a number of problems are encountered. First of all, the vectors will all tend to collapse along the vector associated with the largest Lyapunov exponent. In order to avoid this, it is necessary to orthogonalise the vectors whenever they begin to collapse. The second problem arises if the Lyapunov exponents are either very small (large negative values) or very large, as one might reach the smallest or highest numbers that can be represented numerically. This can be resolved by normalising the vectors and saving the information about their value before the normalisation. Finally, and more subtly than the other issues, you might not get the Lyapunov vector that you are interested in. The procedure outlined here will give what we call the backward singular vectors, while the forward singular vectors and the covariant Lyapunov vectors often have more interesting properties. Under a common name, we use Lyapunov vectors for all of these different vectors.

One approach to define a Lyapunov vector is to generalise the normal modes from the linear stability analysis which are given by \mathbf{u}_i . We will call these vectors covariant Lyapunov vectors, $\boldsymbol{\gamma}_i$, and require that the i 'th vector expands with the rate λ_i when evolved forward in time. Similarly, it must contract at the rate $-\lambda_i$ when evolving the differential equations backwards in time. These two properties mean that one gets the corresponding covariant Lyapunov vectors at different points when a covariant Lyapunov vector is evolved both forward and backward in time [136, 138], a statement which is not true for the singular vectors [135].

The disadvantage of the covariant Lyapunov vector is its inability to give certain information on the divergence or convergence of a given direction by looking at its components. This sounds a bit convoluted, so here is an example: Assume that we have a three dimensional system with the Lyapunov exponents $\lambda_1 = 1$, $\lambda_2 = -1$ and $\lambda_3 = -2$ and the associated covariant Lyapunov vectors $\boldsymbol{\gamma}_1 = \frac{1}{\sqrt{2}}(1, 1, 0)$, $\boldsymbol{\gamma}_2 = \frac{1}{\sqrt{5}}(2, 0, 1)$ and $\boldsymbol{\gamma}_3 = \frac{1}{\sqrt{2}}(1, 0, -1)$. Now, any vector in the $y = 0$ plane can be expressed as a linear combination of $\boldsymbol{\gamma}_2$ and $\boldsymbol{\gamma}_3$, which means that e.g. the vector $(1, 0, 0)$ will converge

towards the solution rather than diverge away from it. This is despite the fact that γ_1 has a non-zero x -component and $\lambda_1 > 0$. If we wish to be able to read of the stability from the Lyapunov vectors in this way, it is actually the forward singular vectors, \mathbf{f}_i , that are interesting. These are defined such that every vector that grows slower than \mathbf{f}_i is in the orthogonal complement to \mathbf{f}_i , and therefore they can be obtained from the covariant Lyapunov vectors by an orthogonalisation starting with the n 'th vector and going backwards. If we do just that for our example, we find that $\mathbf{f}_3 = \gamma_3 = \frac{1}{\sqrt{2}}(1, 0, -1)$, $\mathbf{f}_2 = \frac{1}{\sqrt{2}}(1, 0, 1)$ and $\mathbf{f}_1 = (0, 1, 0)$. From these vectors it is easy to see that the x -axis is actually a stable direction while the positive Lyapunov exponent only expands perturbations along the y -direction. In most of our cases this turns out not to be a problem, but once in Chapter 8, it is an issue and we explicitly will note it. Except for that we will pretend that the problem does not exist.

When it comes to a more stringent definition of the vectors, we need to get back to the definition of the propagator. It turns out that the forward singular vectors can be found as the eigenvectors of $\mathbf{W}_+(t)$ [139]. Similar to the definition of $\mathbf{W}_+(t)$, we can define the matrix

$$\mathbf{W}_-(t) = \lim_{t' \rightarrow -\infty} \left(\mathbf{M}(t, t')^T \mathbf{M}(t, t') \right)^{1/(2(t'-t))}, \quad (6.8)$$

where the backward singular vectors, \mathbf{b}_i , can be found as its eigenvectors, while its eigenvalues are $-\lambda_1, -\lambda_2, \dots, -\lambda_n$.

In order to find a stringent definition of the covariant Lyapunov vectors, it becomes slightly more complicated. First we need to define two sets of what is called Oseledets subspaces. The first set, O_i^+ , is given by all vectors, where Equation (6.4) gives rise to a Lyapunov exponent smaller than or equal to λ_i . This can also be expressed in terms of the forward singular vectors as

$$O_i^+ = \text{span} \{ \mathbf{f}_j(t) | j \geq i \}. \quad (6.9)$$

In a similar fashion, the other Oseledets subspace, O_i^- , consists of all vectors that give rise to a Lyapunov exponent smaller than or equal to $-\lambda_i$ when evolved backwards in time, and this can be expressed in terms of the backwards singular vectors as

$$O_i^- = \text{span} \{ \mathbf{b}_j(t) | j \leq i \}. \quad (6.10)$$

The advantage of working with these two subspaces is, that they respect the dynamics of the system. That is, any vector in $O_i^+(t)$ cannot grow faster than λ_i , which means that $\mathbf{M}(t_1, t_2)O_i^+(t_1) = O_i^+(t_2)$. Correspondingly, any vector in $O_i^-(t)$ must grow slower than $-\lambda_i$ under the inversed dynamics, so $\mathbf{M}^{-1}(t_1, t_2)O_i^-(t_2) = O_i^-(t_1)$. These are the properties that the covariant Lyapunov vectors must also fulfil, and therefore they must be in both subspaces. The exact statement is

$$\text{span} \{ \gamma_i(t) \} = O_i^+(t) \cap O_i^-(t), \quad (6.11)$$

and this relation can be used to compute the covariant Lyapunov vectors as we will see in Chapter 8, where the computational details for the Lyapunov vectors are described along with an improved algorithm for the Lyapunov exponents. As it turns out, the numerically reliable computation of covariant Lyapunov vectors requires a persistently chaotic system as we find in the toy model of Chapter 8. However, more realistic models of neutrino oscillations in very dense media often lack this property, and this will be our next topic.

6.2.3 Transient systems

For many situations, the chaotic behaviour of a system is confined to a specific part of the evolution, where the conditions are special. This is in principle always the case for neutrino oscillations in both the early Universe and in supernovae, and therefore it is very interesting for us.

The problem for such transient systems is that the limit $t \rightarrow \infty$ in Equation (6.4) does not make sense even in principle. Instead, one must use the finite-time Lyapunov exponents

$$\lambda_{i,\text{FT}} = \frac{1}{t_2 - t_1} \ln \frac{\|\mathbf{v}_i(t_2)\|}{\|\mathbf{v}_i(t_1)\|}. \quad (6.12)$$

One way to get a better sampling of the Lyapunov exponents in the system is to repeat the calculation many times with slightly different initial perturbations, but this still does obviously not give the same as a limit towards infinity would in a persistently chaotic system.

Instead, the transient nature means that perturbations only have a limited time to grow, and this can be described as lost information about the initial conditions. The amount of information in the initial conditions depend on how well they are known. More decimal points correspond to more information. This original amount of information should then be compared to the amount of information that is lost during the chaotic phase of the evolution. As the Lyapunov exponents give the divergence per time, the lost information by λ_i is given by

$$I_i = \int_{t_1}^{t_2} \lambda_i dt. \quad (6.13)$$

Combining this with the finite-time definition of λ_i can give an expression for I_i in terms of \mathbf{v}_i , but using Equation (6.12) means that we will measure the lost information in e -foldings, where it is more common to use bits [133]. Therefore we use 2 as a base for the logarithm, and we get the expression

$$I_i = \frac{1}{t_2 - t_1} \log_2 \frac{\|\mathbf{v}_i(t_2)\|}{\|\mathbf{v}_i(t_1)\|}. \quad (6.14)$$

As for the Lyapunov exponents, it might be interesting to know which directions the information loss is associated with, but here the transient nature again gives problems.

Although it is in principle possible to calculate all the various Lyapunov vectors, the results will not be very reliable, as the numerical algorithms (see the appendix in 8.2.A) needs some time to settle into the chaotic behaviour, and this time might not be available for the transient systems.

Chaotic growth of the lepton asymmetry

The intriguing connection between neutrino oscillations and a lepton asymmetry has been around for several decades as already pointed out in Chapter 4, and the story of the chaotic generation of a significant lepton asymmetry also starts two decades ago.

The initial foundation was put down by Foot, Thomson and Volkas [140] when they realised that it is possible to produce a large lepton asymmetry using active-sterile neutrino oscillations. They did, however, not realise that there was a potential for any chaotic behaviour in their findings. Nonetheless, this aspect was soon pointed out by Shi [141] who had solved the momentum averaged quantum rate equations and found that the final sign of the lepton asymmetry was very sensitive to the parameters and initial conditions. This result was later confirmed numerically [142], but at the same time, the validity of the numerics were also questioned by another group [143]. In the end, the issue was settled, concluding that the oscillations are indeed present when solving the quantum rate equations [144, 145], while when employing what has been called the static approximation, no oscillations are present [146]. This leaves the problem of which description is more accurate for describing the physics. The static approximation has the advantage that it attempts to describe the full momentum distribution using various approximations, and hence it should be preferred when its assumptions are applicable. As it turns out, this still leaves an uncovered region where the chaotic behaviour is found using the quantum rate equations. The problem using these is that they describe the full momentum distribution as one state. This is fine when all states behave similarly, but the present case depends crucially on a narrow resonance that only affects a small part of the momentum states at a time, and therefore the validity is questionable.

Nonetheless, there have continued to be an interest in the chaotic behaviour of the quantum rate equations [147, 148], an interest that is mainly motivated by the interesting phenomena that such a chaotic lepton asymmetry production would cause. The most spectacular effect is the creation of domains of varying lepton asymmetry sign throughout the Universe [149]. Apart from having different effect on the primordial nucleosynthesis in different parts of the Universe, such a division into domains would also give rise to a large production of sterile neutrinos at the domain walls, and this should in itself have observable consequences.

In an attempt to resolve the problem for once, the two groups of Di Bari and Foot [150] and Kainulainen and Sorri [46] solved the momentum dependent quantum kinetic equations numerically. Unfortunately, both calculations were haunted by very tough numerics. For Di Bari and Foot it resulted in a lack of convergence that lead them to conclude that the chaotic behaviour was plausible. On the other hand, Kainulainen and Sorri got their calculation under control, and could show something reminiscent of chaotic behaviour, but the means of controlling the calculation was an unphysical

7 Chaotic growth of the lepton asymmetry

term that could invalidate the physics and therefore also the conclusions based on the simulations.

In the paper presented in the next section, we follow the ingenious approach of Kainulainen and Sorri that allows one to resolve the narrow resonance well using a comparatively low number of momentum bins. On top of this, we employ advanced numerical solvers for stiff differential equations and combine them with a multi-threaded high-performance LU-decomposition algorithm that allows us to improve even more upon the momentum resolution in our simulations.

As we look for chaotic behaviour, we use the concept of information loss described in Section 6.2.3, and we confirm that the quantum rate equations are indeed chaotic in the sense that they disperse so much information that random thermal fluctuations in the initial lepton asymmetry become decisive for the final sign of the asymmetry. For the QKEs, we reach the opposite conclusion. Although the numerics are haunted by chaotic behaviour for the low resolution calculations, the chaotics go away when a proper resolution is used in all the relevant parts of momentum space. While some parts of momentum space need a high resolution for physical reasons (a resonance is present), other parts need it for numerical reasons only. Using this scheme, we are able to cover the parameters where Di Bari, Foot, Kainulainen and Sorri found signs of chaotic behaviour. Additionally, we can prove that parameters far away from the domain of the static approximation also lead to a non-chaotic lepton asymmetry.

Compared to the published version, the introduction has been removed, and the section on quantifying chaos has been reduced by referring to Chapter 6. Also a few typos in calculating the resonances has been corrected and the link for LASAGNA has been updated. Most of the equations for the QKEs in Section 7.1.1 are very similar to what was described in Chapter 4, but they have been included for reference as they are used when discussing the QKEs results.

7.1 ¶ Can active-sterile neutrino oscillations lead to chaotic behavior of the cosmological lepton asymmetry?

Steen Hannestad, Rasmus Sloth Hansen, Thomas Tram [1]

While the cosmic baryon asymmetry has been measured at high accuracy to be 6.1×10^{-10} , a corresponding lepton asymmetry could be as large as 10^{-2} if it hides in the neutrino sector. It has been known for some time that such an asymmetry could be generated from a small initial asymmetry given the existence of a sterile neutrino with a mass less than the mass of the active neutrino. While the magnitude of the final lepton asymmetry is deterministic, its sign has been conjectured to be chaotic in nature. This has been proven in the single momentum approximation, also known as the quantum rate equations, but has up to now not been established using the full momentum dependent quantum kinetic equations. Here we investigate this problem by solving the quantum kinetic equations for a system of 1 active and 1 sterile neutrino on an adaptive grid. We show that by increasing the resolution, oscillations in the lepton asymmetry can be eliminated so the sign of the final lepton asymmetry is in fact deterministic. This paper also serves as a launch paper for the adaptive solver LASAGNA which is available at <https://github.com/ThomasTram/LASAGNA>.

7.1.1 Equations of motion

The full system of oscillating and scattering neutrinos in the early Universe at temperatures close to neutrino decoupling is in principle described by the full N -body Hamiltonian which can be followed in time using the Liouville-Von Neumann equation for the density matrix. However, in the system of interest here, correlations introduced by collisions can safely be ignored and the system can be followed using the reduced 1-body density matrix (see [41] for an excellent treatment of this point). The system of equations for the reduced 1-particle density matrix is a generalisation of the 1-particle Boltzmann equation, known as the quantum kinetic equations.

7.1.1.1 Quantum kinetic equations

We use the density matrix formalism to describe the oscillations between active and sterile neutrinos, and we parametrise the density matrix using Bloch vector components:

$$\rho = \frac{1}{2} f_0 (P_0 + \mathbf{P} \cdot \boldsymbol{\sigma}), \quad \bar{\rho} = \frac{1}{2} f_0 (\bar{P}_0 + \bar{\mathbf{P}} \cdot \boldsymbol{\sigma}), \quad (7.1)$$

where $\boldsymbol{\sigma}$ is the Pauli matrix, and $f_0 = (1 + \exp(p/T))^{-1}$ is the Fermi-Dirac distribution function with zero chemical potential.

Since the lepton number is calculated from the asymmetry between neutrinos and anti-neutrinos, it is crucial to know this quantity with good numerical precision. We achieve this by changing to the symmetric and asymmetric variables

$$P_i^\pm = P_i \pm \bar{P}_i \quad , \quad i = 0, x, y, z. \quad (7.2)$$

7 Chaotic growth of the lepton asymmetry

Furthermore, we change to the more physical variables

$$P_a^\pm = P_0^\pm + P_z^\pm = 2\frac{\rho_{aa}^\pm}{f_0}, \quad (7.3)$$

$$P_s^\pm = P_0^\pm - P_z^\pm = 2\frac{\rho_{ss}^\pm}{f_0}. \quad (7.4)$$

With this choice of variables, the quantum kinetic equations finally have the form [46, 48]:

$$\dot{P}_a^\pm = V_x P_y^\pm + \Gamma \left[2f_{\text{eq}}^\pm / f_0 - P_a^\pm \right], \quad (7.5)$$

$$\dot{P}_s^\pm = -V_x P_y^\pm, \quad (7.6)$$

$$\dot{P}_x^\pm = -(V_0 + V_1) P_y^\pm - V_L P_y^\mp - D P_x^\pm, \quad (7.7)$$

$$\dot{P}_y^\pm = (V_0 + V_1) P_x^\pm + V_L P_x^\mp - \frac{1}{2} V_x (P_a^\pm - P_s^\pm) - D P_y^\pm. \quad (7.8)$$

Defining the co-moving momentum $x = p/T$ and the degeneracy parameter $\xi = \mu/T$, f_{eq}^\pm are given by

$$f_{\text{eq}}^\pm = \frac{1}{1 + e^{x-\xi}} \pm \frac{1}{1 + e^{x+\xi}}. \quad (7.9)$$

The potentials are given by

$$V_x = \frac{\delta m_s^2}{2xT} \sin 2\theta_s, \quad (7.10)$$

$$V_0 = -\frac{\delta m_s^2}{2xT} \cos 2\theta_s, \quad (7.11)$$

$$V_1^a = -\frac{7\pi^2}{45\sqrt{2}} \frac{G_F}{M_Z^2} x T^5 [n_{\nu_a} + n_{\bar{\nu}_a}] g_a, \quad (7.12)$$

$$V_L^a = \frac{2\sqrt{2}\zeta(3)}{\pi^2} G_F T^3 L_{(a)}, \quad (7.13)$$

where the x subscript on V denotes the x -direction in the Bloch space, not to be confused with the comoving momentum $x = p/T$. The last two potentials depend on which active neutrino flavour we are considering. For an electron neutrino $g_e = 1 + 4 \sec^2 \theta_W / (n_{\nu_e} + n_{\bar{\nu}_e})$ while for a muon or tau neutrino $g_{\mu,\tau} = 1$. All number densities n_ν are normalised to 1 in thermal equilibrium. The effective asymmetries are given as

$$L_{(e)} = \left(\frac{1}{2} + 2 \sin^2 \theta_W \right) L_e + \left(\frac{1}{2} - 2 \sin^2 \theta_W \right) L_p - \frac{1}{2} L_n + 2L_{\nu_e} + L_{\nu_\mu} + L_{\nu_\tau}, \quad (7.14)$$

$$L_{(\mu)} = L_{(e)} - L_e - L_{\nu_e} + L_{\nu_\mu}, \quad (7.15)$$

$$L_{(\tau)} = L_{(e)} - L_e - L_{\nu_e} + L_{\nu_\tau}, \quad (7.16)$$

where $L_f \equiv (n_f - n_{\bar{f}})N_f/N_\gamma$ and N_f and N_γ are the fermion and photon number densities in equilibrium respectively.

The collisional damping can be approximated as

$$D = \frac{1}{2}\Gamma, \quad (7.17)$$

7.1 \mathfrak{P} Chaotic behaviour of the cosmological lepton asymmetry

where Γ is the scattering rate of neutrinos with other neutrinos and electrons in the electron neutrino case. When evaluating the full matrix element it turns out that it can be well approximated by

$$\Gamma = C_a G_F^2 x T^5. \quad (7.18)$$

The constant C_a depends on the flavour and has the values $C_e \approx 1.27$ and $C_{\mu,\tau} \approx 0.92$.

With all the quantum kinetic equations in place, we could in principle solve the equations now. However, the re-population term of Equation (7.5) breaks lepton number conservation because of the approximate form of the scattering kernel, Equation (7.18). To ensure this conservation, we evolve an additional differential equation for L where we have put the offending term to zero in the integrand by hand:

$$\dot{L}_{(a)} = \frac{1}{8\zeta(3)} \int_0^\infty dx x^2 f_0 V_x P_y^-. \quad (7.19)$$

Finally, we will use the following explicit form for the degeneracy parameter, $\xi = \mu/T$ [48]:

$$\xi = \frac{-2\pi}{\sqrt{3}} \sinh \left(\frac{1}{3} \operatorname{arcsinh} \left[-\frac{18\sqrt{3}\zeta(3)}{\pi^3} L_{(a)} \right] \right). \quad (7.20)$$

7.1.1.2 Parametrisation of momentum space

The quantum kinetic equations, that we want to solve, have now been presented. However, as it was pointed out by Kainulainen and Sorri [46], any attempt of solving the equations using a uniform grid in momentum space is doomed, due to the very narrow resonances that are present for small mixing angles. In order to resolve these features, we will use a slightly modified version of the parametrisation they introduced.

As a first step we map the momentum grid, x , to the variable u such that the extremal point of x gets mapped to $u \approx 1/2$:

$$u(x) = K \frac{x - x_{\min}}{x + x_{\text{ext}}}, \quad K = \frac{x_{\text{ext}} + x_{\max}}{x_{\max} - x_{\min}}. \quad (7.21)$$

This improves the resolution where the distribution has most weight, but it does not help much in resolving the resonances.

The resonances can be found from the condition $V_z \equiv V_0 + V_1 + V_L = 0$. For the inverted hierarchy this results in the resonances [48]

$$x_{r_1} = \sqrt{\left| \frac{V_0}{V_1} \right|} \left(-A + \sqrt{1 + A^2} \right), \quad (7.22)$$

$$x_{r_2} = \sqrt{\left| \frac{V_0}{V_1} \right|} \left(A + \sqrt{1 + A^2} \right), \quad (7.23)$$

where $A = \frac{1}{2} V_L / \sqrt{|V_1 V_0|}$. For the normal hierarchy and $A \geq 1$, the resonances are

$$x_{r_1} = \sqrt{\left| \frac{V_0}{V_1} \right|} \left(A - \sqrt{A^2 - 1} \right), \quad (7.24)$$

$$x_{r_2} = \sqrt{\left| \frac{V_0}{V_1} \right|} \left(A + \sqrt{A^2 - 1} \right). \quad (7.25)$$

7 Chaotic growth of the lepton asymmetry

For the inverted hierarchy there is always one resonance for the neutrinos and one for the anti-neutrinos. For the normal hierarchy there are either two resonances for neutrinos or two for anti-neutrinos.

To improve the resolution at these points, we have used the same polynomials as Kainulainen and Sorri [46]. We have however used a scheme allowing for N_{res} points with improved resolution instead of only the two they used. This was done using the parametrisation

$$u(v) = \begin{cases} \alpha v + a_1 + b_1(v - v_{r_1})^3 & \text{for } v < v_{c_1}, \\ \alpha v + a_i + b_i(v - v_{r_i})^3 & \text{for } v_{c_{i-1}} < v < v_{c_i}, \quad i = 2, \dots, n-1 \\ \alpha v + a_n + b_n(v - v_{r_n})^3 & \text{for } v_{c_{n-1}} < v, \end{cases} \quad (7.26)$$

where $v_{c_i} = \frac{1}{2}(v_{r_i} + v_{r_{i+1}})$.

The differentiability of $u(v)$ requires that

$$\lim_{v \rightarrow v_{c_i}^+} \frac{du(v)}{dv} = \lim_{v \rightarrow v_{c_i}^-} \frac{du(v)}{dv}, \quad (7.27)$$

and thus $b \equiv b_i = b_j \forall i, j$. When we furthermore require continuity, and that $u(0) = 0$, and $u(1) = 1$, we get the equations

$$a_{i-1} + b(v_{c_i} - v_{r_{i-1}})^3 = a_i + b(v_{c_i} - v_{r_i})^3, \quad (7.28)$$

$$0 = a_1 - bv_{r_1}^3, \quad (7.29)$$

$$1 = a_n + b(1 - v_{r_n}^3) + \alpha. \quad (7.30)$$

This system of equations cannot be solved analytically, but using Newton's method we can find b and v_{r_i} . It is also possible to set up differential equations describing v_{r_i} , but given that Newton's method converges to machine precision in just a handful of iterations, we prefer this method due to its increased precision. Note that speed is never an issue here, since the work required in this step is completely negligible compared to the work required to calculate the remaining part of the right hand side.

With this parametrisation in place, the concentration of points close to the resonances can be increased when α is decreased. However the parametrisation introduce a complicated time dependence, and we have to modify QKE in order to take this into account [46]:

$$\left(\frac{\partial \rho(T, x(T, v))}{\partial T} \right)_v = \left(\frac{\partial \rho}{\partial T} \right)_x + \left(\frac{\partial u}{\partial T} \right)_v \left(\frac{\partial v}{\partial u} \right)_T \left(\frac{\partial \rho}{\partial v} \right)_T. \quad (7.31)$$

The first term is the ordinary QKE, while the last describes the changes induced by the parametrisation. The three factors in the last term can be found using different methods [46]. $(\partial \rho / \partial v)_T$ can be found using a simple stencil method. $(\partial v / \partial u)_T = (\partial u / \partial v)_T^{-1}$ can be found by differentiation of Equation (7.26). Finally, $(\partial u / \partial T)_v$ can also be found by differentiating Equation (7.26), but the result depends on $\partial v_{r_i} / \partial T$. To determine these, Equation (7.28) and (7.30) must be differentiated, and then the resulting system of linear equations can be solved. Note that Equation (7.29) can be used to find an expression for b and $\partial b / \partial T$, but it could also be included as another equation if $\partial b / \partial T$ was included as an unknown.




Figure 7.1: The difference between v and x . Note that x runs from only 0 to 1, while v covers x -values from 0 to 100.

7.1.1.3 Quantum rate equations

The QRE can be derived from QKE by applying two approximations. The first and most important approximation is to replace the momentum by its mean value $\langle p \rangle = 7\pi^4/(180\zeta(3))T \simeq 3.15T$, assuming a thermal distribution. This greatly reduces the number of equations in the system and gives a much simpler system to solve. Next, we neglect re-population of active neutrinos from the plasma. This is a good approximation provided that [151]

$$|\delta m^2| \sin^4(2\theta) \lesssim 10^{-9}. \quad (7.32)$$

This is the case for the parameters we are interested in. As a consequence $\dot{P}_0 = 0$, and the change from P_0 and P_z to P_a and P_s does not give any advantages. Finally, L can be calculated directly from $P_z - \bar{P}_z$, and since we are neglecting re-population, it is not necessary to evolve an independent variable to ensure lepton number conservation. In the next paragraphs we will sum up the result of the considerations above.

The density matrices of the system can be written as

$$\rho = \frac{1}{2}(1 + \mathbf{P} \cdot \boldsymbol{\sigma}), \quad \bar{\rho} = \frac{1}{2}(1 + \bar{\mathbf{P}} \cdot \boldsymbol{\sigma}). \quad (7.33)$$

Since the interesting quantity is the lepton asymmetry, we will still use the variables

7 Chaotic growth of the lepton asymmetry

$P_i^\pm = P_i \pm \bar{P}_i$ to improve numerical accuracy. The time derivatives of these are:

$$\dot{P}_x^\pm = -(V_0 + V_1)P_y^\pm - V_L P_y^\mp - D P_x^\pm, \quad (7.34)$$

$$\dot{P}_y^\pm = (V_0 + V_1)P_x^\pm + V_L P_x^\mp - V_x P_z^\pm - D P_y^\pm, \quad (7.35)$$

$$\dot{P}_z^\pm = V_x P_y^\pm. \quad (7.36)$$

The potentials V_0 , V_1 , and V_L are still given by equations (7.10)-(7.13) and the damping D by Equation (7.17), but now $x = 3.15$.

The lepton asymmetry can be calculated from the z-component of \mathbf{P} : [58]

$$L = 2 \cdot \frac{3}{8}(n_{\nu_\alpha} - n_{\bar{\nu}_\alpha}) + 2 \cdot L_{\text{initial}} = \frac{3}{8}P_z^- + 2L_{\text{initial}}. \quad (7.37)$$

While the above considerations give the correct equations, it might not be obvious why the approximations are justified. For a more rigorous derivation see e.g. Enqvist, Kainulainen, and Thomson [58] or McKellar and Thomson [45].

7.1.2 Quantifying chaos

The stochastic behaviour of the lepton asymmetry sign has been demonstrated qualitatively in QRE [142, 148] and for a few points in parameter space in QKE [46]. To quantify this we need to use some tools from the mathematical theory of chaos as it was done for QRE in [147]. These were introduced in Chapter 6, and since the Universe expands and the chaotic behaviour is confined to a limited range of temperatures, this is an example of a transient system which was discussed in Section 6.2.3.

The information present in the system depends on the thermal fluctuations in the early Universe which give a random background with $\Delta L \approx 10^{-18}$. The initial lepton asymmetry is assumed to be 10^{-10} , and the difference corresponds to about 25 bits of information [147, 150].

Apart from the information, we would also like to investigate some of the Lyapunov vectors, but the calculation of all Lyapunov vectors requires the solution of $\mathcal{O}(64v_{\text{res}}^2)$ differential equations which is unfeasible. Therefore, we limit our calculation to the largest Lyapunov exponent and the corresponding backwards singular vector.

Finally, there are some numerical issues to consider. The calculation of \mathbf{J} in every time step requires $\mathcal{O}(v_{\text{res}})$ evaluations of the QKEs, making the calculation unfeasible again. Instead, we make a compromise in accuracy by using the Jacobian matrix computed by the solver although it is not calculated as often¹. Another performance issue arises regarding the initial value of \mathbf{v}_i . The direction of \mathbf{v}_i must be chosen randomly to ensure the presence of a leading eigenvector component. However, this initial random direction gives rise to a very sensitive system and forces the solver to take very small steps. To circumvent this problem we suppress the numerical importance of the initial \mathbf{v}_i by choosing $|\mathbf{v}_i| \ll 1$. Both of these simplifications influence the detailed development of $I(t)$, but it does not change the overall patterns.

¹The criterion for recalculating the Jacobian is based on the convergence speed of the linearised system. Thus, it is reasonable to believe that the Jacobian will be recomputed if it is too different from the correct Jacobian.




Figure 7.2: The lepton asymmetry for $\delta m^2 = -10^{-2} \text{eV}^2$ and $\sin^2(2\theta) = 10^{-7}$ using quantum rate equations. Sign changes are marked with dots.

7.1.3 Results

7.1.3.1 Quantum rate equations

We have solved the quantum rate equations numerically for the muon neutrino case, and our results are similar to what others have found [142, 147, 148]. We have furthermore followed the evolution of the Lyapunov vector and calculated the total information loss. An example that demonstrates the type of oscillations we see is found in Figure 7.2. Since we only follow one momentum state, we can predict the expected position of the oscillations by solving $V_0 + V_1 + V_L = 0$, and we find that this prediction holds. In Figure 7.2 we see a significant number of oscillations, but if $\sin^2(2\theta)$ is decreased, the number of oscillations decreases as well while an increased value of $\sin^2(2\theta)$ results in more oscillations. The number of oscillations and the resulting loss of information also depends on δm^2 , but in a less straight forward way.

The amount of information that is lost for a given set of parameters is seen in Figure 7.3. On top of this, the final sign of the lepton asymmetry is shown as a bright or dark shading, and the agreement between the two patterns is quite good. The area where the lepton asymmetry sign seems to be very sensitive to changes in δm^2 and $\sin^2(2\theta)$ are the same areas where the information loss exceeds the limit for stochasticity, $I = 25$. The slightly high values of information loss that show up in the lower left corner is likely due to a loss of information that is not related to the lepton asymmetry since the associated Lyapunov vector, which indicates the variables that are responsible for the information loss, is mainly pointing in the P_x and P_y direction.

The levels of information loss in Figure 7.3 show the same tendency as the Lyapunov




Figure 7.3: The Information is shown in colour. A bright shading of the colour indicates a positive final lepton asymmetry and a dark shading of the colour indicates a negative final lepton asymmetry. All $I > 25$ have the same colour as this is the limit of stochasticity.

analysis that was done previously [147]. The numbers are not exactly the same, but the exact results depend on the numerical scheme that is used as well as the initial temperature and other parameter choices. This means that the loss of information cannot be used as an exact gauge to distinguish stochastic and non-stochastic areas, but it can help to deepen our understanding of the quantum rate equations.

7.1.3.2 Quantum kinetic equations

The quantum kinetic equations have been solved numerically using a numerical differentiation formulae of order 1-5 (`ndf15`) devised by Shampine and Reichelt [62], and again we have focused on the muon neutrino case. Since the system is expected to be stiff, we have reduced the maximal order from 5 to 2, and thereby the solver becomes L-stable. The `ndf15` method is implicit and uses a modified Newton's method to solve the implicit equation. This imply the solution of matrix equations, and due to the need of many momentum bins, they take up most of the computation time. To improve the performance and take advantage of multicore CPU's we used the `SuperLU_MT` package [152, 153] to solve the matrix equations. `SuperLU_MT` takes advantage of a very clever work distribution mechanism, and it depends on efficient BLAS-implementations for speed. It is possible to use a multithreaded BLAS implementation in conjunction with `SuperLU`, but we found it most efficient to use a sequential BLAS and let `SuperLU_MT` do all the threading. We tested three different BLAS-libraries and found Intel's MKL BLAS to be the fastest for our purposes.

7.1 \mathfrak{B} Chaotic behaviour of the cosmological lepton asymmetry




Figure 7.4: The lepton asymmetry for $\delta m^2 = -10^{-2} \text{ eV}^2$ and $\sin^2(2\theta) = 10^{-7}$. Stars mark sign changes. Note that $v_{\text{res}} = 1600, 2400,$ and 3200 are indistinguishable. The parameters used here are also used in Figure 7.2.




Figure 7.5: The lepton asymmetry for $\delta m^2 = -10^{-2} \text{ eV}^2$ and $\sin^2(2\theta) = 10^{-6}$. Stars mark sign changes.

7 Chaotic growth of the lepton asymmetry




Figure 7.6: The lepton asymmetry for $\delta m^2 = -1\text{eV}^2$ and $\sin^2(2\theta) = 10^{-6}$. Stars mark sign changes.

Our main results are shown in Figure 7.4-7.6. All three cases show that the oscillations present for a small number of momentum bins disappear when a sufficiently large number of bins are used. The first figure also confirms that the solution has converged as the three solutions with most bins are indistinguishable. These results are also unchanged when we modify the parametrisation parameters within reasonable limits, and the patterns are similar for all the points in parameter space that we have examined. This is contrary to the findings using quantum rate equations(QRE) [142, 147, 148] and to the results of two earlier studies using the full quantum kinetic equations(QKE) [46, 146]. The chosen oscillation parameters in Figure 7.4-7.6 are all in the region where QRE show many oscillations, and the parameters of Figure 7.5 are in a part of the parameter space where the final sign of the lepton asymmetry appears to be stochastic [142, 148].

The lack of oscillations can be understood from the quantum kinetic equations. Consider the equations for P_y^+ , P_x^+ and P_y^- .

$$\dot{P}_y^+ = (V_0 + V_1)P_x^+ + V_L P_x^- - \frac{1}{2}V_x(P_a^+ - P_s^+) - DP_y^+ \quad (7.38)$$

$$\dot{P}_x^+ = -(V_0 + V_1)P_y^+ - V_L P_y^- - DP_x^+ \quad (7.39)$$

$$\dot{P}_y^- = (V_0 + V_1)P_x^- + V_L P_x^+ - \frac{1}{2}V_x(P_a^- - P_s^-) - DP_y^- \quad (7.40)$$

Let us take a point in momentum space above the resonant value as it is done in Figure 7.7. Here the time derivatives are very close to zero before the resonance has passed. As the temperature decreases, the resonance approaches the point, and $(V_0 + V_1)P_x^+$ becomes less dominant in the equation for \dot{P}_y^+ . Since $-1/2V_x P_a^+$ is unchanged, this leads to




Figure 7.7: P_y^+ , P_x^+ , and P_y^- as a function of T for $x = 0.3$. $\delta m^2 = -10^{-2} \text{eV}^2$, and $\sin^2(2\theta) = 10^{-7}$ as in Figure 7.2 and Figure 7.4.

a rise in P_y^+ . The growing value of P_y^+ effects \dot{P}_x^+ , and P_x^+ becomes positive as well. Finally, this affects \dot{P}_y^- which becomes positive due to the term $V_L P_x^+$. When this initial mechanism has excited the different parts of the density matrix, the complexity of the full equations dictates the detailed evolution, but eventually the damping terms will dominate, and P_y^- becomes zero.

To sum up the mechanism: The amplitude of P_y^- depends directly on the value of L , and as L decreases the amplitude, P_y^- does the same. Since dL/dt is the integral of P_y^- , this mechanism means that as L approaches zero so does dL/dt , and there will never be a sign change in L .

Despite the feedback mechanism just described, we see oscillations when the number of momentum bins is too small. This happens because the mechanism is a gross simplification, and different effects that are not considered can easily make the approximations break down. For the case in Figure 7.4 we can try to understand why the solutions show oscillation for 800 momentum bins but no oscillations for 1600 and 3200 bins.

In Figure 7.8 the evolution of dL/dt is shown as a function of temperature. As it could be expected, there are oscillations for 800 bins. What might be more surprising is that the solutions with 1600 and 3200 bins show oscillations as well, however on a much smaller scale. Since dL/dt is the integral of P_y^- , we consider this variable in Figure 7.9. At a temperature of 20 MeV there is no difference between the three solutions. In the first panel some small oscillations start to emerge, and they become dominant for the solutions of the 800 and 1600 bin cases at a temperature of 12.5 MeV. The key difference between 800 and 1600 bins is that the oscillations are symmetric around the correct solution in the 1600 bin case and asymmetric in the 800 bin case. This explains the difference in Figure 7.8. Notice also that the scale of the y-axis decreases dramatically as the temperature evolves. This is to be expected due to the feedback mechanism

7 Chaotic growth of the lepton asymmetry




Figure 7.8: The oscillations in L are also present in dL/dt as it could be expected. The parameters are $\delta m^2 = -10^{-2}\text{eV}^2$ and $\sin^2(2\theta) = 10^{-7}$ as in Figure 7.4, and the system has been solved for the same values of T . Stars mark sign changes.




Figure 7.9: $\delta m^2 = -10^{-2}\text{eV}^2$ and $\sin^2(2\theta) = 10^{-7}$ as in Figure 7.4 and 7.8. P_y^- is dominated by prominent oscillations as the temperature falls, and this leads to the oscillations seen in Figure 7.8.




Figure 7.10: Compare this Figure to Figure 7.9. The oscillations in P_y^- come from the V_x -term in Equation (7.40). The oscillations in this term are present even for high temperatures, but they do not dominate until the remaining terms shrink to $\mathcal{O}(10^{-3})$.

described earlier.

When we want to pinpoint the source of these rapid oscillation further, we need to look at dP_y^-/dT . This is a quite complicated quantity where the different terms cancel each other to a very high degree, and the oscillations could in principle originate in a complicated interplay between the different terms. Fortunately, the $V_x P_s^-$ -term seems to contribute a lot more to the oscillations than the other terms as it is seen in Figure 7.9. This means that the oscillations come from P_s^- , and as Figure 7.11 shows the oscillations grow gradually from some ill resolved features. P_s^- has the advantage that its derivative is very simple. It only depends on P_y^- which has no oscillations at that high temperatures, and thus the only possibility is that the oscillations come from the term added to account for the parametrisation. This is also to be expected since an insufficient resolution in momentum space results in an inaccurate estimate of $\partial P_s^-/\partial v$.

The most straightforward way to eliminate the error in $\partial P_s^-/\partial v$ is to increase the number of momentum bins in the calculation, but another option is to change the parametrisation. When increasing the α parameter, less points are placed close to the resonance and more away from it. This is shown in Figure 7.12 where the convergence is achieved by increasing α rather than increasing v_{res} . The drawback of this approach is that the resonance becomes ill resolved at some point, and this does also result in oscillations. A hint of this effect can be seen in Figure 7.12 for high temperatures, but it is too small to influence the final result. A third method to improve the sampling of the negative feature in Figure 7.11 is to change the parametrisation. This can easily be done by including a third v_{r_i} in Equation (7.26). Since the feature starts at the original position of the resonance, and $P_s^- = 0$ when the resonance has moved away, it




Figure 7.11: $\delta m^2 = -10^{-2}\text{eV}^2$ and $\sin^2(2\theta) = 10^{-7}$ as in Figure 7.4, 7.8, 7.9 and 7.10. There are no oscillations at $T = 25\text{MeV}$, but they emerge as the negative feature is moving away from the resonance as the first panel shows. These oscillations are purely numerical in origin, and they arise because the region below the resonance is undersampled. In the second panel the oscillations extend across the resonance and this gives rise to the oscillations seen in Figure 7.10.

remains at this position at least until the lepton asymmetry starts to grow, and this is the period we are interested in. With this modification to the code we can show the convergence to no oscillations even for $\delta m^2 = -1\text{eV}^2$ and $\sin^2(2\theta) = 10^{-5}$, which is in the middle of the oscillating region found by Di Bari and Foot [146]. In general, many different effects may result in oscillations, and we shall not try to describe all the cases here. However, all the oscillations we have seen disappear when the number of momentum bins is high enough and all features are well resolved.

We have tried to include a smoothing term for the sterile neutrinos to confirm the oscillations seen by Kainulainen and Sorri [46]. In order to conserve the number density we added

$$R_{\nu_s}^{\pm} = r_s \Gamma (n_{\nu_s} f_{\text{eq}}(p, \mu) \pm n_{\bar{\nu}_s} f_{\text{eq}}(p, -\mu) - \rho_{ss}^{\pm}) \quad (7.41)$$

to \dot{P}_s^{\pm} . This did give rise to some oscillations with amplitudes a lot smaller than 10^{-10} , but the solution converged to our original result when r_s was gradually reduced. It is worth noting that the computation time was reduced with a factor of ~ 15 when the distributions of the sterile neutrinos were smoothed out. Unfortunately the amount of smoothing that is necessary to remove the numerical noise in P_s^- is so large that it significantly alters the behaviour of L , even when the added term conserves number density. The reason for the high sensitivity of L is the large cancellations that occur in several different terms, one of them being for dL/dt which is the integral of an oscillating function as shown in Figure 7.9.




Figure 7.12: $\delta m^2 = -10^{-2}\text{eV}^2$ and $\sin^2(2\theta) = 10^{-7}$ as in Figure 7.4 and Figures 7.8-7.11. Stars mark sign changes. The oscillations disappear as alpha rises and the negative feature from Figure 7.11 becomes better sampled. Note however that some other oscillations start to emerge for high temperatures, as alpha rises and the sampling of the resonance becomes worse.

7.1.3.3 Lyapunov analysis of quantum kinetic equations

In the previous section we have shown that the oscillating behaviour found in QRE disappear in QKE when an adequate number of momentum bins is used. However, it could still be interesting to see how this disappearance manifests itself in the Lyapunov analysis of the system.

When calculating the information loss, the initial direction of the Lyapunov vector is chosen at random. As the computation is started, this initial direction changes, and the length of it grows rapidly leading to an information loss of approximately 20 bits. Since this information loss can not be related to a physical phenomenon, but has a purely numerical origin, we reset the scale of information loss accordingly. The parameters we have chosen to investigate are once again those of Figure 7.4. Since the Lyapunov analysis introduces n additional equations, we have chosen to vary α instead of the number of bins. The results are shown in Figure 7.13.

The oscillating case is seen to the left, and the information loss is substantial here as could be expected. When $\alpha = 0.1$, the information loss is associated with a Lyapunov vector pointing in the direction of P_x^\pm and P_y^\pm . This loss of information is obviously not related to an unpredictable sign of the lepton number, but it is rather related to the decoherence as the temperature decreases. It is interesting to note that the calculation of Lyapunov exponents affects the convergence of the solution. When the Lyapunov exponents were not calculated a $\alpha = 0.5$ was needed, but when the Lyapunov

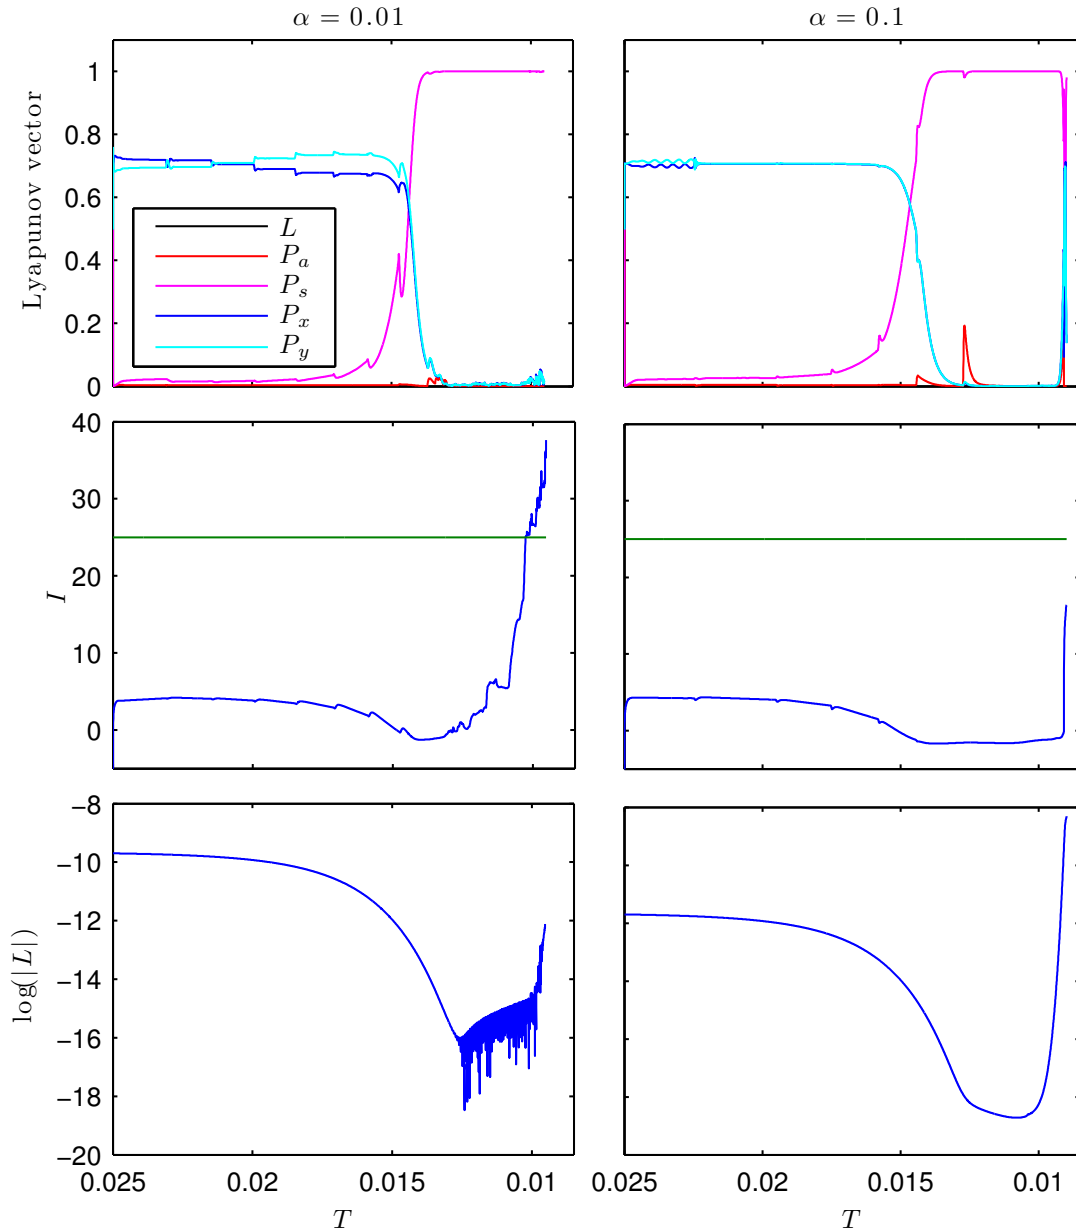


Figure 7.13: $\delta m^2 = -10^{-2}\text{eV}^2$ and $\sin^2(2\theta) = 10^{-7}$ as in Figure 7.4 and 7.8-7.11. The Lyapunov analysis has been performed for $\alpha = 0.01$ and $\alpha = 0.1$. The first plot for each α shows the relative components of the Lyapunov vector, the second plot shows the amount of lost information, and the last plot shows the logarithm lepton asymmetry.

exponents were included, it only took a value of 0.1 to converge to the non-oscillating solution. This is probably due to some slight changes in numerical errors while solving the equations, and it further illustrates that the oscillations are non-physical.

7.1.4 Conclusions

We have studied the evolution of the lepton asymmetry in models with active sterile neutrino oscillations. First, we confirmed that using the quantum rate equations and assuming the hierarchy to be inverted there are certain regions of the mixing parameter space where the lepton asymmetry can undergo sign changes and where the sign of the final asymmetry is chaotic in nature.

Once the full momentum dependent quantum kinetic equations are used these oscillations vanish as soon as the resolution is sufficiently high and the system no longer exhibits any signs of chaotic behaviour. We quantified the behaviour of the system using a Lyapunov analysis and found that there is indeed a very significant amount of information loss in the regions where L becomes oscillatory, to the point where the final sign does become chaotic in nature. Using this analysis it is also clear that this information loss is not a physical phenomenon, but rather an artifact associated with lack of numerical resolution in momentum space. This settles a long standing open question in early Universe neutrino phenomenology.

A final, interesting point is that the Lyapunov analysis can also be used to study kinematical decoherence of the system at temperatures below resonance. At low temperatures the most significant information loss is no longer associated with L , but rather with P_x, P_y , i.e. in the transverse direction associated with normal vacuum oscillations.

Acknowledgements

We would like to thank Kimmo Kainulainen, Georg Raffelt, and Irene Tamborra for helpful discussions and comments on the manuscript.

7.1.A LASAGNA

Code available at <https://github.com/ThomasTram/LASAGNA>. This appendix gives a short overview of the numerical code that we have developed.

7.1.A.1 Overview

LASAGNA² is a code written in C for solving stiff systems of ordinary differential equations (ODEs). We developed the code for solving the Quantum Kinetic Equations for active-sterile oscillations in the early Universe, but we stress that due to its modular structure makes, it is easy to implement any other set of differential equations. Input parameters are read from a text file³, and output is handled by a user specified output function. This could be a simple function that just writes to the screen, but **LASAGNA** also contains custom routines for creating, displaying and modifying its own binary file format. The

²Possibly an acronym for Lepton Asymmetric Sterile-Active momentum-Grid Neutrino Analyser.

³Many thanks to Julien Lesgourgues for letting us use the parser he wrote for his Boltzmann-code **CLASS**.

LASAGNA binary files are compatible with MATLAB for fast and easy visualisation of output data.

The user can choose from 3 different ODE-solvers, `ndf15`, `RADAU5` and an implementation of the embedded Runge-Kutta formula due to Dormand and Prince of order 4 and 5. All three solvers are plug-compatible, so it is easy to try different solvers and also possible to implement another solver. `ndf15` and `RADAU5` are both implicit solvers, and they both use a modified Newton's method for solving the (possibly) non-linear system of equations. The implicit methods also require access to a linear algebra solver. As the most optimal linear algebra solver for a given problem depends on the size of the system and the sparsity of the Jacobian, we have implemented one dense and two sparse solvers. Like the ODE-solvers, all the linear algebra solvers are plug-compatible.

7.1.A.2 ODE-solvers

We have included 3 ODE-solvers in LASAGNA, `ndf15`, `RADAU5` and an explicit Runge-Kutta solver which is not optimised and is primarily included for reference. The first two are both suitable for stiff systems. `ndf15` is a variable order, adaptive step size linear multistep solver, based on the Numerical Differentiation Formulae of Shampine [62] which are of order 1-5. Since the formulae are only L-stable for order 1 and 2, we often found it necessary to reduce the maximal order to 2⁴.

`RADAU5` employs an implicit Runge-Kutta method of order 5 which is L-stable. It is a 3-stage method, so the resulting system of algebraic equations are naively $3N \times 3N$. However, by doing clever transformations on the Runge-Kutta matrix⁵, Hairer and Wanner showed [154] that the equations separate to a $N \times N$ system and a $2N \times 2N$ system. Moreover, the $2N \times 2N$ system are in fact equivalent to a $N \times N$ system of complex numbers. The amount of linear algebra work in each step of `RADAU5` is then just 5 times the work in each step of `ndf15`. Which one is better will depend on size, stiffness and tolerance parameters.

7.1.A.3 Linear algebra solvers

A good linear algebra solver is very important for implicit solvers, but the best solver is problem dependent as we have shown in Table 7.1. It is important to have an interface which is general enough to accommodate a range of solvers, but without sacrificing performance. A wrapper to a linear algebra solver in LASAGNA consists of the following 4 subroutines which are passed to the evolvers:

```
int linalg_initialise(MultiMatrix *A,
                    void **linalg_workspace,
                    ErrorMessage error_message);
int linalg_finalise(void *linalg_workspace,
                  ErrorMessage error_message);
int linalg_factorise(MultiMatrix *A,
                   EvolverOptions *options,
                   void *linalg_workspace,
```

⁴This is done by setting `int maxk=2` in `evolver_ndf15.c`

⁵One finds the transformation matrices that brings the Runge-Kutta matrix to its real Jordan form.

7.1 \mathfrak{B} Chaotic behaviour of the cosmological lepton asymmetry

	dense_NR	sparse	SuperLU
Type of method	Dense	Sparse	Sparse
Range of sizes	1 – 100	$\sim \mathcal{O}(10^1 - 10^2)$	$\sim \mathcal{O}(10^2 - 10^5)$
Threaded	No	No	Yes (good scaling up to 16 cores)

Table 7.1: The implicit ODE-solvers can utilise different linear algebra solvers through a common interface.

```

                                ErrorMessage error_message);
int linalg_solve(MultiMatrix *B,
                MultiMatrix *X,
                void *linalg_workspace,
                ErrorMessage error_message);

```

The `MultiMatrix` structure is used to represent any matrix used by the evolvers, and it is defined in `multimatrix.h`. One notable exception from table 7.1 is LAPACK for large dense systems, which could easily be implemented using the above scheme.



7.2 Perspectives on the chaotic lepton asymmetry

Our results in the previous section demonstrates that the idea of chaotic generation of large lepton asymmetries using active-sterile neutrino oscillations is dead for the parameter space where it was claimed. This also means that the idea of generating domains of positive and negative lepton numbers cannot be fulfilled using this mechanism. However, other mechanisms that can give rise to such domains might still work [155, 156], although this particular model using isocurvature perturbations might be severely constrained by the recent results from Planck [21].

Apart from the direct implications of our results, the paper also demonstrates an interesting use of Lyapunov analysis for neutrino oscillations. While the chaotic behaviour for neutrino oscillation in the early Universe has now been proven to be absent, it might well be that e.g. supernovae provide sites where the nonlinear feedback of neutrinos on their own oscillations give rise to chaotic behaviour.

Currently, the understanding of collective neutrino oscillations in supernovae is rapidly evolving, and it has been realised that there are several crucial complications that have previously been neglected. Some of the most important are: The neutrino were assumed to evolve identically in different angular directions [52] (they evolve in different ways); neutrinos scattered backwards in the outer part of the supernova were assumed to be subdominant [157] (they often contribute as much as the other neutrinos); the initial approximate homogeneous emission of neutrinos were assumed to be stable [131] (they are unstable). All of these issues present a huge challenge for simulating collective neutrino oscillations in supernovae, and it seems unfeasible to perform a calculation where all of the problems are taken into account at the same time. However, some of the simulations that have been presented for toy-models describing one of the problems at a time, suggest that chaotic behaviour could play a part [158, 159]. If this is the case, it could well be that the information loss approach will be relevant at some point as a tool for characterising which models can be trusted in detail and which are merely a result of the detailed choice of initial conditions.

In the next chapter we will consider a very simple example of such a toy-model and explore the use of Lyapunov exponents and covariant Lyapunov vectors in a study of the model.

Chaotic supernova neutrinos

A core-collapse supernova ends the life of a star more massive than $\sim 8M_{\odot}$. At the end of its life cycle, it collapses until the core reaches nuclear density. At this point most of the in-falling material bounces on the newly formed protoneutron star, and a shock front travels out through the plasma. At the same time, the protoneutron star starts emitting vast amounts of neutrinos first from the conversion of protons and electrons to neutrons, and later from cooling. The luminosity is estimated to be about 10^{53} erg/s, which corresponds to more than 99% of the total energy emitted by the star. This huge fraction of energy means that the neutrinos are very important for the dynamical evolution of the explosion, and that neutrino flavour oscillations are important to understand not only for the interpretation of a supernova neutrino signal, but potentially also for the explosion itself.

The neutrinos travel through the dense envelope of the protoneutron star, and this environment gives rise to several different effects that enhance and inhibit neutrino oscillations. Since the star consists of matter, there is a non-negligible matter potential from electrons. Furthermore, there are so large amounts of neutrinos that neutrino-neutrino self-interactions are important. Finally, the geometry is not at all homogeneous and isotropic, but generically three dimensional. Throughout the past years, the hydrodynamic simulations of supernova explosions have evolved from spherically symmetric one-dimensional models to advanced three-dimensional simulations where convective overturn and other genuinely multidimensional effects are important. For the neutrino oscillations in comparison, we have only just started to break the spherical symmetry, and we still lack a good understanding of some of the caveats that might be hidden in a multidimensional treatment. One of the most important aspects of this is the existence of instabilities that can lead to significant flavour conversion and also pose a problem for the numerical stability of simulations.

8.1 Instabilities in collective supernova neutrino oscillations

Due to the huge fraction of energy emitted as neutrinos from a supernova, they have been one of the favourite environments for neutrino oscillations. In the outer part of the supernova envelope, the MSW-resonance is crossed, but at shorter distances from the centre, the neutrinos themselves constitute an important part of the background. On top of this, the anisotropic environment means that the factor $(1 - v_{\mathbf{p}} \cdot v_{\mathbf{q}})$ must be included for the background potential, which means neutrinos from different directions couple at different strengths.

Besides the anisotropy term, there are also a few other differences compared to the treatment of neutrino oscillations in the early Universe that we have used so far. As

already mentioned in Section 2.2, we define the polarisation vectors for the antineutrinos in terms of the density matrix $\bar{\rho}$ instead of $\bar{\rho}$. Because $\bar{\rho} = \bar{\rho}^*$, our new definition of the antineutrino polarisation vector has the opposite y -component compared to the previous definition. This does not change much, but since there is a cross product involved in the QKEs, we will see a change of sign for the new antineutrino polarisation vector. Furthermore, we now consider active-active neutrino oscillations, and it is necessary to account for the effect of oscillated neutrino in the background. Starting from the asymmetric term in Equation (2.17), this gives rise to a potential of the form

$$\mathbf{V}_{\text{int}, \nu}(\mathbf{p}) \propto \int d\mathbf{q} \left(\mathbf{P}(\mathbf{q}) - \bar{\mathbf{P}}(\mathbf{q}) \right) (1 - v_{\mathbf{q}} \cdot v_{\mathbf{p}}), \quad (8.1)$$

for both the neutrinos and the antineutrinos.

The important contribution from background neutrinos lead to non-linear equations of motion, and the results are some peculiar oscillation phenomena that are commonly described as collective neutrino oscillations. Two of the most interesting patterns that were found [49, 50] are the synchronised oscillations, which give rise to collective small amplitude oscillations of all momentum modes, and bipolar oscillations where a full conversion from one flavour to another can happen even for very small mixing angles, somewhat similar to a spherical pendulum starting in an upwards position.

This simple picture becomes more complicated once more details of the supernova density and geometry are included, and the first attempt at a reasonable numerical description used the bulb model [47], which assumes that the neutrino emission is spherically symmetric, but does allow for a dependence on the neutrino emission angle relative to the proto-neutron star surface.

This approach has since been expanded using three-neutrino mixing, more realistic energy and density profiles as well as angle distributions, and a instability related to the angle relative to the surface was identified. However, at the same time it was realised that even an approximately symmetric neutrino gas might be unstable towards small perturbations in the angular distribution. The issue was raised for the anisotropic environment of a supernova by Raffelt, Sarikas and Seixas [52] by doing a stability analysis of a distribution that depends on the azimuthal angle as well as the elevational. The result has been confirmed by several numerical studies [159–161], and as a result, one needs to be very careful when using the bulb model, as it might not capture all the relevant physics, although it is also possible that the instability is suppressed in some cases [128].

As an orthogonal approach to the more realistic simulations, it was also realised that the instability can be found in very simple systems only consisting of two neutrino beams [51]. The idea is that the simplest azimuthal distribution you can think of occurs in one dimension, where neutrinos going in the two directions are allowed to have different flavour compositions. It turns out that if this is not permitted, it is only the inverted neutrino mass hierarchy that allows for flavour conversion (the bipolar oscillations mentioned above). This is in contrast to the case where the evolution is independent, and conversion is observed for both hierarchies. On top of the possibility of bipolar oscillations in the normal mass hierarchy, it was also found that the flavour evolution seems to enter a chaotic phase which we will now analyse further.

The goal of the paper presented in the next section is to analyse the stability of the two beam model and investigate if it shows signs of chaos in the stringent sense that we

8.1 Instabilities in collective supernova neutrino oscillations

defined in Section 6.2. We find that the behaviour is indeed chaotic by calculating the largest Lyapunov exponent, but we also extend the calculation to the other Lyapunov exponents and the covariant Lyapunov vectors. This shows that the normal mass hierarchy in general have larger Lyapunov exponents than the inverted hierarchy, which means that it is more unstable towards perturbations in this model. From the Lyapunov vectors, we get a detailed characteristic of the unstable directions, and hence knowledge of the type of perturbations that will grow.

Compared to the published version of the paper, the introduction has been omitted, and the sections on Lyapunov analysis have been removed as they were already covered in Chapter 6. The exception is the description of the numerical algorithms used, where the details are quite specific to the model considered here, although similar approaches can be used on other system.

8.2 \wp Chaotic flavor evolution in an interacting neutrino gas

Rasmus Sloth Hansen, Steen Hannestad [2]

Neutrino-neutrino refraction can lead to non-periodic flavor oscillations in dense neutrino gases, and it has been hypothesized that some solutions are chaotic in nature. This is of particular interest in the case of neutrino emission from core-collapse supernovae where the measurement of the spectral shape for different flavors can provide crucial information about both neutrino physics and the physical conditions close to the proto-neutron star. Whether a system is chaotic or not can be assessed by the Lyapunov exponents which quantify the rate of divergence of nearby trajectories in the system. We have done a numerical case study for a simple toy model of two neutrino flavors with two momentum states traveling against each other which is known to exhibit flavor transition instabilities. We find the leading Lyapunov exponent to be positive in all cases, confirming the chaoticity of the system for both the normal and the inverted neutrino mass hierarchy. However, more Lyapunov exponents were approximately zero in the inverted hierarchy compared to the normal which has implications for the stability of the system. To investigate this, we have calculated a generalized set of normal modes, the so-called covariant Lyapunov vectors. The covariant Lyapunov vectors associated with vanishing Lyapunov exponents showed the existence of marginally stable directions in phase space for some cases. While our analysis was done for a toy model example, it should work equally well for more realistic cases of neutrinos streaming from a proto-neutron star and provide valuable insight into the nature of the flavor instability. We finally stress that our approach captures many more properties of the physical system than the linear stability analyses which have previously been performed.

8.2.1 The two beam model

Our model contains only two momentum states, but before we specialize to that case, we will consider the more general case of N momentum modes.

8.2.2 N momentum modes

In the general case we consider a neutrino gas of oscillating ν_e and ν_x ($x = \mu$ or τ) consisting of N momentum modes. Using the polarization vector parameterization of the density matrices, the oscillation equations without a matter background can be written as [51]

$$\begin{aligned}\dot{\mathbf{P}}_i &= \left(\omega_i \mathbf{B} + \frac{\mu}{2} \sum_{j=1}^N (\mathbf{P}_j - \bar{\mathbf{P}}_j) (1 - \vec{v}_j \cdot \vec{v}_i) \right) \times \mathbf{P}_i, \\ \dot{\bar{\mathbf{P}}}_i &= \left(-\omega_i \mathbf{B} + \frac{\mu}{2} \sum_{j=1}^N (\mathbf{P}_j - \bar{\mathbf{P}}_j) (1 - \vec{v}_j \cdot \vec{v}_i) \right) \times \bar{\mathbf{P}}_i,\end{aligned}\tag{8.2}$$

where $\bar{\mathbf{P}}_i$ refers to antineutrinos, $\mu \sim 2\sqrt{2}G_F n_\nu$ ¹, $\omega_i = \delta m^2/2E_i$, \mathbf{B} is the mass unit vector in flavor space, and $\vec{v}_i = \vec{p}_i/E_i$ is the direction of the momentum. We use arrows

¹Note that we define μ slightly different than in [51] to absorb a factor of 2 in Equation (8.6).

to indicate vectors in real space while bold faces refer to vectors in polarization space.

There are two obvious choices for the coordinate system in polarization space. The first takes the z -direction to coincide with the pure electron neutrino state, and the second lets \mathbf{B} determine the orientation of the z -axis and exploits the symmetries of the equations. Since we do not aim to calculate any oscillation probabilities, we will adopt the latter convention and set $\mathbf{B} = (0, 0, -1)$. This choice for \mathbf{B} ensures that $\omega > 0$ corresponds to the normal hierarchy while $\omega < 0$ corresponds to the inverted hierarchy.

Since the two-flavor oscillation is a two level system, it has many similarities to spins, and this is highlighted by the formulation in terms of polarization vectors. In the isospin convention, the isospin vectors can be identified with angular momenta, but this means that neutrinos and antineutrinos with similar flavor content will be associated with isospin vectors pointing in opposite directions. To avoid this, we choose the opposite sign for $\bar{\mathbf{P}}_i$, and therefore \mathbf{P}_i and $-\bar{\mathbf{P}}_i$ correspond to the angular momenta.

As we assume no dissipation in our equations of motion, the system is Hamiltonian, and it turns out to be enlightening to consider a classical Hamiltonian formulation of the equations. For this kind of motion confined to a set of spheres (the lengths of \mathbf{P}_i and $\bar{\mathbf{P}}_i$ are constant), the canonical coordinates and momenta are φ_i and $P_{iz} = P_i \cos \theta_i$ for the neutrinos and $\bar{\varphi}_i$ and $-\bar{P}_{iz} = \bar{P}_i \cos \bar{\theta}_i$ for the antineutrinos. From these variables we can define the polarization vectors as

$$\begin{aligned}\mathbf{P}_i &= P_i(\cos \varphi_i \sin \theta_i, \sin \varphi_i \sin \theta_i, \cos \theta_i), \\ \bar{\mathbf{P}}_i &= -\bar{P}_i(\cos \bar{\varphi}_i \sin \bar{\theta}_i, \sin \bar{\varphi}_i \sin \bar{\theta}_i, \cos \bar{\theta}_i).\end{aligned}\tag{8.3}$$

Furthermore, we can derive the Poisson brackets $\{P_{ia}, P_{ib}\} = \varepsilon_{abc} P_{ic}$ for $a, b, c = x, y, z$ and $\{\bar{P}_{ia}, \bar{P}_{ib}\} = -\varepsilon_{abc} \bar{P}_{ic}$ for $a, b, c = x, y, z$.² Since \mathbf{P}_i and $-\bar{\mathbf{P}}_i$ correspond to the angular momenta, we can define the total angular momentum to be $\mathbf{P} = \sum_{i=1}^N \mathbf{P}_i - \bar{\mathbf{P}}_i$. Similarly, we can identify $\omega_i \mathbf{P}$ and $\omega_i \bar{\mathbf{P}}$ with the magnetic moments, and we can define the total magnetic moment to be $\mathbf{M} = \sum_{i=1}^N \omega_i \mathbf{P}_i + \omega_i \bar{\mathbf{P}}_i$. Let us define

$$\mathcal{H} \equiv \mathbf{B} \cdot \mathbf{M} + \frac{\mu}{4} \sum_{i,j=1}^N (\mathbf{P}_i - \bar{\mathbf{P}}_i) \cdot (\mathbf{P}_j - \bar{\mathbf{P}}_j) (1 - \vec{v}_i \cdot \vec{v}_j).\tag{8.4}$$

With this Hamiltonian, we can recover Equation (8.2) from the Poisson bracket formulation of Hamilton's equations, $\dot{f} = \{f, \mathcal{H}\}$.

With the formulation in Hamiltonian mechanics, we have identified \mathcal{H} as a conserved quantity. Knowing \mathcal{H} it is also easy to show that the projection of the total angular momentum on the mass vector, $\mathbf{P} \cdot \mathbf{B}$, is conserved. There is, however, one caveat when considering the system as Hamiltonian. Naively, the system seems to be $6N$ dimensional since there is a three dimensional polarization vector for each neutrino and antineutrino, but as we saw, the phase space is actually only $4N$ dimensional, and this becomes important when we later interpret the Lyapunov exponents.

8.2.3 Two momentum modes

The specific model, we will consider, has only two momentum states $\vec{p}_1 = -\vec{p}_2$ [51], and for each of these momentum states, we define the sums $\mathbf{S}_i = \mathbf{P}_i + \bar{\mathbf{P}}_i$ and the differences

² ε_{abc} is the Levi-Civita symbol




Figure 8.1: Illustrations of different solutions to Equation (8.6). 1a and 1b show the normal hierarchy with symmetrical initial conditions. 2a and 2b show the inverted hierarchy with symmetric initial conditions. 3a and 3b show the normal hierarchy with anti-symmetric initial conditions.

$\mathbf{D}_i = \mathbf{P}_i - \bar{\mathbf{P}}_i$. In this model the total angular momentum, the total magnetic moment, and the conserved quantities from the Hamiltonian formulation are

$$\begin{aligned}
 \mathbf{P} &= \mathbf{D}_+ \equiv \mathbf{D}_1 + \mathbf{D}_2, \\
 \mathbf{M} &= \omega \mathbf{S}_+ \equiv \omega (\mathbf{S}_1 + \mathbf{S}_2), \\
 \mathcal{H} &= \omega \mathbf{B} \cdot \mathbf{S}_+ + \mu \mathbf{D}_1 \mathbf{D}_2, \\
 \mathbf{B} \cdot \mathbf{P} &= \mathbf{B} \cdot \mathbf{D}_+ = D_{+z}.
 \end{aligned} \tag{8.5}$$

The equations of motion can now be found from either \mathcal{H} or from Equation (8.2), and we get

$$\begin{aligned}
 \dot{\mathbf{S}}_1 &= \omega \mathbf{B} \times \mathbf{D}_1 + \mu \mathbf{D}_2 \times \mathbf{S}_1, \\
 \dot{\mathbf{S}}_2 &= \omega \mathbf{B} \times \mathbf{D}_2 + \mu \mathbf{D}_1 \times \mathbf{S}_2, \\
 \dot{\mathbf{D}}_1 &= \omega \mathbf{B} \times \mathbf{S}_1 + \mu \mathbf{D}_2 \times \mathbf{D}_1, \\
 \dot{\mathbf{D}}_2 &= \omega \mathbf{B} \times \mathbf{S}_2 + \mu \mathbf{D}_1 \times \mathbf{D}_2,
 \end{aligned} \tag{8.6}$$

which are the equations we will solve numerically.

In all of the following, we use $\omega = \pm 1$ for the two different mass hierarchies and $\mu = 6$ in some arbitrary units, and we choose to describe a pure electron neutrino beam by $\mathbf{S}_i = 2(\sin(2\theta), 0, \cos(2\theta))$, where we use the mixing angle $\sin(2\theta) = 0.1$. However, before we come to the numerical results, we will briefly review the simplest solutions.

8.2 \mathfrak{P} Chaotic flavor evolution in an interacting neutrino gas

If we assume that the two momentum states have the same initial conditions, the equations reduce to

$$\begin{aligned}\dot{\mathbf{S}} &= (\omega\mathbf{B} - \mu\mathbf{S}) \times \mathbf{D}, \\ \dot{\mathbf{D}} &= \omega\mathbf{B} \times \mathbf{S}.\end{aligned}\tag{8.7}$$

This system is equivalent to the isotropic case [49], and it supports two different simple solutions depending on the sign of ω , assuming that $\mu > |\omega|$ which corresponds to neutrino-neutrino interactions dominating vacuum oscillations. For the normal hierarchy where $\omega > 0$, a small initial S_x -value will make D_y negative which in turn decreases the value of S_x as it is seen in Figure 8.1, panel 1a and 1b. This results in an oscillatory motion around the z -axis for \mathbf{S} and oscillations on the y -axis for \mathbf{D} . Since this solution only deviates slightly from the initial state, we will call it the stationary solution.

For the inverted hierarchy where $\omega < 0$, the same initial condition will make D_y positive. The positive value of D_y will enhance the growth of S_x since $\mu > |\omega|$ making the configuration unstable, and we get a full inversion of the polarization vector, which is illustrated in panel 2a and 2b of Figure 8.1. In the literature this type of motion has been compared to an inverted pendulum, and the oscillations are called bipolar since D_y attains a significant value and separates the polarization vectors describing neutrinos and antineutrinos [49].

Let us now go beyond the symmetry assumption and consider the two momentum states separately in the normal hierarchy. For an anti-symmetric initial condition where all components are zero except for $S_{1x} = -S_{2x} > 0$ and $S_{1z} = S_{2z}$, we see that D_{1y} will become negative whereas D_{2y} will become positive. Since $\mu > |\omega|$, we find that D_{2y} dominates $\dot{\mathbf{S}}_1$ while D_{1y} dominates $\dot{\mathbf{S}}_2$ making both unstable (Figure 8.1, panel 3a and 3b) as it happened in the inverted hierarchy for the symmetric initial conditions. All of this results in bipolar oscillations as it did for the inverted hierarchy with symmetric initial conditions. In a similar way one can see that the inverted hierarchy will give an approximately stationary solution with the anti-symmetric initial conditions.

While these simple trajectories are solutions to the system, we will also investigate how modifications to the initial conditions turn out to give much more complicated trajectories in polarization space.

8.2.4 Numerical results

We have solved the equations of motion along with the equations describing the Lyapunov exponents and covariant Lyapunov vectors for four different types of trajectories. However, before we present these results, we will present some of the results which are common for all the different trajectories.

For all the cases, we find numerically that the spectrum of Lyapunov exponents has the form

$$(\lambda_1, \lambda_2, 0, 0, 0, 0, 0, 0, 0, -\lambda_2, -\lambda_1)\tag{8.8}$$

for the two beam model. This is also what one would expect due to the Hamiltonian nature of the system. We expect to see a symmetric spectrum as we already mentioned in Section 6.2, but it is a little more involved to argue for all of the zeros. Since we do our calculations using the polarization vectors, we have 3/2 times as many variables as the canonical Hamiltonian formulation. Therefore, four of the zeros in the spectrum actually relate to the constraints from the constant lengths of \mathbf{P}_i and $\bar{\mathbf{P}}_i$ rather than to

No δ added		
	$S_{1x} = -S_{2x}$	$S_{1x} = S_{2x}$
$\lambda_{1,\text{NH}}$	$0.99697 \pm 2 \cdot 10^{-6}$	$3.3124 \pm 3 \cdot 10^{-5}$
$\lambda_{2,\text{NH}}$	$0.5448 \pm 6 \cdot 10^{-6}$	$3.3054 \pm 3 \cdot 10^{-5}$
$\lambda_{1,\text{IH}}$	$3.3124 \pm 6 \cdot 10^{-6}$	$0.5448 \pm 2 \cdot 10^{-5}$
$\lambda_{2,\text{IH}}$	$3.3026 \pm 6 \cdot 10^{-6}$	0.0006 ± 0.0003

δS_z		
	$S_{1x} = -S_{2x}$	$S_{1x} = S_{2x}$
$\lambda_{1,\text{NH}}$	1.19 ± 0.02	1.234 ± 0.005
$\lambda_{2,\text{NH}}$	0.76 ± 0.02	0.794 ± 0.004
$\lambda_{1,\text{IH}}$	0.76 ± 0.02	0.753 ± 0.02
$\lambda_{2,\text{IH}}$	0.045 ± 0.003	0.054 ± 0.002

δS_y		
	$S_{1x} = -S_{2x}$	$S_{1x} = S_{2x}$
$\lambda_{1,\text{NH}}$	0.97 ± 0.01	0.973 ± 0.004
$\lambda_{2,\text{NH}}$	0.53 ± 0.01	0.521 ± 0.005
$\lambda_{1,\text{IH}}$	0.68 ± 0.03	0.720 ± 0.005
$\lambda_{2,\text{IH}}$	0.062 ± 0.003	0.082 ± 0.002

Table 8.1: Lyapunov exponents for the normal hierarchy (NH) and the inverted hierarchy (IH) were calculated as described in Appendix 8.2.A.

any conserved quantity. The other four zeros, however, correspond to our two conserved quantities; \mathcal{H} and $\mathbf{P} \cdot \mathbf{B}$.

The many zeros and the symmetry reduce the Lyapunov spectrum to only two interesting numbers; λ_1 and λ_2 . The values of these depend on ω , μ , and θ , but in this paper our goal is not to map out this dependence. We would also like to remark that we do not need to know the values of λ_i with very high precision. It is the order of magnitude we are interested in, and therefore it is not crucial to have a very stringent error estimate either.

Our calculated Lyapunov exponents are seen in Table 8.1, where we show λ_1 and λ_2 . We also give an estimate of the uncertainty on our numbers, but note that these are not stringent standard deviations due to some issues with correlated data which we discuss further in Appendix 8.2.A.

For all the cases we have studied, the leading Lyapunov exponent is positive, indicating chaotic behavior. This might seem strange for the stationary and periodic orbits, but for these orbits it is merely a statement of instability.

An example of the covariant Lyapunov vectors is shown in Figure 8.2. For each coordinate, the average magnitude of that component is shown for every vector. We have computed the covariant Lyapunov vectors for 100000 time steps, but in order to ensure that the computation have actually converged both forward and backward, we skip the first and last 20% when doing the averages.³ From the Figure it is clear that the first covariant Lyapunov vector (γ_1) and the last (γ_{12}) as well as the second (γ_2)

³We have also tried to skip 40% which gives the same result, so 20% is sufficient to ensure convergence.

8.2 \mathfrak{B} Chaotic flavor evolution in an interacting neutrino gas

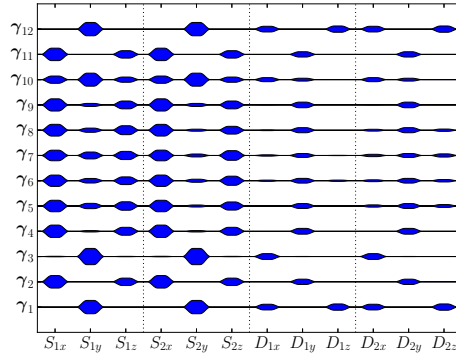


Figure 8.2: All covariant Lyapunov vectors for the normal hierarchy with no modification added to the initial conditions and $S_{1x} = -S_{2x}$ which is the bipolar case. The average magnitude of each component is shown for every vector. The height of the colored area shows how much of the component given on the first axis is present in the vector given on the second axis. e.g. γ_1 has equally large components in the $S_{1,y}$ - and $S_{2,y}$ -directions while the components of γ_1 in the $S_{1,x}$ -, $S_{1,z}$ -, $S_{2,x}$ -, and $S_{2,z}$ -directions are zero.




Figure 8.3: The trajectory for the normal hierarchy in the stationary case projected on \mathbf{P}_1 is shown in black. Three sets of covariant Lyapunov vectors are also shown. Blue is γ_1 , red is γ_2 , greens are γ_3 to γ_{10} , magenta is γ_{11} , and cyan is γ_{12} .

and the second last (γ_{11}) point in similar directions. We find this to be the case for all our calculations, so we will only be interested in γ_1 and γ_2 from now on.

8.2.4.1 Stationary solutions

The simplest trajectories, we will consider, are the stationary solutions where the polarization vectors stay very close to the z-axis as seen in Figure 8.3. Although we call this the stationary case, we must remember that the solution is only approximately




Figure 8.4: Covariant Lyapunov vectors for the static case (two upper and two lower vectors) and the bipolar case (central four vectors) in the $\mathbf{S}_{1,2}, \mathbf{D}_{1,2}$ coordinates. The average magnitude of each component is shown for every vector. Consult Figure 8.2 for notes on how to read the figure.




Figure 8.5: Covariant Lyapunov vectors for the static case (two upper and two lower vectors) and the bipolar case (central four vectors) in the $\mathbf{S}_{\pm}, \mathbf{D}_{\pm}$ coordinates. The average magnitude of each component is shown for every vector. Consult Figure 8.2 for notes on how to read the figure.

stationary, and it turns out that the small deviation from a genuinely stationary solution will effect some of the quantities we calculate.

For the stationary cases (that is $S_{1x} = S_{2x}$ for the normal hierarchy and $S_{1x} = -S_{2x}$ for the inverted hierarchy), we find that $\lambda_1 \approx \lambda_2 = 3.31$. This means that any linear combination of γ_1 and γ_2 should be considered a covariant Lyapunov vector for the highest Lyapunov exponent. A small sample of covariant Lyapunov vectors is shown in Figure 8.3, but it is hard to find any structures when depicting them in this way. Instead we will consider averages of the lengths of each component over time. These averages can be seen in Figure 8.4 for all the perfectly symmetric and anti-symmetric cases, and due to the symmetry between the two momentum states, we also transform the vectors to the $\{+, -\}$ -basis where $\mathbf{S}_{\pm} = \mathbf{S}_1 \pm \mathbf{S}_2$ and $\mathbf{D}_{\pm} = \mathbf{D}_1 \pm \mathbf{D}_2$ in Figure 8.5.

8.2 \mathfrak{B} Chaotic flavor evolution in an interacting neutrino gas

The stationary case is the lower two and the upper two vectors in each figure. In the $\{1, 2\}$ -basis, the two sets of vectors look very similar, but in the $\{+, -\}$ -basis, the normal hierarchy has only a S_- and D_- component while the inverted hierarchy has only a S_+ and D_+ component. It is also possible to get some analytical insight since the covariant Lyapunov vectors coincide with the normal modes of an ordinary stability analysis when the trajectory is stationary.

Inspired by the stability analysis Duan [162] did on the two beam model, we transform Equation (8.6) to the $\{+, -\}$ -basis:

$$\begin{aligned}
 \dot{\mathbf{S}}_+ &= \omega \mathbf{B} \times \mathbf{D}_+ + \frac{\mu}{2} \mathbf{D}_+ \times \mathbf{S}_+ - \frac{\mu}{2} \mathbf{D}_- \times \mathbf{S}_- \\
 &\approx (\omega + \mu) \mathbf{B} \times \mathbf{D}_+, \\
 \dot{\mathbf{S}}_- &= \omega \mathbf{B} \times \mathbf{D}_- + \frac{\mu}{2} \mathbf{D}_+ \times \mathbf{S}_- - \frac{\mu}{2} \mathbf{D}_- \times \mathbf{S}_+ \\
 &\approx (\omega - \mu) \mathbf{B} \times \mathbf{D}_-, \\
 \dot{\mathbf{D}}_+ &= \omega \mathbf{B} \times \mathbf{S}_+, \\
 \dot{\mathbf{D}}_- &= \omega \mathbf{B} \times \mathbf{S}_- + \mu \mathbf{D}_+ \times \mathbf{D}_- \\
 &\approx \omega \mathbf{B} \times \mathbf{S}_-,
 \end{aligned} \tag{8.9}$$

where we have used the approximations $\mathbf{S}_+/2 \approx \mathbf{S}_1 \approx \mathbf{S}_2 \approx -\mathbf{B}$, and \mathbf{D}_+ , \mathbf{D}_- , and \mathbf{S}_- are small, so some quadratic terms can be neglected.

From these equations it is clear that $+$ and $-$ decouple, and we find

$$\begin{aligned}
 \ddot{\mathbf{S}}_+ &\approx (\omega + \mu) \omega \mathbf{B} \times (\mathbf{B} \times \mathbf{S}_+) = -\omega(\omega + \mu) \mathbf{S}_+, \\
 \ddot{\mathbf{D}}_+ &\approx -\omega(\omega + \mu) \mathbf{D}_+, \\
 \ddot{\mathbf{S}}_- &\approx -\omega(\omega - \mu) \mathbf{S}_-, \\
 \ddot{\mathbf{D}}_- &\approx -\omega(\omega - \mu) \mathbf{D}_-.
 \end{aligned} \tag{8.10}$$

For the normal hierarchy, we get the solutions

$$\begin{aligned}
 \mathbf{S}_+ &= \mathbf{a}_1 e^{\pm i t k_+}, \quad \mathbf{D}_+ = \mathbf{a}_2 e^{\pm i t k_+}, \quad k_+ = \sqrt{\omega(\mu + \omega)} \\
 \mathbf{S}_- &= \mathbf{a}_3 e^{\pm i t k_-}, \quad \mathbf{D}_- = \mathbf{a}_4 e^{\pm i t k_-}, \quad k_- = \sqrt{\omega(\mu - \omega)}.
 \end{aligned} \tag{8.11}$$

This suggest that $\lambda_1 = \lambda_2 = -\lambda_{11} = -\lambda_{12}$, and that the covariant Lyapunov vectors should point towards \mathbf{S}_- and \mathbf{D}_- in the normal hierarchy as it is seen for the two upper vectors in Figure 8.5. S_{-z} and D_{-z} are almost zero since both their derivatives are approximated by $\propto \mathbf{B} \times \mathbf{X}$. As $\mathbf{B} = (0, 0, -1)$, the derivative in the z -direction is zero. The small deviations from zero are due to the fact that the simulated system is not perfectly stationary. If we set $\sin^2(2\theta) = 0$, we find the two z -components to be exactly zero.

A similar analysis can be done for the inverted hierarchy. Here the sign of ω is opposite, and the solutions are

$$\begin{aligned}
 \mathbf{S}_+ &= a_1 e^{\pm i t k_+}, \quad \mathbf{D}_+ = a_2 e^{\pm i t k_+}, \quad k_+ = \sqrt{-\omega(\mu + \omega)} \\
 \mathbf{S}_- &= a_3 e^{\pm i t k_-}, \quad \mathbf{D}_- = a_4 e^{\pm i t k_-}, \quad k_- = \sqrt{-\omega(\mu - \omega)}.
 \end{aligned} \tag{8.12}$$

Again this is consistent with the numerical result in Figure 8.5.

From a more intuitive point of view, we notice that we have $S_{1x} = S_{2x}$ in the initial condition for the normal hierarchy, and this turns out to hold true for all times. Therefore, any perturbation acting symmetrically on S_{1x} and S_{2x} or S_{1y} and S_{2y} will conserve the symmetry of the system. On the contrary perturbations acting anti-symmetrically will break the symmetry. These two cases correspond to perturbations in S_{+x} and S_{+y} versus S_{-x} and S_{-y} respectively, so γ_1 and γ_2 must point in the directions of S_{-x} and S_{-y} as we also find. In the inverted hierarchy $S_{1x} = -S_{2x}$, and anti-symmetric perturbations will conserve the symmetry whereas symmetric perturbations will break the symmetry of the system, so γ_1 and γ_2 must point in the directions of S_{+x} and S_{+y} . All of this is consistent with the numerical and analytical results.

With a better understanding of the covariant Lyapunov vectors, we will now consider the stability of the system. For the positive and negative Lyapunov exponents, we note that the stable and unstable directions in polarization space are coincident according to the covariant Lyapunov vectors. This is also the result in our algebraic analysis where we find both positive and negative exponentials to solve the differential equations. Consequently, it is impossible to find a set of converging solutions as the diverging solution will always dominate. Apart from the positive and negative Lyapunov exponents, we found eight vanishing exponents which we can interpret in terms of constraints and conserved quantities. With regard to stability, however, their associated covariant Lyapunov vectors indicate marginally stable directions in which perturbation will neither shrink nor grow on average. From the intuitive point of view, we find that the marginally stable directions correspond to perturbations which do respect the symmetry of the system. Also, it turns out that all the z -components relate to vanishing Lyapunov exponents in the perfectly stationary case. For the more realistic case where $\theta \neq 0$, however, we find that the directions of S_{-z} and D_{-z} are unstable as well.

A more physical interpretation of the perturbations is slightly hampered by the fact that we have chosen the z -axis along the direction of \mathbf{B} and not in the direction of the pure flavor state. In order to recover the more interpretable coordinate system, we would have to rotate all the vectors with the angle θ in the (x, z) -plane. This means that whenever we encounter an x -component it actually contains a little of the flavor z -component and vice versa. Apart from this minor complication, a perturbation in any z -coordinate corresponds to a perturbation in the flavor content of the neutrinos while perturbations in the x - and y -coordinates correspond to perturbations in the phase of the neutrino oscillations. From this point of view, a perfectly stationary system with vanishing mixing angle would be marginally stable towards perturbations in the flavor content but unstable towards anything that could shift the phase. For the case of a non-vanishing mixing angle, the inverted hierarchy will be unstable towards flavor perturbations although the components of the covariant Lyapunov vectors are small. The normal hierarchy will also be unstable towards anti-symmetric perturbations but marginally stable towards symmetric flavor perturbations as it is seen in Figure 8.5.

8.2.4.2 Bipolar solutions

For the bipolar case, we get the well known periodic solutions where the polarization vectors oscillate from $P_z > 0$ to $P_z < 0$ [49–51] as seen in Figure 8.6. In the normal hierarchy, this solution is obtained when $S_{1x} = -S_{2x}$, and in the inverted hierarchy, we




Figure 8.6: The trajectory for the normal hierarchy in the bipolar case projected on \mathbf{P}_1 . Three sets of covariant Lyapunov vectors are also shown. Blue is γ_1 , red is γ_2 , greens are γ_3 to γ_{10} , magenta is γ_{11} , and cyan is γ_{12} .

find it for $S_{1x} = S_{2x}$ as we described in Section 8.2.1. For the Lyapunov exponents, we get the values $\lambda_1 = 0.997$ and $\lambda_2 = 0.545$ in the normal hierarchy, while the inverted hierarchy gives the values $\lambda_1 = 0.545$ and $\lambda_2 = 0.0006 \approx 0$. This is remarkable since it suggests that $\lambda_{2,\text{NH}} = \lambda_{1,\text{IH}}$, and it shows the existence of two more vanishing Lyapunov exponents for the inverted hierarchy. The vanishing Lyapunov exponents could suggest that there is another conserved quantity, but since they only vanish for the symmetric and not for the anti-symmetric initial conditions, it is probably rather an artifact of the specific bipolar solution. When the values are compared to the stationary case, we note that the first two Lyapunov exponents are not degenerate any more and that $\lambda_{\text{bipolar}} < \lambda_{\text{stationary}}$.

Regarding stability analysis, this suggests the stationary cases to be more unstable than the bipolar ones. This also means that a perturbation in the normal hierarchy will need three times longer to grow by the same factor in the bipolar case than in the stationary cases. For the inverted hierarchy, it will need six times as long. If this result transfers to real physical systems, it can have an important impact since these perturbations will grow only while μ is large. In a supernova, μ becomes smaller as you go away from the center of the supernova, and in the early Universe, μ decays with the expansion of the Universe. This limits the time a perturbation has to grow, and the value of the Lyapunov exponents can thus determine if a small perturbation becomes large and makes the trajectory non-periodic.

We will now turn to the covariant Lyapunov vectors. Again we see a sample of the trajectory in Figure 8.6, but we still find the averages to be more interesting. When we consider the four central covariant Lyapunov vectors in Figure 8.4, the pattern from the Lyapunov exponents is repeated as $\gamma_{1,\text{IH}}$ is very similar to $\gamma_{2,\text{NH}}$. On the other hand, there is no information in $\gamma_{2,\text{IH}}$ since its Lyapunov exponent is 10 times degenerate. Going to Figure 8.5, it is only the perturbations breaking the symmetry which actually grow as we also saw for the stationary solution. We see that the normal hierarchy with the initial condition $S_{1x} = -S_{2x}$ is stable towards anti-symmetric perturbations (S_{-x} and S_{-y}) but unstable with regards to symmetric perturbations (S_{+x} and S_{+y}). In the




Figure 8.7: The trajectory for the normal hierarchy in the mildly chaotic case projected on \mathbf{P}_1 . Notice how the trajectory to some degree follows the bipolar solution. The change from $P_{1y} > 0$ to $P_{1y} < 0$ happens very rarely, but it is seen in this example. Three sets of covariant Lyapunov vectors are also shown. Blue is γ_1 , red is γ_2 , greens are γ_3 to γ_{10} , magenta is γ_{11} , and cyan is γ_{12} .

same way, the inverted hierarchy with $S_{1x} = S_{2x}$ is stable towards perturbations in S_{+x} and S_{+y} but unstable with regards to perturbations in S_{-x} and S_{-y} .

As for the stationary case, we can interpret the missing components of γ_1 and γ_2 (in the normal hierarchy) as directions in polarization space far more stable against perturbations than the other directions. We see that perturbations in S_{1y} , S_{2y} , D_{1x} , D_{1z} , D_{2x} , and D_{2z} are marginally stable in the inverted hierarchy, while their exponential growth is approximately twice as fast as that of other perturbations in the normal hierarchy.

In a physical interpretation, this is interesting since it shows that small symmetric perturbations in the flavor content or the phase will not be important in the inverted hierarchy if the initial conditions are approximately symmetrical. For the normal hierarchy, it is tempting to draw the same conclusion regarding the flavor content, but here we must remember that the non-zero x -component also contains some of the flavor z -component.

8.2.4.3 Non-periodic solutions

While the bipolar case and the stationary case have been studied for about a decade [47, 49, 50, 163, 164], the interest in the chaotic, non-periodic solutions is quite recent [51, 162].

The non-periodic solutions we have considered are obtained by taking the initial conditions corresponding to the stationary and bipolar solutions and add $\delta = 2 \cdot 10^{-3}$ to one of the coordinates. This modification is large enough to make the marginally stable directions non-periodic, but if there is a difference between $S_{1x} = S_{2x}$ and $S_{1x} = -S_{2x}$, we still expect to see it. With this approach, we have found two different types of non-periodicity. For some small modifications of the stationary and periodic cases, we get a trajectory which is not recurrent but stays close to the periodic solution known




Figure 8.8: The trajectory for the normal hierarchy in the highly chaotic case projected on \mathbf{P}_1 . Three sets of covariant Lyapunov vectors are also shown. Blue is γ_1 , red is γ_2 , greens are γ_3 to γ_{10} , magenta is γ_{11} , and cyan is γ_{12} .

from the bipolar case as it is seen in Figure 8.7. A property of this group of solutions is that the trajectory stays in the S_{ix} - S_{iz} planes and along D_{iy} . These coordinates are also the ones where we can add our δ without making the trajectory even more non-periodic. If we modify S_{iy} , D_{ix} , or D_{iz} , we get a trajectory which eventually covers the full polarization space fulfilling that $|\mathbf{P}_i|$ and $|\bar{\mathbf{P}}_i|$ are conserved as seen in Figure 8.8. How fast it will deviate significantly from the bipolar oscillations depends on which mass hierarchy we consider, and we will return to this point when discussing the covariant Lyapunov vectors.

When doing the Lyapunov analysis, we have chosen to modify δS_z to represent the mildly non-periodic case and modify δS_y to represent the most chaotic case. Modifying all the other coordinates give results similar to either one or the other. We find the covariant Lyapunov vectors to be very similar within each group while the Lyapunov exponents are within $\sim 20\%$ for each group.

The first Lyapunov exponents for the non-periodic cases, λ_1 , range from 0.68 to 1.23, so there is no large difference in how fast perturbations grow in the dominantly unstable directions. The second Lyapunov exponents, λ_2 , however, shows a significant difference. In the normal hierarchy, the second Lyapunov exponents range from 0.52 to 0.79, but in the inverted hierarchy, the largest λ_2 is 0.082. This is not much larger than zero, and it indicates that there are directions which are almost marginally stable in the inverted hierarchy.

Comparing λ_1 to the stationary and bipolar cases, we find them to be approximately one third of the leading Lyapunov exponents in the stationary cases but quite similar to the bipolar case. As in the bipolar case, this means that a certain perturbation would need about three times longer to affect the solution in the non-periodic cases than it would in the stationary cases.

For the non-periodic cases it is even harder to digest the covariant Lyapunov vectors when shown along the trajectory due to its complicated nature. A few examples are shown in Figure 8.7 and Figure 8.8, but again we need to consider the averages to find patterns. The averages of the first and second covariant Lyapunov vectors are




Figure 8.9: Covariant Lyapunov vectors for modifications in the δS_z direction in the $\mathbf{S}_{1,2}, \mathbf{D}_{1,2}$ coordinates. The average magnitude of each component is shown for every vector. Consult Figure 8.2 for notes on how to read the figure.




Figure 8.10: Covariant Lyapunov vectors for modifications in the δS_y direction in the $\mathbf{S}_{1,2}, \mathbf{D}_{1,2}$ coordinates. The average magnitude of each component is shown for every vector. Consult Figure 8.2 for notes on how to read the figure.

seen in Figure 8.9 for δS_z and in Figure 8.10 for δS_y . As in the previous cases, we have also tried to plot the vectors in the $\{+, -\}$ -basis, but it does not provide any additional information except for the fact that the D_{+z} -component is zero. This is what we expect for generalized normal modes since $-D_{+z}$ is the projection of the total angular momentum on \mathbf{B} which is a conserved quantity. If we however plot the forward singular vectors in the $\{+, -\}$ -basis, we find the D_{+z} -component to be non-zero for the δS_y cases. For δS_z we find the same for γ_1 in normal hierarchy and γ_2 in inverted hierarchy. This means that a small modification in the D_{+z} -direction will lead to a diverging solution. Not that it will diverge in the D_{+z} -direction, which is conserved, but the non-linear evolution will transfer the difference to other non-conserved coordinates.

If we now go back to δS_z in Figure 8.9, we see some structure. λ_1 and λ_2 are comparable in size and of order one for the normal hierarchy, and the directions associated with γ_1 will only diverge a little faster than directions associated with γ_2 .

As a result, the difference between γ_1 and γ_2 will not have large consequences in the normal hierarchy. For the inverted hierarchy, on the other hand, $\lambda_2 \sim 0$ while $\lambda_1 \sim 1$. This means that directions associated with γ_2 will diverge much slower than directions associated with γ_1 . If we go back and compare the non-zero coordinates of γ_2 to the coordinates which give rise to the more chaotic solution exemplified by the δS_y case, we find them to be identical. Interestingly enough, we also find these highly non-periodic trajectories to be slower filling out the phase space for the inverted hierarchy than for the normal hierarchy. We believe this to be a remnant from the low value of λ_2 in the less chaotic δS_z case. For the stability of the system, this vanishing components of γ_1 in the inverted hierarchy indicate marginally stable directions. From a physical point of view, however, it is not clear how to perturb only S_{1x} and not S_{1y} , so all the phases must be regarded unstable. On the other hand, a flavor perturbation with opposite effects on neutrinos and anti neutrinos would be marginally stable since the D_{ix} -components are vanishing as well as the D_{iz} -components, and this type of perturbation might be possible.

We will now turn to δS_y in Figure 8.10 where we see that there is very little information to be gained. All components are present in all vectors. This is understandable since the trajectories cover all of the allowed polarization space which means we average over vectors pointing in all directions.

8.2.5 Conclusions

Our analysis has shown a way to generalize the linear stability analysis to periodic and even non-periodic solutions of a set of differential equations. The Lyapunov exponents quantifies how fast a small perturbation to a known solution can be expected to grow, and the covariant Lyapunov vectors generalize the normal modes from the stationary case and contain information about the dynamics of a given trajectory.

In the simple two beam model we have considered here, we have shown that the stationary case with very little flavor conversion is more unstable than the bipolar flavor changing case since the Lyapunov exponents are larger. Furthermore the covariant Lyapunov vectors show that some directions are marginally stable, so that perturbations confined to these directions will need a very long time to grow significantly if they will grow at all. This is for example the case for any perturbations in S_{+z} where the flavor content of all neutrinos and anti-neutrinos are perturbed by the same amount.

For the non-periodic trajectories, we have investigated two different types of variations to the stationary and bipolar cases which lead to quite different behavior of the polarization vectors. In the less chaotic case, the trajectory resembles the bipolar solution to some degree, and some of the symmetries in the equations are unbroken. Again we found that the second Lyapunov exponent was quite small in the inverted hierarchy, and for the less chaotic case this means that perturbations in some directions grow very slowly. In the more chaotic case we struggled to find any structure in the covariant Lyapunov vectors at all, although they do have a tendency to be orthogonal to the polarization vectors which would also be expected since the lengths of the polarization vectors are conserved.

Physically the consequence of a large Lyapunov exponent is that a small perturbation introduced by a thermal fluctuation, an anisotropy, or an inhomogeneity might grow very fast. This is the case if the perturbation has a component of the Lyapunov vector

corresponding to the large Lyapunov exponent. Therefore, the difference we have found between the Lyapunov exponents for the inverted and normal hierarchy is interesting. Most notably for the unperturbed symmetric case, where we find the Lyapunov exponent in the normal hierarchy to be six times as large as the Lyapunov exponent in the inverted hierarchy. This is of course also comparing a stationary to a bipolar case which really are two different solutions, but even for the non-periodic cases there is a difference. The consequence is that perturbations in the normal hierarchy grow faster than similar perturbations in the inverted hierarchy. If this pattern is also found in more realistic models, it might lead to significant differences between the normal and the inverted hierarchy in the early Universe or a supernova. In these environments we find a decaying neutrino background potential with time or radius respectively, and this decay limits the region where perturbations can grow. Therefore the growth rate will determine if the perturbation becomes large and changes the observable signatures.

Finally, while the results for this two beam model is of limited use when considering real physical systems, the Lyapunov analysis highlights that there is a lot of information about the stability of such a system to be found beyond a simple stationary linearization.

Acknowledgments

We would like to thank Yvonne Wong for comments on the manuscript and Georg Raffelt for numerous valuable discussions and critical comments on earlier versions of the manuscript.

8.2.A Numerical calculation of Lyapunov exponents and covariant Lyapunov vectors

The first numerical algorithm to calculate the full spectrum of Lyapunov exponents was proposed by Benettin et al. [134, 165]. The leading exponent can easily be calculated by simply choosing a random vector in the tangent space and evolving it according to Equation (6.2). The only complication is that the components of the vector can exceed the value admissible for a float or a double. The ease of finding the leading Lyapunov exponent is also the curse for finding any of the other exponents. If the equations can be inverted, it is possible to find the smallest Lyapunov exponent from the inverted dynamics, but everything in between needs a trick.

The trick is to orthogonalize the vectors before they collapse into the most unstable direction. This can be done using a standard **QR**-decomposition where the unitary **Q**-matrix contains the orthonormalized vectors, and the diagonal of the upper triangular **R**-matrix contains the lengths of the orthogonalized vectors before they are normalized. These lengths are exactly what is needed in order to compute the Lyapunov exponents. For a trajectory divided in k sections, the Lyapunov exponents are given by

$$\lambda_i \approx \frac{1}{t} \ln \left(\prod_{j=1}^k r_{jj} \right) = \frac{1}{t} \sum_{j=1}^k \ln r_{jj}. \quad (8.13)$$

This is a finite time version of Equation (6.4) where $\|v(t_0)\| = 1$ since we orthonormalize the vectors in each step.

It turns out that in many cases, the convergence of λ_i is quite slow. In order to improve this, Goldhirsch et al. [139] showed that the error depends on time as $1/t$. This




Figure 8.11: The asymptotic value of λ_2 is found at $1/t = 0$. The blue line shows the Lyapunov exponents obtained from the forward calculation for different time intervals while the red line shows the Lyapunov exponents obtained from the backward calculation. The black lines are fits where the first 300 time steps have been excluded. This plot is for the inverted hierarchy with the modification δS_z and $S_{1x} = -S_{2x}$, and with the procedure for calculating the error described in the text, we find $\lambda_2 = 0.045 \pm 0.003$.




Figure 8.12: The asymptotic value of λ_2 is found at $1/t = 0$. The blue line shows the Lyapunov exponents obtained from the forward calculation for different time intervals while the red line shows the Lyapunov exponents obtained from the backward calculation. The black lines are fits where the first 2500 time steps have been excluded. This plot is for the inverted hierarchy with the modification δS_y and $S_{1x} = -S_{2x}$, and with the procedure for calculating the error described in the text, we find $\lambda_2 = 0.062 \pm 0.003$ marked by a green point in the figure.

means that plotting $\lambda_i(t)$ versus $1/t$ should yield a straight line where the asymptotic value for λ_i is found at $1/t = 0$. An example of this is seen in Figure 8.11 where we show two sets of calculated Lyapunov exponents. As it can be seen in the Figure, the estimated asymptotic value will depend somewhat on the amount of data that is used when fitting the straight line. This is even more clear in Figure 8.12 where we show a smaller segment of the time axis. We have used this difference as a measure of the error in the computed Lyapunov exponents. The computed set of values contain 101000 time steps, and we have constructed a sample of different estimates of λ_i by calculating λ_i 500 times using the last 100900 to 51000 time steps. This is done for both the sets of values shown in Figure 8.11, and, finally, we calculated the mean and the standard deviation of this full sample to get $\lambda_i \pm \varepsilon$. For Figure 8.11 this gives 0.0455 ± 0.003 , and for Figure 8.12 it gives 0.062 ± 0.003 . Note that the different values in our samples are not independent, and thus the error we calculate should not be interpreted as a stringent standard deviation but rather as an indication of our level of precision. The lack of independence is partly due to the origin of the two different sets of values in Figure 8.11, but before we describe that, we need to understand how the covariant Lyapunov vectors are calculated.

The algorithm used to find the covariant Lyapunov vectors is a slight modification of the one presented by Kuptsov and Parlitz [135]. The idea builds mainly on what was proposed by Wolfe and Samelson [136] while Ginelli et al. [138] has an alternative approach.

The method we use to compute covariant Lyapunov vectors is to find the forward and backward singular vectors first and then calculate the covariant Lyapunov vectors. To find the backward singular vectors, we can consider an arbitrary vector \mathbf{v} in the tangent space. Almost any such vector will grow with the average rate λ_1 giving $|\mathbf{v}(t_2)| = |\mathbf{v}(t_1)| \exp(\lambda_1(t_2 - t_1))$. Using the time-reversed dynamics on $\mathbf{v}(t_2)$, we find

$$\mathbf{M}^{-1}(t_1, t_2)\mathbf{v}(t_2) = \mathbf{M}^{-1}(t_1, t_2)\mathbf{M}(t_1, t_2)\mathbf{v}(t_1) = \mathbf{v}(t_1), \quad (8.14)$$

so $\mathbf{v}(t_2)$ contracts at a rate $-\lambda_1$ under the backwards dynamics and therefore approaches $\mathbf{b}_1(t_2)$ as we use longer time intervals. We assume that all Lyapunov exponents are non-degenerate, but the generalization to the degenerate case is straight forward. Considering an area spanned by \mathbf{v} and another arbitrary vector \mathbf{u} , again for almost any choice of \mathbf{u} , it will grow with the rate $\lambda_1 + \lambda_2$. This means that the component of \mathbf{u} orthogonal to \mathbf{b}_1 must approach \mathbf{b}_2 by an argument similar to the one above. This process can be repeated, and we can find all the backward singular vectors in this way by induction. Going back to our method for calculating the Lyapunov exponents, it should be clear that a byproduct of the calculation is that the backward singular vectors end up as the columns in the \mathbf{Q} -matrix of the \mathbf{QR} -decomposition.

With a more involved argument, it is also possible to show that the forward singular vectors can be obtained by using $\mathbf{M}^T(t_1, t_2)$ and going backwards in time [135].

Having obtained the forward and backward singular vectors, we need a method to find the covariant Lyapunov vectors. Since the covariant Lyapunov vectors must respect both forward and backward dynamics, it must be possible to write the matrix $\mathbf{\Gamma}(t) = [\gamma_1 \dots \gamma_n]$ as

$$\mathbf{\Gamma}(t) = \mathbf{B}(t)\mathbf{A}^-(t) = \mathbf{F}(t)\mathbf{A}^+(t), \quad (8.15)$$

where $\mathbf{B}(t)$ and $\mathbf{F}(t)$ are the matrices of backward and forward singular vectors respectively. As the i 'th covariant Lyapunov vector must grow with only $-\lambda_i$ in the backwards

dynamics, it can only have components from \mathbf{b}_j with $\lambda_j \geq \lambda_i$. This means that $\mathbf{A}^-(t)$ can be chosen to be upper diagonal. Similarly $\mathbf{A}^+(t)$ can be chosen to be lower diagonal. Multiplying by $\mathbf{F}^T(t)$, we get the equation

$$\mathbf{F}^T(t)\mathbf{B}(t)\mathbf{A}^-(t) = \mathbf{A}^+(t) \Leftrightarrow \quad (8.16)$$

$$\mathbf{F}^T(t)\mathbf{B}(t) = \mathbf{A}^+(t)(\mathbf{A}^-(t))^{-1}, \quad (8.17)$$

which is a LU-factorization of $\mathbf{F}^T(t)\mathbf{B}(t)$. To find $\mathbf{A}^-(t)$, we can restrict our attention to the upper left j times j submatrix of $\mathbf{F}^T(t)\mathbf{B}(t)$ in Equation (8.16) and focusing on the j 'th column of $\mathbf{A}^\pm(t)$. This gives an equation of the form

$$\begin{pmatrix} x_{11} & x_{12} & \dots & x_{1j} \\ x_{21} & x_{22} & \dots & x_{2j} \\ \dots & \dots & \dots & \dots \\ x_{(j-1)1} & x_{(j-1)2} & \dots & x_{(j-1)j} \\ x_{j1} & x_{j2} & \dots & x_{jj} \end{pmatrix} \begin{pmatrix} a_{1j}^- \\ a_{2j}^- \\ \dots \\ a_{(j-1)j}^- \\ a_{jj}^- \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ \dots \\ 0 \\ a_{jj}^+ \end{pmatrix}. \quad (8.18)$$

Since the LU-factorization is unique only up to the diagonal of one of the matrices, we can eliminate the j 'th row in the above matrix equation resulting in a homogeneous system. Solving this to find $\mathbf{A}^-(t)$, we can find $\mathbf{\Gamma}(t)$ from Equation (8.15).

We now have all the tools to find the covariant Lyapunov vectors, but let us go back and get the full overview of the numerical algorithm. To control the divergence of the singular vectors, we solve the differential equations for short time steps dt (we use $dt = 0.05$) and use the result as the initial conditions for the next step. The calculation goes through four different phases:

1. A random unitary matrix initialize \mathbf{B} . The trajectory and propagator is found for each time step, and \mathbf{B} is evolved by using the propagator and finding the **QR**-factorization of the result. The \mathbf{R} -matrix diagonal is saved. (We used 1000 steps)
2. The trajectory and propagator is found for each time step, and \mathbf{B} is evolved as before. The \mathbf{R} -matrix diagonal, the trajectory, and \mathbf{B} are saved. (We used 100000 steps)
3. The trajectory is found and saved for each time step as preparation. A random unitary matrix initialize \mathbf{F} . Starting with the last time step, the trajectory and propagator is found from the previous time step and evolved to the current one. \mathbf{F} is then evolved backwards in time by using \mathbf{M}^T on \mathbf{F} and finding the **QR**-factorization of the result. The \mathbf{R} -matrix diagonal is saved. (We used 1000 steps)
4. Starting with the last time step from phase 2, \mathbf{F} is evolved backwards in time as described in phase 3. Knowing both \mathbf{B} and \mathbf{F} , Equation (8.15) and (8.18) gives the covariant Lyapunov vectors. The \mathbf{R} -matrix diagonal and the covariant Lyapunov exponents are saved. (We used 100000 steps as in phase 2)

In this way, we find the covariant Lyapunov vector, and it is possible to save the forward singular vectors if we are interested in those. Furthermore, we obtain two sets

8 Chaotic supernova neutrinos

of data from which we can estimate the Lyapunov exponents. The caveat here is that these two sets of data are not independent as they originate in the same trajectory. In order to eliminate this dependency, we would need to skip one of the data sets, and thereby we would lose some of our precision.



8.3 Further developments in the multi-azimuthal-angle instability and inhomogeneities in collective neutrino oscillations

Due to the highly simplified nature of the model analysed in the previous section, the implications of the results for our knowledge of real supernova neutrinos are limited. However, some interesting features are highlighted by the analysis such as the relative rates of divergence and the type of perturbations that will set off such divergences. Yet the most important message to take, is the applicability of Lyapunov analysis on supernova neutrinos. One interesting exercise is to compare the results obtained by Mirizzi [159] for a more realistic toy-model to the results from the previous section.

The model Mirizzi studies can basically be described as a neutrino emitting ring, which is a two dimensional analogue for a supernova. It is assumed that the neutrinos are only emitted in the two directions parallel to the surface of the ring, and in this way it is quite similar to the system we discussed. This is especially the case when the neutrinos are assumed to be emitted homogeneously from every point of the ring, whereas the calculations allowing for an inhomogeneity has distinctively different features. The homogeneous results show that for the normal hierarchy, a perturbation of the size 10^{-3} becomes important at a radius of 80km, while the larger perturbation of 10^{-2} needs 180km to become important in the inverted hierarchy. This means that the growth rate in the normal hierarchy was ~ 3.5 times larger than it was the case in the inverted hierarchy. This result follows the lines of our Lyapunov exponents, where we found that the normal hierarchy give rise to a value ~ 6 times larger than the inverted hierarchy. The difference can well be due to the quite different geometry that is present in the two different cases, but it might also be due to the type of perturbation used. We found that the unstable Lyapunov vectors for the normal hierarchy (which almost coincide with the normal modes) mainly point in the P_x - and P_y -directions⁴, and the small component in the P_z -direction is due to the mixing. The perturbation used by Mirizzi in introduced in the flux and the x - and y -components of the perturbation would hence also be reduced by the mixing angle. As the mixing angle used was 10^{-2} , it is quite possible that this actually suppresses the perturbation to an effective 10^{-5} . A value which would mean that the growth rate is actually ~ 6 times higher in the normal hierarchy than in the inverted hierarchy just as we found.

Assuming that the essence of this discussion is actually correct, it would mean that a similar perturbation in the oscillation phase would lead to conversion at a significantly lower radius, but in doing this, we neglect the important phenomenon of synchronised oscillations that could well suppress the conversion. Only a detailed calculation will be able to show if that is the case.

The next instability in collective oscillations that has been characterised is present when the assumption of homogeneity is abandoned. Although the initial studies introduced the inhomogeneity in one coordinate while the velocity was defined in a two-dimensional space, they were effectively one-dimensional [131, 158, 166]. The interesting feature of this instability is that it exists at very high neutrino density for very small scales. The implication for simulation is that one would need to follow

⁴Notice that we use the coordinate system from the previous section rather than the one used by Mirizzi in [159].

almost arbitrarily small scales when going to the very centre of the supernova where the neutrino density is highest. The work by Mirizzi discussed above is the latest development on the numerical side, where the one dimensional model has been wrapped into a ring [159], while the task of analysing a two dimensional model and including the matter effect has been done in [6].

The results in the two-dimensional model are very similar to what was found in the 1D case, although a new unstable mode appears. The inclusion of the matter effect turns out to be more interesting, as it suppresses the small scale modes at high densities to the point where the homogeneous mode is practically always the first one to become unstable. This means that the problem of simulating such a system can probably be overcome using a finite resolution, although the details will depend on the density profile of the supernova. Despite these first steps towards a better understanding of the inhomogeneity instability, there is still a lot to learn. The full system is described by a non-linear partial differential equation for the mixing in three momenta for every point in the supernova that one wishes to model, and is impossible to handle. Some of the issues that are known to have an important impact and that has not been included in the inhomogeneity studies until now are neutrinos from the halo [157] and some of the more nasty details of supernova explosions [167], but on top of that, there might well be new and unexplored problems hidden somewhere in the 7 dimensional phase space.

Part III

Future developments

*Det er svært at spå, især om fremtiden.
(Predicting is hard, especially the future.)*

— Storm P.

Although predicting the future is always dangerous, there are several interesting experiments that we can expect to give results in the coming years, and some of the major efforts for the more distant future are also planned already. This said, the progress often presents itself in unexpected ways.

Experiments elucidating the elusive neutrino

Some of the near future experiments that have the tightest connection to this thesis are the efforts for testing the neutrino anomalies. There are several different approaches towards this goal. The most popular method is to install a detector at a nuclear reactor, which is either segmented, can move, has good sensitivity to the spectrum or a combination of several of these. The second option is to take a dedicated neutrino detector, that has an accurate determination of the point where the neutrino is detected and place a strong neutrino source in the near vicinity or even inside the detector. The third approach is to use accelerator neutrinos just as the LSND and MiniBooNE experiments did and observe either appearance or disappearance of the neutrinos.

A few of the nuclear reactor experiments are already running. These include NUCIFER [20] and DANSS [168], but their original design was intended for reactor monitoring and not to search for sterile neutrino oscillations, and therefore they are only sensitive to part of the interesting parameter space. Some of the future experiments are instead dedicated to the search for oscillations into sterile neutrino e.g. Stereo [20] in France, Neutrino4 [169] in Russia and Prospect [170] in the United States. Taking the STEREO experiment as an example, the goal is to determine the neutrino energy spectrum at different distances of the reactor core. To achieve this, a detector with a gadolinium doped liquid scintillator is used. The antineutrinos hit a proton resulting in an inverse beta decay: $\bar{\nu}_e + p \rightarrow e^+ + n$. The positron is then detected using the emitted light from the liquid scintillator, while the neutron is captured by gadolinium which subsequently decays emitting a γ -ray. Typical challenges involved in these experiment are the proximity of the surface imposed by the location of the nuclear reactor; neutrons and γ -rays from the reactor itself as well as the radioactivity of the surroundings. Depending on the time needed for data analysis, NUCIFER could release results soon, while DANSS is still taking data, and the time frame for the remaining experiments is slightly longer.

The dedicated neutrino detectors that have been suggested as promising sites for testing the neutrino anomalies using strong radioactive sources are the Borexino [171] and the KamLAND detectors [172]. The advantage of these experiments is that both the detector and the source are expected to be well understood, and as for the nuclear reactor experiments, it should be possible to see neutrino oscillations both as a function of energy and of position. The major challenge is to first produce the strong source and afterwards transport it from the production plant to the site of the experiment. The transport is complicated by the requirement of fast deployment for the highest possible activity combined with the fact that many countries have placed heavy restrictions on radioactive nuclear material due to safety concerns. The current schedule for SOX places the first experiments with a cerium source in 2016, while details of the time plan for CeLAND are unknown.

Various accelerator experiments are also running and being planned. The ICARUS

experiment [18], which has stopped data taking, has been mentioned earlier in the thesis, but the OPERA experiment which is also using the CNGS beam can also give some limits on sterile neutrino oscillations [173]. Others of the major running experiments are NO ν A, MINOS+ and the T2K experiment. These tend to be multipurpose neutrino experiments in that they both aim to measure neutrino cross sections as well as putting constraints on various aspects of neutrino oscillations, including sterile neutrinos. One of the future efforts is the MicroBooNE experiment which will continue the work that originally started with LSND and was continued with MiniBooNE. It will even use the same beam that was also used for MiniBooNE. All of these experiments use large existing accelerator facilities to produce a beam which travels between 100's of meters and 1000's of kilometers, but there is also an alternative approach. This is to make a small accelerator and place it near an existing detector, a method proposed by the IsoDAR collaboration.

As has probably been demonstrated by now, there are a lot of experiments that aim to resolve the neutrino anomaly problem within the coming years, but apart from this, there are also many other interesting experiments that are progressing forwards.

After the Daya Bay experiment successfully determined the θ_{13} mixing angle by measuring reactor neutrinos, the collaboration behind it has worked on an improved detector and site, which has resulted in the plans for JUNO. The new detector will allow for an improved measurement of the mixing parameters, and ultimately a determination of the neutrino mass hierarchy. Due to the large and sensitive detector, it will furthermore be possible to measure neutrinos from the Sun, from supernova and geo-neutrinos.

Although the JUNO project is ambitious, it will by no means be the largest neutrino detector. That prize goes to ICECUBE at the South Pole, which consist of 1km³ of instrumented ice. The main goal of ICECUBE is to measure neutrinos from outside our solar system, but in doing this, it also measures a lot of atmospheric neutrinos. The large volume of the detector allows it to detect extremely high energy events with energies above 1PeV, but a large distance between the photon detectors limits its efficiency for low energy events. This also means that the prospect for determining the mass hierarchy using ICECUBE itself are not good. In order to contribute towards this goal as well, the collaboration will be installing PINGU, which is a far more densely packed grid of detectors, making it capable of measuring low energy atmospheric neutrinos.

Another approach towards the mass hierarchy is to determine the masses themselves. This has been the goal for the KATRIN experiments for more than a decade now, but as the various components are assembled, results are moving closer. The idea is to use the decay of tritium from a strong source to measure the detailed electron spectrum. Knowing what the maximal available energy to the neutrino should be in the case of a massless neutrino, it is possible to determine the neutrino mass from the shape of the electron spectrum end point. Using similar techniques, the neutrino mass has been constrained to $\leq 2.3\text{eV}$, but the goal of KATRIN is to put a limit of $\leq 0.2\text{eV}$ which is estimated to be the ultimately achievable precision with this type of experiment.

Finally, one of the other fundamental and unknown properties of the neutrino is its particle type. If it is its own antiparticle, it is the first known fundamental Majorana fermion and otherwise, it is a Dirac fermion just as the quarks and the charged leptons. If it is a Majorana particle, it should be possible to observe double beta decay without any emitted neutrinos, an experimental signature which both the Gerda and the Majorana experiments look for using germanium as both sources and detectors. Other experiments

such and EXO and NEMO use other elements that decay through double beta decay, but so far none of the experiments have been successful in observing a reproducible signal.

Given the large number of experiments, it is impossible to predict where the next interesting results are going to come from. Even with the many experiments mentioned here, a lot of interesting projects have been omitted, which fortunately means that the prospects for learning more about neutrinos are bright.

Other particular experiments

Although a great effort is put towards understanding neutrinos, they do not feature prominently at the largest particle experiment of them all, the Large Hadron Collider (LHC). After run-1 with collisions at 7 and 8TeV has produced evidence for a Higgs (like) particle with a mass of 125GeV, the collider is now running again with an energy of 13TeV. This higher energy will of course be used for further characterising the Higgs boson, but there are also a few interesting hints in the data from run-1 that could be new physics (or just the usual statistical fluctuations).

Currently, both Atlas and CMS see an excess in diboson channels at or just below 2TeV. The local significance for Atlas is $\sim 3\sigma$, and at CMS it is even lower, but the fact that both experiments see something around the same energy makes it slightly more interesting than the standard 3σ fluctuation. The problem is that the signal is hard to explain using simple models as the Atlas excess requires a rather large coupling between a hypothetical new particle and the light quarks, but still a decay channel that is dominated by diboson events. On top of this, the cross sections seen by Atlas and CMS are inconsistent although the processes they measure are somewhat different. Considering this, the most probable outcome is that the signal is gone in the run-2 data, and otherwise we will need some rather convoluted theoretical models to explain this.

Both Atlas and CMS have also reported other excesses. In Atlas' case e.g. an excess in a search for squarks and gluinos, and for CMS a signal that could be a neutralino. The common property of all these cases is that the other experiment did not see anything, which makes it quite unlikely that it is an actual signal.

One of the more persistent stories throughout the first run, is the anomalies observed by LHCb and CMS in B-mesons decaying to Kaons. Although there could be quite complicated physics involved, some of the deviations are seen in quite clean channels such as the ratio of $B^+ \rightarrow K^+ + e^+ + e^-$ and $B^+ \rightarrow K^+ + \mu^+ + \mu^-$, where most uncertainties should cancel, and lepton flavour universality gives a solid prediction of the ratio. While we currently do not have well established models that would predict this sort of signals, it is allegedly not too hard to think of new particles and Lagrangian terms that would give rise to just such a signature. Fortunately, there will soon be more data, and as for the other hints, there is a good chance to know for sure within some years.

Despite these few hints for new interesting physics, it is hard to know what to expect from run-2 of LHC, and the common wisdom seems to be that there is a significant chance that nothing beyond the Standard Model will show up. Theoretically there are not any really good reasons to believe that new physics should be just around the corner, and most of the experimental data in particle physics are already well understood.

Another field where the risk of no signals is non-negligible is the search for dark matter. Especially the WIMP-paradigm is under pressure as the interesting cross sections are being probed by larger and larger detectors such as LUX and Xenon-1000 using xenon both as a target nucleus for interactions with the dark matter and as a liquid scintillator. While these experiments are leading the field at higher masses, germanium and silicon detectors such as CDMS are dominating the lower mass range. It is not that there have been no possible detections, but unless very specialised models are considered, all of these are ruled out by more sensitive experiments.

Yet another way to search for dark matter is to look for indirect signals from annihilation and decay. This was where the possible 7keV sterile neutrino signal was seen which we discussed in Part I. As already mentioned, this signal will be tested by new observations [36] which can hopefully enlighten the discussion, but there are also other signals that could come from dark matter. The Fermi-satellite has provided one of these signals [174], where an excess was observed in the galactic centre with a shape consistent with a NFW-profile. This can be interpreted as an annihilating WIMP, but the required cross section is being pushed by the latest results from LUX and Xenon. Another interesting excess was observed already by the PAMELA satellite that found a rising positron fraction for larger energies [175]. These positrons could come from dark matter, but it is also very possible that they are of astrophysical origin. The latest chapter in this story was written when the AMS-02 [176] experiment reported that the positron fraction levels out at even higher energies. As it turns out, this can actually be explained if the positrons come from secondary production [177]. If future measurement were to find a sharp decline, the secondary origin would be rejected, and the most probable source would be pulsars or dark matter [178].

To summarise, there are several interesting excesses and signals both in particle and astroparticle experiments, although only few of them seem to have a reasonable chance of being anything but experimental effects, random fluctuations or complex astrophysics.

Megaparsec milestones

On the cosmological side of things, Planck has just released the latest version of their data [179], and it will take the community a while to digest the results. However, the room for surprises is quite limited as the cosmological parameters inferred by the official analysis gave very few deviations from what was expected for Λ CDM [21]. With these results, the power spectrum of CMB anisotropies has been measured as well as it is possible up to reasonably high multipoles given that we only have one sky to look at. Future improvements will come from even higher multipoles that are observed from telescopes, and possibly also from improved measurements of the CMB spectrum. Most of the major current efforts are however pointing in other directions.

Both the Euclid space mission and the LSST telescope will aim to map structures in the universe using weak gravitational lensing. The idea is that the observed shapes of galaxies are deformed by objects between the observer and the galaxy. Thus it is possible to infer the distribution of matter and properties of space given statistical knowledge of the original shapes of the galaxies. While this technique is already being used, the scale and coverage of Euclid and LSST are unprecedented and the results will

improve constraints on cosmological models significantly [72].

Another promising approach for observing the Universe is to use the 21cm line from neutral hydrogen which comes from a forbidden transition between two hyperfine levels. The unique property of this line is that it can probe the “Dark ages” of the Universe when the Universe was filled with neutral hydrogen after recombination and before the first stars were formed. This era ended as soon as the first stars and other radiating objects started to ionize hydrogen again during reionization. The impact of such observations will of course be a better understanding of the reionization itself, but it will also give an invaluable insight into the early parts of structure formation that cannot be probed in any other way. The challenge for such observations is that the foregrounds are as large as ~ 1000 times the signal we would like to observe, and hence a great deal of work and ingenuity will be required in order to reliably subtract foregrounds and obtain a clean signal.

The next galactic supernova

The final part of this thesis has been focussed on supernova neutrino oscillations, and although there might be the possibility of observing a diffuse background of supernova neutrinos from other galaxies, the only real test of the theory would come if we got a galactic supernova. The estimated rate is around one supernova per 30-50 years, if we look for radioactive Al-26 which is mainly produced by massive stars [180] or look for the supernova rate in other similar galaxies. Looking at the historic observations however, there are far fewer recorded supernovae than this which is probably explained by obscuring galactic dust. For neutrinos this is not an issue, although it would be advantageous to have the supernova light curve as well as the neutrino events for an improved understanding of the explosion. Given the rate, there should occur 1-2 galactic supernovae in a lifetime, and as I just missed SN1987A with a few months (which does arguably not count anyway as it was located in the Large Magellanic Cloud), I am still looking forward to my first one.

The possibility of observing neutrinos from a galactic supernova relies on the existence of large neutrino detectors that are continuously monitored. These will usually be constructed for other purposes, and currently there are a few around the Earth. The most important are the Super-Kamiokande, ICECUBE, Borexino, KamLAND and LVD. The last one is actually a dedicated detector for neutrinos from stellar collapses, while all of the others mainly look for neutrinos from other sources. Here the ICECUBE neutrino detector stands a bit out, as it was not designed for so low energy neutrinos. However, the sheer number of low energy neutrinos would give a significantly increased background that could be observed, although the detector would not be capable of resolving individual events. The highest event rate is expected for Super-Kamiokande due to its large volume, and a supernova at the galactic centre is estimated to give ~ 4000 events. With so large statistics, it is conceivable that even relatively fine features of the neutrino signal can be resolved. In the future even more powerful detectors such as the Hyper-Kamiokande and Juno will join the club, and we can expect an even better signal. The beauty is that we cannot know if they will make it in time. The next galactic supernova might be observed just now, or we might have to wait several decades.

Bibliography

- [1] Steen Hannestad, **Rasmus Sloth Hansen**, and Thomas Tram. “Can active-sterile neutrino oscillations lead to chaotic behavior of the cosmological lepton asymmetry?” *JCAP* 1304 (2013), p. 032. DOI: [10.1088/1475-7516/2013/04/032](https://doi.org/10.1088/1475-7516/2013/04/032). arXiv: [1302.7279](https://arxiv.org/abs/1302.7279) [[astro-ph.CO](#)].
- [2] **Rasmus Sloth Hansen** and Steen Hannestad. “Chaotic flavor evolution in an interacting neutrino gas”. *Phys.Rev.* D90.2 (2014), p. 025009. DOI: [10.1103/PhysRevD.90.025009](https://doi.org/10.1103/PhysRevD.90.025009). arXiv: [1404.3833](https://arxiv.org/abs/1404.3833) [[hep-ph](#)].
- [3] Steen Hannestad, **Rasmus Sloth Hansen**, and Thomas Tram. “How Self-Interactions can Reconcile Sterile Neutrinos with Cosmology”. *Phys.Rev.Lett.* 112.3 (2014), p. 031802. DOI: [10.1103/PhysRevLett.112.031802](https://doi.org/10.1103/PhysRevLett.112.031802). arXiv: [1310.5926](https://arxiv.org/abs/1310.5926) [[astro-ph.CO](#)].
- [4] Maria Archidiacono, Steen Hannestad, **Rasmus Sloth Hansen**, and Thomas Tram. “Cosmology with self-interacting sterile neutrinos and dark matter - A pseudoscalar model”. *Phys.Rev.* D91.6 (2015), p. 065021. DOI: [10.1103/PhysRevD.91.065021](https://doi.org/10.1103/PhysRevD.91.065021). arXiv: [1404.5915](https://arxiv.org/abs/1404.5915) [[astro-ph.CO](#)].
- [5] Steen Hannestad, **Rasmus Sloth Hansen**, Thomas Tram, and Yvonne Y. Y. Wong. “Active-sterile neutrino oscillations in the early Universe with full collision terms”. *JCAP* (2015), forthcoming. arXiv: [1506.05266](https://arxiv.org/abs/1506.05266) [[hep-ph](#)].
- [6] Sovan Chakraborty, Rasmus Sloth Hansen, Ignacio Izaguirre, and Georg Raffelt. “Self-induced flavor conversion of supernova neutrinos on small scales” (2015). arXiv: [1507.07569](https://arxiv.org/abs/1507.07569) [[hep-ph](#)].
- [7] O. Yu. Smirnov et al. “Measurement of neutrino flux from the primary proton-proton fusion process in the Sun with Borexino detector” (2015). arXiv: [1507.02432](https://arxiv.org/abs/1507.02432) [[hep-ex](#)].
- [8] B. Pontecorvo. “Mesonium and anti-mesonium”. *Sov. Phys. JETP* 6 (1957). [*Zh. Eksp. Teor. Fiz.*33,549(1957)], p. 429.
- [9] L. Wolfenstein. “Neutrino Oscillations in Matter”. *Phys.Rev.* D17 (1978), pp. 2369–2374. DOI: [10.1103/PhysRevD.17.2369](https://doi.org/10.1103/PhysRevD.17.2369).
- [10] S. P. Mikheev and A. Yu. Smirnov. “Resonant amplification of neutrino oscillations in matter and solar neutrino spectroscopy”. *Nuovo Cim.* C9 (1986), pp. 17–26. DOI: [10.1007/BF02508049](https://doi.org/10.1007/BF02508049).
- [11] S. Schael et al. “Precision electroweak measurements on the Z resonance”. *Phys.Rept.* 427 (2006), pp. 257–454. DOI: [10.1016/j.physrep.2005.12.006](https://doi.org/10.1016/j.physrep.2005.12.006). arXiv: [hep-ex/0509008](https://arxiv.org/abs/hep-ex/0509008) [[hep-ex](#)].

Bibliography

- [12] A. Aguilar-Arevalo et al. “Evidence for neutrino oscillations from the observation of anti- ν /e appearance in a anti- ν / μ beam”. *Phys. Rev.* D64 (2001), p. 112007. DOI: [10.1103/PhysRevD.64.112007](https://doi.org/10.1103/PhysRevD.64.112007). arXiv: [hep-ex/0104049](https://arxiv.org/abs/hep-ex/0104049).
- [13] A.A. Aguilar-Arevalo et al. “Improved Search for $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ Oscillations in the MiniBooNE Experiment”. *Phys.Rev.Lett.* 110 (2013), p. 161801. DOI: [10.1103/PhysRevLett.110.161801](https://doi.org/10.1103/PhysRevLett.110.161801). arXiv: [1207.4809](https://arxiv.org/abs/1207.4809) [[hep-ex](#)].
- [14] Th.A. Mueller et al. “Improved Predictions of Reactor Antineutrino Spectra”. *Phys.Rev.* C83 (2011), p. 054615. DOI: [10.1103/PhysRevC.83.054615](https://doi.org/10.1103/PhysRevC.83.054615). arXiv: [1101.2663](https://arxiv.org/abs/1101.2663) [[hep-ex](#)].
- [15] J.N. Abdurashitov et al. “Measurement of the solar neutrino capture rate with gallium metal. III: Results for the 2002–2007 data-taking period”. *Phys.Rev.* C80 (2009), p. 015807. DOI: [10.1103/PhysRevC.80.015807](https://doi.org/10.1103/PhysRevC.80.015807). arXiv: [0901.2200](https://arxiv.org/abs/0901.2200) [[nucl-ex](#)].
- [16] F. Kaether et al. “Reanalysis of the GALLEX solar neutrino flux and source experiments”. *Phys.Lett.* B685 (2010), pp. 47–54. DOI: [10.1016/j.physletb.2010.01.030](https://doi.org/10.1016/j.physletb.2010.01.030). arXiv: [1001.2731](https://arxiv.org/abs/1001.2731) [[hep-ex](#)].
- [17] B. Armbruster et al. “Upper limits for neutrino oscillations muon-anti-neutrino \rightarrow electron-anti-neutrino from muon decay at rest”. *Phys. Rev.* D65 (2002), p. 112001. DOI: [10.1103/PhysRevD.65.112001](https://doi.org/10.1103/PhysRevD.65.112001). arXiv: [hep-ex/0203021](https://arxiv.org/abs/hep-ex/0203021) [[hep-ex](#)].
- [18] M. Antonello et al. “Some conclusive considerations on the comparison of the ICARUS ν_μ to ν_e oscillation search with the MiniBooNE low-energy event excess” (2015). arXiv: [1502.04833](https://arxiv.org/abs/1502.04833) [[hep-ph](#)].
- [19] J. Gaffiot. “Reactor experiments to test sterile neutrinos”. *Nucl.Phys.Proc.Suppl.* 237-238 (2013), pp. 326–328. DOI: [10.1016/j.nuclphysbps.2013.04.118](https://doi.org/10.1016/j.nuclphysbps.2013.04.118).
- [20] K.N. Abazajian et al. “Light Sterile Neutrinos: A White Paper” (2012). arXiv: [1204.5379](https://arxiv.org/abs/1204.5379) [[hep-ph](#)].
- [21] P.A.R. Ade et al. “Planck 2015 results. XIII. Cosmological parameters” (2015). arXiv: [1502.01589](https://arxiv.org/abs/1502.01589) [[astro-ph.CO](#)].
- [22] Anatoly A. Klypin, Andrey V. Kravtsov, Octavio Valenzuela, and Francisco Prada. “Where are the missing Galactic satellites?” *Astrophys.J.* 522 (1999), pp. 82–92. DOI: [10.1086/307643](https://doi.org/10.1086/307643). arXiv: [astro-ph/9901240](https://arxiv.org/abs/astro-ph/9901240) [[astro-ph](#)].
- [23] Michael Boylan-Kolchin, James S. Bullock, and Manoj Kaplinghat. “Too big to fail? The puzzling darkness of massive Milky Way subhaloes”. *Mon.Not.Roy.Astron.Soc.* 415 (2011), p. L40. arXiv: [1103.0007](https://arxiv.org/abs/1103.0007) [[astro-ph.CO](#)].
- [24] W.J.G. de Blok. “The Core-Cusp Problem”. *Adv.Astron.* 2010 (2010), p. 789293. DOI: [10.1155/2010/789293](https://doi.org/10.1155/2010/789293). arXiv: [0910.3538](https://arxiv.org/abs/0910.3538) [[astro-ph.CO](#)].
- [25] Matteo Viel, George D. Becker, James S. Bolton, and Martin G. Haehnelt. “Warm dark matter as a solution to the small scale crisis: New constraints from high redshift Lyman- α forest data”. *Phys.Rev.* D88 (2013), p. 043502. DOI: [10.1103/PhysRevD.88.043502](https://doi.org/10.1103/PhysRevD.88.043502). arXiv: [1306.2314](https://arxiv.org/abs/1306.2314) [[astro-ph.CO](#)].

- [26] Esra Bulbul et al. “Detection of An Unidentified Emission Line in the Stacked X-ray spectrum of Galaxy Clusters”. *Astrophys.J.* 789 (2014), p. 13. DOI: [10.1088/0004-637X/789/1/13](https://doi.org/10.1088/0004-637X/789/1/13). arXiv: [1402.2301](https://arxiv.org/abs/1402.2301) [[astro-ph.CO](#)].
- [27] Alexey Boyarsky, Oleg Ruchayskiy, Dmytro Iakubovskiy, and Jeroen Franse. “Unidentified Line in X-Ray Spectra of the Andromeda Galaxy and Perseus Galaxy Cluster”. *Phys.Rev.Lett.* 113 (2014), p. 251301. DOI: [10.1103/PhysRevLett.113.251301](https://doi.org/10.1103/PhysRevLett.113.251301). arXiv: [1402.4119](https://arxiv.org/abs/1402.4119) [[astro-ph.CO](#)].
- [28] Alexey Boyarsky, Jeroen Franse, Dmytro Iakubovskiy, and Oleg Ruchayskiy. “Checking the dark matter origin of 3.53 keV line with the Milky Way center” (2014). arXiv: [1408.2503](https://arxiv.org/abs/1408.2503) [[astro-ph.CO](#)].
- [29] D. Malyshev, A. Neronov, and D. Eckert. “Constraints on 3.55 keV line emission from stacked observations of dwarf spheroidal galaxies”. *Phys.Rev.* D90.10 (2014), p. 103506. DOI: [10.1103/PhysRevD.90.103506](https://doi.org/10.1103/PhysRevD.90.103506). arXiv: [1408.3531](https://arxiv.org/abs/1408.3531) [[astro-ph.HE](#)].
- [30] O. Urban et al. “A Suzaku Search for Dark Matter Emission Lines in the X-ray Brightest Galaxy Clusters” (2014). arXiv: [1411.0050](https://arxiv.org/abs/1411.0050) [[astro-ph.CO](#)].
- [31] Eric Carlson, Tesla Jeltema, and Stefano Profumo. “Where do the 3.5 keV photons come from? A morphological study of the Galactic Center and of Perseus”. *JCAP* 1502.02 (2015), p. 009. DOI: [10.1088/1475-7516/2015/02/009](https://doi.org/10.1088/1475-7516/2015/02/009). arXiv: [1411.1758](https://arxiv.org/abs/1411.1758) [[astro-ph.HE](#)].
- [32] Takayuki Tamura et al. “An X-ray Spectroscopic Search for Dark Matter in the Perseus Cluster with Suzaku”. *Publ.Astron.Soc.Jap.* 67.2 (2015), p. 23. DOI: [10.1093/pasj/psu156](https://doi.org/10.1093/pasj/psu156). arXiv: [1412.1869](https://arxiv.org/abs/1412.1869) [[astro-ph.HE](#)].
- [33] Norio Sekiya, Noriko Y. Yamasaki, and Kazuhisa Mitsuda. “A Search for a keV Signature of Radiatively Decaying Dark Matter with Suzaku XIS Observations of the X-ray Diffuse Background” (2015). arXiv: [1504.02826](https://arxiv.org/abs/1504.02826) [[astro-ph.HE](#)].
- [34] Signe Riemer-Sorensen. “Questioning a 3.5 keV dark matter emission line” (2014). arXiv: [1405.7943](https://arxiv.org/abs/1405.7943) [[astro-ph.CO](#)].
- [35] Tesla E. Jeltema and Stefano Profumo. “Discovery of a 3.5 keV line in the Galactic Center and a Critical Look at the Origin of the Line Across Astronomical Targets”. *Mon.Not.Roy.Astron.Soc.* 450 (2015), pp. 2143–2152. DOI: [10.1093/mnras/stv768](https://doi.org/10.1093/mnras/stv768). arXiv: [1408.1699](https://arxiv.org/abs/1408.1699) [[astro-ph.HE](#)].
- [36] Mark R. Lovell et al. “Decaying dark matter: the case for a deep X-ray observation of Draco” (2014). arXiv: [1411.0311](https://arxiv.org/abs/1411.0311) [[astro-ph.CO](#)].
- [37] M. Fukugita and T. Yanagida. *Physics of Neutrinos: And Applications to Astrophysics*. Physics and astronomy online library. Springer, 2003. ISBN: 9783540438007.
- [38] K.A. Olive et al. “Review of Particle Physics”. *Chin.Phys.* C38 (2014), p. 090001. DOI: [10.1088/1674-1137/38/9/090001](https://doi.org/10.1088/1674-1137/38/9/090001).
- [39] Alexey Vlasenko, George M. Fuller, and Vincenzo Cirigliano. “Neutrino Quantum Kinetics”. *Phys.Rev.* D89.10 (2014), p. 105004. DOI: [10.1103/PhysRevD.89.105004](https://doi.org/10.1103/PhysRevD.89.105004). arXiv: [1309.2628](https://arxiv.org/abs/1309.2628) [[hep-ph](#)].

Bibliography

- [40] Alexey Vlasenko, George M. Fuller, and Vincenzo Cirigliano. “Prospects for Neutrino-Antineutrino Transformation in Astrophysical Environments” (2014). arXiv: [1406.6724 \[astro-ph.HE\]](#).
- [41] Cristina Volpe, Daavid Väänänen, and Catalina Espinoza. “Extended evolution equations for neutrino propagation in astrophysical and cosmological environments” (2013). arXiv: [1302.2374 \[hep-ph\]](#).
- [42] Julien Serreau and Cristina Volpe. “Neutrino-antineutrino correlations in dense anisotropic media”. *Phys.Rev.* D90.12 (2014), p. 125040. DOI: [10.1103/PhysRevD.90.125040](#). arXiv: [1409.3591 \[hep-ph\]](#).
- [43] A. Kartavtsev, G. Raffelt, and H. Vogel. “Neutrino propagation in media: Flavor-, helicity-, and pair correlations”. *Phys.Rev.* D91.12 (2015), p. 125020. DOI: [10.1103/PhysRevD.91.125020](#). arXiv: [1504.03230 \[hep-ph\]](#).
- [44] Cristina Volpe. “Neutrino Quantum Kinetic Equations” (2015). arXiv: [1506.06222 \[astro-ph.SR\]](#).
- [45] Bruce H. J. McKellar and Mark J. Thomson. “Oscillating doublet neutrinos in the early universe”. *Phys. Rev.* D49 (1994), pp. 2710–2728. DOI: [10.1103/PhysRevD.49.2710](#).
- [46] Kimmo Kainulainen and Antti Sorri. “Oscillation induced neutrino asymmetry growth in the early universe”. *JHEP* 02 (2002), p. 020. arXiv: [hep-ph/0112158](#).
- [47] Huaiyu Duan, George M. Fuller, J Carlson, and Yong-Zhong Qian. “Simulation of Coherent Non-Linear Neutrino Flavor Transformation in the Supernova Environment: Correlated Neutrino Trajectories”. *Phys.Rev.* D74 (2006), p. 105014. DOI: [10.1103/PhysRevD.74.105014](#). arXiv: [astro-ph/0606616 \[astro-ph\]](#).
- [48] Steen Hannestad, Irene Tamborra, and Thomas Tram. “Thermalisation of light sterile neutrinos in the early universe”. *JCAP* 1207 (2012), p. 025. DOI: [10.1088/1475-7516/2012/07/025](#). arXiv: [1204.5861 \[astro-ph.CO\]](#).
- [49] Steen Hannestad, Georg G. Raffelt, Gunter Sigl, and Yvonne Y.Y. Wong. “Self-induced conversion in dense neutrino gases: Pendulum in flavour space”. *Phys.Rev.* D74 (2006), p. 105010. DOI: [10.1103/PhysRevD.74.105010](#), [10.1103/PhysRevD.76.029901](#). arXiv: [astro-ph/0608695 \[astro-ph\]](#).
- [50] Huaiyu Duan, George M. Fuller, and Yong-Zhong Qian. “Collective neutrino flavor transformation in supernovae”. *Phys.Rev.* D74 (2006), p. 123004. DOI: [10.1103/PhysRevD.74.123004](#). arXiv: [astro-ph/0511275 \[astro-ph\]](#).
- [51] Georg Raffelt and David de Sousa Seixas. “Neutrino flavor pendulum in both mass hierarchies”. *Phys.Rev.* D88 (2013), p. 045031. DOI: [10.1103/PhysRevD.88.045031](#). arXiv: [1307.7625 \[hep-ph\]](#).
- [52] Georg Raffelt, Srdjan Sarikas, and David de Sousa Seixas. “Axial symmetry breaking in self-induced flavor conversion of supernova neutrino fluxes”. *Phys.Rev.Lett.* 111 (2013), p. 091101. DOI: [10.1103/PhysRevLett.111.091101](#). arXiv: [1305.7140 \[hep-ph\]](#).
- [53] Basudeb Dasgupta, Amol Dighe, Alessandro Mirizzi, and Georg G. Raffelt. “Collective neutrino oscillations in non-spherical geometry”. *Phys.Rev.* D78 (2008), p. 033014. DOI: [10.1103/PhysRevD.78.033014](#). arXiv: [0805.3300 \[hep-ph\]](#).

- [54] Alessandro Mirizzi, Ninetta Saviano, Gennaro Miele, and Pasquale Dario Serpico. “Light sterile neutrino production in the early universe with dynamical neutrino asymmetries”. *Phys.Rev.* D86 (2012), p. 053009. DOI: [10.1103/PhysRevD.86.053009](https://doi.org/10.1103/PhysRevD.86.053009). arXiv: [1206.1046](https://arxiv.org/abs/1206.1046) [[hep-ph](#)].
- [55] Ninetta Saviano et al. “Multi-momentum and multi-flavour active-sterile neutrino oscillations in the early universe: role of neutrino asymmetries and effects on nucleosynthesis”. *Phys.Rev.* D87 (2013), p. 073006. DOI: [10.1103/PhysRevD.87.073006](https://doi.org/10.1103/PhysRevD.87.073006). arXiv: [1302.1200](https://arxiv.org/abs/1302.1200) [[astro-ph.CO](#)].
- [56] G. Sigl and G. G. Raffelt. “General kinetic description of relativistic mixed neutrinos”. *Nucl. Phys.* B406 (1993), pp. 423–451. DOI: [10.1016/0550-3213\(93\)90175-0](https://doi.org/10.1016/0550-3213(93)90175-0).
- [57] Dirk Notzold and Georg G. Raffelt. “Neutrino Dispersion at Finite Temperature and Density”. *Nucl. Phys.* B307 (1988), p. 924. DOI: [10.1016/0550-3213\(88\)90113-7](https://doi.org/10.1016/0550-3213(88)90113-7).
- [58] K. Enqvist, K. Kainulainen, and Mark J. Thomson. “Stringent cosmological bounds on inert neutrino mixing”. *Nucl.Phys.* B373 (1992), pp. 498–528. DOI: [10.1016/0550-3213\(92\)90442-E](https://doi.org/10.1016/0550-3213(92)90442-E).
- [59] A.I. Akhiezer and S.V. Peletminskii. “Chapter 5 - Kinetic Equations for Quantum Systems”. *Methods of Statistical Physics*. Ed. by A.I. Akhiezer S.V. Peletminskii. Vol. 104. International Series in Natural Philosophy. Pergamon, 1981, pp. 246–371. ISBN: 9780080250403. DOI: [10.1016/B978-0-08-025040-3.50010-9](https://doi.org/10.1016/B978-0-08-025040-3.50010-9).
- [60] Nicole F. Bell, Raymond R. Volkas, and Yvonne Y. Y. Wong. “Relic neutrino asymmetry evolution from first principles”. *Phys. Rev.* D59 (1999), p. 113001. DOI: [10.1103/PhysRevD.59.113001](https://doi.org/10.1103/PhysRevD.59.113001). arXiv: [hep-ph/9809363](https://arxiv.org/abs/hep-ph/9809363).
- [61] S. Dodelson. *Modern Cosmology*. Academic Press. Academic Press, 2003. ISBN: 9780122191411.
- [62] Lawrence F. Shampine and Mark W. Reichelt. “The MATLAB ODE Suite”. *SIAM J. Sci. Comput.* 18.1 (1997), pp. 1–22. ISSN: 1064-8275. DOI: [10.1137/S1064827594276424](https://doi.org/10.1137/S1064827594276424). URL: <http://dx.doi.org/10.1137/S1064827594276424>.
- [63] Thomas Tram. “Dark Relics in Cosmology: from production to prediction”. PhD thesis. Aarhus University, 2012.
- [64] Steen Hannestad and Jes Madsen. “Neutrino decoupling in the early universe”. *Phys.Rev.* D52 (1995), pp. 1764–1769. DOI: [10.1103/PhysRevD.52.1764](https://doi.org/10.1103/PhysRevD.52.1764). arXiv: [astro-ph/9506015](https://arxiv.org/abs/astro-ph/9506015) [[astro-ph](#)].
- [65] A.D. Dolgov, S.H. Hansen, and D.V. Semikoz. “Nonequilibrium corrections to the spectra of massless neutrinos in the early universe”. *Nucl.Phys.* B503 (1997), pp. 426–444. DOI: [10.1016/S0550-3213\(97\)00479-3](https://doi.org/10.1016/S0550-3213(97)00479-3). arXiv: [hep-ph/9703315](https://arxiv.org/abs/hep-ph/9703315) [[hep-ph](#)].
- [66] F. Hahn-Woernle, M. Plumacher, and Y.Y.Y. Wong. “Full Boltzmann equations for leptogenesis including scattering”. *JCAP* 0908 (2009), p. 028. DOI: [10.1088/1475-7516/2009/08/028](https://doi.org/10.1088/1475-7516/2009/08/028). arXiv: [0907.0205](https://arxiv.org/abs/0907.0205) [[hep-ph](#)].

Bibliography

- [67] James M. Cline. “Constraints on almost Dirac neutrinos from neutrino - anti-neutrino oscillations”. *Phys.Rev.Lett.* 68 (1992), pp. 3137–3140. DOI: [10.1103/PhysRevLett.68.3137](https://doi.org/10.1103/PhysRevLett.68.3137).
- [68] Yi-Zen Chu and Marco Cirelli. “Sterile neutrinos, lepton asymmetries, primordial elements: How much of each?” *Phys.Rev.* D74 (2006), p. 085015. DOI: [10.1103/PhysRevD.74.085015](https://doi.org/10.1103/PhysRevD.74.085015). arXiv: [astro-ph/0608206](https://arxiv.org/abs/astro-ph/0608206) [[astro-ph](#)].
- [69] Kimmo Kainulainen. “Light Singlet Neutrinos And The Primordial Nucleosynthesis”. *Phys.Lett.* B244 (1990), pp. 191–195. DOI: [10.1016/0370-2693\(90\)90054-A](https://doi.org/10.1016/0370-2693(90)90054-A).
- [70] Alessandro Mirizzi, Gianpiero Mangano, Ofelia Pisanti, and Ninetta Saviano. “Collisional production of sterile neutrinos via secret interactions and cosmological implications”. *Phys.Rev.* D91.2 (2015), p. 025019. DOI: [10.1103/PhysRevD.91.025019](https://doi.org/10.1103/PhysRevD.91.025019). arXiv: [1410.1385](https://arxiv.org/abs/1410.1385) [[hep-ph](#)].
- [71] Xiaoyong Chu, Basudeb Dasgupta, and Joachim Kopp. “Sterile Neutrinos with Secret Interactions - Lasting Friendship with Cosmology” (2015). arXiv: [1505.02795](https://arxiv.org/abs/1505.02795) [[hep-ph](#)].
- [72] Tobias Basse et al. “Dark energy properties from large future galaxy surveys”. *JCAP* 1405 (2014), p. 021. DOI: [10.1088/1475-7516/2014/05/021](https://doi.org/10.1088/1475-7516/2014/05/021). arXiv: [1304.2321](https://arxiv.org/abs/1304.2321) [[astro-ph.CO](#)].
- [73] Jochen Flaig. “Neutrino-Oszillationen in thermischer Umgebung”. *Diploma Thesis, Max-Planck-Institute for Physics and Astrophysics* (1989).
- [74] Gianpiero Mangano et al. “Relic neutrino decoupling including flavour oscillations”. *Nucl. Phys.* B729 (2005), pp. 221–234. DOI: [10.1016/j.nuclphysb.2005.09.041](https://doi.org/10.1016/j.nuclphysb.2005.09.041). arXiv: [hep-ph/0506164](https://arxiv.org/abs/hep-ph/0506164).
- [75] M. Yu. Khlopov and S.T. Petcov. “Possible cosmological effect of CP violation in neutrino oscillations”. *Phys.Lett.* B99 (1981), p. 117. DOI: [10.1016/0370-2693\(81\)90963-1](https://doi.org/10.1016/0370-2693(81)90963-1).
- [76] Martin J. Savage, Robert A. Malaney, and George M. Fuller. “Neutrino Oscillations and the Leptonic Charge of the Universe”. *Astrophys.J.* 368 (1991), pp. 1–11. DOI: [10.1086/169665](https://doi.org/10.1086/169665).
- [77] K. Enqvist, K. Kainulainen, and J. Maalampi. “Neutrino Asymmetry and Oscillations in the Early Universe”. *Phys.Lett.* B244 (1990), pp. 186–190. DOI: [10.1016/0370-2693\(90\)90053-9](https://doi.org/10.1016/0370-2693(90)90053-9).
- [78] Robert Foot and R.R. Volkas. “Reconciling sterile neutrinos with big bang nucleosynthesis”. *Phys.Rev.Lett.* 75 (1995), p. 4350. DOI: [10.1103/PhysRevLett.75.4350](https://doi.org/10.1103/PhysRevLett.75.4350). arXiv: [hep-ph/9508275](https://arxiv.org/abs/hep-ph/9508275) [[hep-ph](#)].
- [79] D.P. Kirilova and M.V. Chizhov. “Neutrino degeneracy effect on neutrino oscillations and primordial helium yield”. *Nucl.Phys.* B534 (1998), pp. 447–463. DOI: [10.1016/S0550-3213\(98\)00641-5](https://doi.org/10.1016/S0550-3213(98)00641-5). arXiv: [hep-ph/9806441](https://arxiv.org/abs/hep-ph/9806441) [[hep-ph](#)].
- [80] A.D. Dolgov and F.L. Villante. “BBN bounds on active sterile neutrino mixing”. *Nucl.Phys.* B679 (2004), pp. 261–298. DOI: [10.1016/j.nuclphysb.2003.11.031](https://doi.org/10.1016/j.nuclphysb.2003.11.031). arXiv: [hep-ph/0308083](https://arxiv.org/abs/hep-ph/0308083) [[hep-ph](#)].

- [81] Kevork Abazajian, Nicole F. Bell, George M. Fuller, and Yvonne Y. Y. Wong. “Cosmological lepton asymmetry, primordial nucleosynthesis, and sterile neutrinos”. *Phys. Rev. D* 72 (2005), p. 063004. DOI: [10.1103/PhysRevD.72.063004](https://doi.org/10.1103/PhysRevD.72.063004). arXiv: [astro-ph/0410175](https://arxiv.org/abs/astro-ph/0410175).
- [82] Chad T. Kishimoto, George M. Fuller, and Christel J. Smith. “Coherent Active-Sterile Neutrino Flavor Transformation in the Early Universe”. *Phys.Rev.Lett.* 97 (2006), p. 141301. DOI: [10.1103/PhysRevLett.97.141301](https://doi.org/10.1103/PhysRevLett.97.141301). arXiv: [astro-ph/0607403](https://arxiv.org/abs/astro-ph/0607403) [[astro-ph](https://arxiv.org/abs/astro-ph)].
- [83] Carlo Giunti and Marco Laveder. “3+1 and 3+2 Sterile Neutrino Fits”. *Phys.Rev.* D84 (2011), p. 073008. DOI: [10.1103/PhysRevD.84.073008](https://doi.org/10.1103/PhysRevD.84.073008). arXiv: [1107.1452](https://arxiv.org/abs/1107.1452) [[hep-ph](https://arxiv.org/abs/hep-ph)].
- [84] Scott Dodelson and Lawrence M. Widrow. “Sterile-neutrinos as dark matter”. *Phys.Rev.Lett.* 72 (1994), pp. 17–20. DOI: [10.1103/PhysRevLett.72.17](https://doi.org/10.1103/PhysRevLett.72.17). arXiv: [hep-ph/9303287](https://arxiv.org/abs/hep-ph/9303287) [[hep-ph](https://arxiv.org/abs/hep-ph)].
- [85] Xiang-Dong Shi and George M. Fuller. “A new dark matter candidate: Non-thermal sterile neutrinos”. *Phys. Rev. Lett.* 82 (1999), pp. 2832–2835. DOI: [10.1103/PhysRevLett.82.2832](https://doi.org/10.1103/PhysRevLett.82.2832). arXiv: [astro-ph/9810076](https://arxiv.org/abs/astro-ph/9810076).
- [86] Mikhail S. Bilenky and Arcadi Santamaria. “‘Secret’ neutrino interactions” (1999), pp. 50–61. arXiv: [hep-ph/9908272](https://arxiv.org/abs/hep-ph/9908272) [[hep-ph](https://arxiv.org/abs/hep-ph)].
- [87] Yasaman Farzan. “Bounds on the coupling of the Majoron to light neutrinos from supernova cooling”. *Phys.Rev.* D67 (2003), p. 073015. DOI: [10.1103/PhysRevD.67.073015](https://doi.org/10.1103/PhysRevD.67.073015). arXiv: [hep-ph/0211375](https://arxiv.org/abs/hep-ph/0211375) [[hep-ph](https://arxiv.org/abs/hep-ph)].
- [88] Maria Archidiacono and Steen Hannestad. “Updated constraints on non-standard neutrino interactions from Planck”. *JCAP* 1407 (2014), p. 046. DOI: [10.1088/1475-7516/2014/07/046](https://doi.org/10.1088/1475-7516/2014/07/046). arXiv: [1311.3873](https://arxiv.org/abs/1311.3873) [[astro-ph.CO](https://arxiv.org/abs/astro-ph.CO)].
- [89] K.S. Babu and I.Z. Rothstein. “Relaxing nucleosynthesis bounds on sterile-neutrinos”. *Phys.Lett.* B275 (1992), pp. 112–118. DOI: [10.1016/0370-2693\(92\)90860-7](https://doi.org/10.1016/0370-2693(92)90860-7).
- [90] Kari Enqvist, Kimmo Kainulainen, and Mark J. Thomson. “Cosmological bounds on Dirac-Majorana neutrinos”. *Phys.Lett.* B280 (1992), pp. 245–250. DOI: [10.1016/0370-2693\(92\)90062-9](https://doi.org/10.1016/0370-2693(92)90062-9).
- [91] Francis-Yan Cyr-Racine and Kris Sigurdson. “Limits on Neutrino-Neutrino Scattering in the Early Universe”. *Phys.Rev.* D90.12 (2014), p. 123533. DOI: [10.1103/PhysRevD.90.123533](https://doi.org/10.1103/PhysRevD.90.123533). arXiv: [1306.1536](https://arxiv.org/abs/1306.1536) [[astro-ph.CO](https://arxiv.org/abs/astro-ph.CO)].
- [92] Riccardo Barbieri and A. Dolgov. “Neutrino oscillations in the early universe”. *Nucl.Phys.* B349 (1991), pp. 743–753. DOI: [10.1016/0550-3213\(91\)90396-F](https://doi.org/10.1016/0550-3213(91)90396-F).
- [93] K. Enqvist, K. Kainulainen, and J. Maalampi. “Refraction and oscillations of neutrinos in the early universe”. *Nucl. Phys.* B349 (1991), pp. 754–790. DOI: [10.1016/0550-3213\(91\)90397-G](https://doi.org/10.1016/0550-3213(91)90397-G).
- [94] Leo Stodolsky. “On the Treatment of Neutrino Oscillations in a Thermal Environment”. *Phys. Rev.* D36 (1987), p. 2273. DOI: [10.1103/PhysRevD.36.2273](https://doi.org/10.1103/PhysRevD.36.2273).

Bibliography

- [95] Steen Hannestad, H. Thomas Janka, Georg G. Raffelt, and Gunter Sigl. “Electron, mu, and tau number conservation in a supernova core”. *Phys.Rev.* D62 (2000), p. 093021. DOI: [10.1103/PhysRevD.62.093021](https://doi.org/10.1103/PhysRevD.62.093021). arXiv: [astro-ph/9912242](https://arxiv.org/abs/astro-ph/9912242) [[astro-ph](#)].
- [96] Basudeb Dasgupta and Joachim Kopp. “Cosmologically Safe eV-Scale Sterile Neutrinos and Improved Dark Matter Structure”. *Phys.Rev.Lett.* 112.3 (2014), p. 031803. DOI: [10.1103/PhysRevLett.112.031803](https://doi.org/10.1103/PhysRevLett.112.031803). arXiv: [1310.6337](https://arxiv.org/abs/1310.6337) [[hep-ph](#)].
- [97] Ninetta Saviano, Ofelia Pisanti, Gianpiero Mangano, and Alessandro Mirizzi. “Unveiling secret interactions among sterile neutrinos with big-bang nucleosynthesis”. *Phys.Rev.* D90.11 (2014), p. 113009. DOI: [10.1103/PhysRevD.90.113009](https://doi.org/10.1103/PhysRevD.90.113009). arXiv: [1409.1680](https://arxiv.org/abs/1409.1680) [[astro-ph.CO](#)].
- [98] Torsten Bringmann, Jasper Hasenkamp, and Jörn Kersten. “Tight bonds between sterile neutrinos and dark matter”. *JCAP* 1407 (2014), p. 042. DOI: [10.1088/1475-7516/2014/07/042](https://doi.org/10.1088/1475-7516/2014/07/042). arXiv: [1312.4947](https://arxiv.org/abs/1312.4947) [[hep-ph](#)].
- [99] P. Ko and Yong Tang. “ $\nu\Lambda$ MDM: A Model for Sterile Neutrino and Dark Matter Reconciles Cosmological and Neutrino Oscillation Data after BICEP2”. *Phys.Lett.* B739 (2014), pp. 62–67. DOI: [10.1016/j.physletb.2014.10.035](https://doi.org/10.1016/j.physletb.2014.10.035). arXiv: [1404.0236](https://arxiv.org/abs/1404.0236) [[hep-ph](#)].
- [100] Yong Tang. “More Is Different: Reconciling eV Sterile Neutrinos with Cosmological Mass Bounds” (2014). arXiv: [1501.00059](https://arxiv.org/abs/1501.00059) [[hep-ph](#)].
- [101] Y. Chikashige, Rabindra N. Mohapatra, and R.D. Peccei. “Are There Real Goldstone Bosons Associated with Broken Lepton Number?” *Phys.Lett.* B98 (1981), p. 265. DOI: [10.1016/0370-2693\(81\)90011-3](https://doi.org/10.1016/0370-2693(81)90011-3).
- [102] M. Kachelriess, R. Tomas, and J.W.F. Valle. “Supernova bounds on Majoron emitting decays of light neutrinos”. *Phys.Rev.* D62 (2000), p. 023004. DOI: [10.1103/PhysRevD.62.023004](https://doi.org/10.1103/PhysRevD.62.023004). arXiv: [hep-ph/0001039](https://arxiv.org/abs/hep-ph/0001039) [[hep-ph](#)].
- [103] T. Bernatowicz et al. “Neutrino mass limits for a precise determination of beta-beta decay rates of Te-128 and Te-130”. *Phys.Rev.Lett.* 69 (1992), pp. 2341–2344. DOI: [10.1103/PhysRevLett.69.2341](https://doi.org/10.1103/PhysRevLett.69.2341).
- [104] Joachim Kopp, Pedro A. N. Machado, Michele Maltoni, and Thomas Schwetz. “Sterile Neutrino Oscillations: The Global Picture”. *JHEP* 1305 (2013), p. 050. DOI: [10.1007/JHEP05\(2013\)050](https://doi.org/10.1007/JHEP05(2013)050). arXiv: [1303.3011](https://arxiv.org/abs/1303.3011) [[hep-ph](#)].
- [105] C. Giunti, M. Laveder, Y.F. Li, and H.W. Long. “Short-baseline electron neutrino oscillation length after troitsk”. *Phys.Rev.* D87.1 (2013), p. 013004. DOI: [10.1103/PhysRevD.87.013004](https://doi.org/10.1103/PhysRevD.87.013004). arXiv: [1212.3805](https://arxiv.org/abs/1212.3805) [[hep-ph](#)].
- [106] A.D. Dolgov, S. Pastor, J.C. Romao, and J.W.F. Valle. “Primordial nucleosynthesis, Majorons and heavy tau neutrinos”. *Nucl.Phys.* B496 (1997), pp. 24–40. DOI: [10.1016/S0550-3213\(97\)00213-7](https://doi.org/10.1016/S0550-3213(97)00213-7). arXiv: [hep-ph/9610507](https://arxiv.org/abs/hep-ph/9610507) [[hep-ph](#)].
- [107] Steen Hannestad. “Structure formation with strongly interacting neutrinos - Implications for the cosmological neutrino mass bound”. *JCAP* 0502 (2005), p. 011. DOI: [10.1088/1475-7516/2005/02/011](https://doi.org/10.1088/1475-7516/2005/02/011). arXiv: [astro-ph/0411475](https://arxiv.org/abs/astro-ph/0411475) [[astro-ph](#)].

- [108] Georg Raffelt and Joseph Silk. “Light neutrinos as cold dark matter”. *Phys.Lett.* B192 (1987), p. 65. DOI: [10.1016/0370-2693\(87\)91143-9](https://doi.org/10.1016/0370-2693(87)91143-9).
- [109] Fernando Atrio-Barandela and Sacha Davidson. “Interacting hot dark matter”. *Phys.Rev.* D55 (1997), pp. 5886–5894. DOI: [10.1103/PhysRevD.55.5886](https://doi.org/10.1103/PhysRevD.55.5886). arXiv: [astro-ph/9702236](https://arxiv.org/abs/astro-ph/9702236) [[astro-ph](#)].
- [110] John F. Beacom, Nicole F. Bell, and Scott Dodelson. “Neutrinoless universe”. *Phys.Rev.Lett.* 93 (2004), p. 121302. DOI: [10.1103/PhysRevLett.93.121302](https://doi.org/10.1103/PhysRevLett.93.121302). arXiv: [astro-ph/0404585](https://arxiv.org/abs/astro-ph/0404585) [[astro-ph](#)].
- [111] Marco Cirelli and Alessandro Strumia. “Cosmology of neutrinos and extra light particles after WMAP3”. *JCAP* 0612 (2006), p. 013. DOI: [10.1088/1475-7516/2006/12/013](https://doi.org/10.1088/1475-7516/2006/12/013). arXiv: [astro-ph/0607086](https://arxiv.org/abs/astro-ph/0607086) [[astro-ph](#)].
- [112] Alexander Friedland, Kathryn M. Zurek, and Sergei Bashinsky. “Constraining Models of Neutrino Mass and Neutrino Interactions with the Planck Satellite” (2007). arXiv: [0704.3271](https://arxiv.org/abs/0704.3271) [[astro-ph](#)].
- [113] Anders Basboll, Ole Eggers Bjaelde, Steen Hannestad, and Georg G. Raffelt. “Are cosmological neutrinos free-streaming?” *Phys.Rev.* D79 (2009), p. 043512. DOI: [10.1103/PhysRevD.79.043512](https://doi.org/10.1103/PhysRevD.79.043512). arXiv: [0806.1735](https://arxiv.org/abs/0806.1735) [[astro-ph](#)].
- [114] Antony Lewis and Sarah Bridle. “Cosmological parameters from CMB and other data: A Monte Carlo approach”. *Phys.Rev.* D66 (2002), p. 103511. DOI: [10.1103/PhysRevD.66.103511](https://doi.org/10.1103/PhysRevD.66.103511). arXiv: [astro-ph/0205436](https://arxiv.org/abs/astro-ph/0205436) [[astro-ph](#)].
- [115] P.A.R. Ade et al. “Planck 2013 results. XV. CMB power spectra and likelihood”. *Astron.Astrophys.* 571 (2014), A15. DOI: [10.1051/0004-6361/201321573](https://doi.org/10.1051/0004-6361/201321573). arXiv: [1303.5075](https://arxiv.org/abs/1303.5075) [[astro-ph.CO](#)].
- [116] Adam G. Riess et al. “A 3% Solution: Determination of the Hubble Constant with the Hubble Space Telescope and Wide Field Camera 3”. *Astrophys.J.* 730 (2011), p. 119. DOI: [10.1088/0004-637X/732/2/129](https://doi.org/10.1088/0004-637X/732/2/129), [10.1088/0004-637X/730/2/119](https://doi.org/10.1088/0004-637X/730/2/119). arXiv: [1103.2976](https://arxiv.org/abs/1103.2976) [[astro-ph.CO](#)].
- [117] Lotty Ackerman, Matthew R. Buckley, Sean M. Carroll, and Marc Kamionkowski. “Dark Matter and Dark Radiation”. *Phys.Rev.* D79 (2009), p. 023519. DOI: [10.1103/PhysRevD.79.023519](https://doi.org/10.1103/PhysRevD.79.023519). arXiv: [0810.5126](https://arxiv.org/abs/0810.5126) [[hep-ph](#)].
- [118] Brando Bellazzini, Mathieu Cliche, and Philip Tanedo. “Effective theory of self-interacting dark matter”. *Phys.Rev.* D88.8 (2013), p. 083506. DOI: [10.1103/PhysRevD.88.083506](https://doi.org/10.1103/PhysRevD.88.083506). arXiv: [1307.1129](https://arxiv.org/abs/1307.1129).
- [119] Francis-Yan Cyr-Racine, Roland de Putter, Alvis Raccanelli, and Kris Sigurdson. “Constraints on Large-Scale Dark Acoustic Oscillations from Cosmology”. *Phys.Rev.* D89.6 (2014), p. 063517. DOI: [10.1103/PhysRevD.89.063517](https://doi.org/10.1103/PhysRevD.89.063517). arXiv: [1310.3278](https://arxiv.org/abs/1310.3278) [[astro-ph.CO](#)].
- [120] Georg G Raffelt. *Stars as laboratories for fundamental physics: The astrophysics of neutrinos, axions, and other weakly interacting particles*. University of Chicago press, 1996.

Bibliography

- [121] Mark Vogelsberger, Jesus Zavala, and Abraham Loeb. “Subhaloes in Self-Interacting Galactic Dark Matter Haloes”. *Mon.Not.Roy.Astron.Soc.* 423 (2012), p. 3740. DOI: [10.1111/j.1365-2966.2012.21182.x](https://doi.org/10.1111/j.1365-2966.2012.21182.x). arXiv: [1201.5892](https://arxiv.org/abs/1201.5892) [[astro-ph.CO](#)].
- [122] Mark Vogelsberger et al. “Properties of galaxies reproduced by a hydrodynamic simulation”. *Nature* 509 (2014), pp. 177–182. DOI: [10.1038/nature13316](https://doi.org/10.1038/nature13316). arXiv: [1405.1418](https://arxiv.org/abs/1405.1418) [[astro-ph.CO](#)].
- [123] Joop Schaye et al. “The EAGLE project: Simulating the evolution and assembly of galaxies and their environments”. *Mon.Not.Roy.Astron.Soc.* 446 (2015), pp. 521–554. DOI: [10.1093/mnras/stu2058](https://doi.org/10.1093/mnras/stu2058). arXiv: [1407.7040](https://arxiv.org/abs/1407.7040) [[astro-ph.GA](#)].
- [124] Till Sawala et al. “Local Group galaxies emerge from the dark” (2014). arXiv: [1412.2748](https://arxiv.org/abs/1412.2748) [[astro-ph.GA](#)].
- [125] Kyle A. Oman et al. “The unexpected diversity of dwarf galaxy rotation curves” (2015). arXiv: [1504.01437](https://arxiv.org/abs/1504.01437) [[astro-ph.GA](#)].
- [126] Shunsaku Horiuchi et al. “Sterile neutrino dark matter bounds from galaxies of the Local Group”. *Phys.Rev.* D89.2 (2014), p. 025017. DOI: [10.1103/PhysRevD.89.025017](https://doi.org/10.1103/PhysRevD.89.025017). arXiv: [1311.0282](https://arxiv.org/abs/1311.0282) [[astro-ph.CO](#)].
- [127] Arka Banerjee, Amol Dighe, and Georg Raffelt. “Linearized flavor-stability analysis of dense neutrino streams”. *Phys.Rev.* D84 (2011), p. 053013. DOI: [10.1103/PhysRevD.84.053013](https://doi.org/10.1103/PhysRevD.84.053013). arXiv: [1107.2308](https://arxiv.org/abs/1107.2308) [[hep-ph](#)].
- [128] Sovan Chakraborty, Alessandro Mirizzi, Ninetta Saviano, and David de Sousa Seixas. “Suppression of the multi-azimuthal-angle instability in dense neutrino gas during supernova accretion phase” (2014). arXiv: [1402.1767](https://arxiv.org/abs/1402.1767) [[hep-ph](#)].
- [129] Sovan Chakraborty, Georg Raffelt, Hans-Thomas Janka, and Bernhard Mueller. “Supernova deleptonization asymmetry: Impact on self-induced flavor conversion” (2014). arXiv: [1412.0670](https://arxiv.org/abs/1412.0670) [[hep-ph](#)].
- [130] Srdjan Sarikas, David de Sousa Seixas, and Georg Raffelt. “Spurious instabilities in multi-angle simulations of collective flavor conversion”. *Phys.Rev.* D86 (2012), p. 125020. DOI: [10.1103/PhysRevD.86.125020](https://doi.org/10.1103/PhysRevD.86.125020). arXiv: [1210.4557](https://arxiv.org/abs/1210.4557) [[hep-ph](#)].
- [131] Huaiyu Duan and Shashank Shalgar. “Flavor instabilities in the neutrino line model”. *Phys.Lett.* B747 (2015), pp. 139–143. DOI: [10.1016/j.physletb.2015.05.057](https://doi.org/10.1016/j.physletb.2015.05.057). arXiv: [1412.7097](https://arxiv.org/abs/1412.7097) [[hep-ph](#)].
- [132] Francesco Ginelli, Hugues Chaté, Roberto Livi, and Antonio Politi. “Covariant Lyapunov vectors”. *Journal of Physics A: Mathematical and Theoretical* 46.25 (2013), p. 254005. DOI: [10.1088/1751-8113/46/25/254005](https://doi.org/10.1088/1751-8113/46/25/254005).
- [133] A. Wolf. “Quantifying chaos with Lyapunov exponents”. *Chaos*. Ed. by A. Wolf. Princeton University Press, 1986. Chap. 13.
- [134] Giancarlo Benettin, Luigi Galgani, Antonio Giorgilli, and Jean-marie Strelcyn. “Lyapunov Characteristic Exponents for smooth dynamical systems and for hamiltonian systems; a method for computing all of them. Part 1: Theory”. *Meccanica* 15.1 (1980), pp. 9–20. DOI: [10.1007/BF02128236](https://doi.org/10.1007/BF02128236).

- [135] Pavel V. Kuptsov and Ulrich Parlitz. “Theory and Computation of Covariant Lyapunov Vectors”. *Journal of Nonlinear Science* 22.5 (2012), pp. 727–762. DOI: [10.1007/s00332-012-9126-5](https://doi.org/10.1007/s00332-012-9126-5).
- [136] Christopher L. Wolfe and Roger M. Samelson. “An efficient method for recovering Lyapunov vectors from singular vectors”. *Tellus A* 59.3 (2007), pp. 355–366. DOI: [10.1111/j.1600-0870.2007.00234.x](https://doi.org/10.1111/j.1600-0870.2007.00234.x).
- [137] V.I. Oseledets. “A multiplicative ergodic theorem. Lyapunov characteristic numbers for dynamical systems.” *Trans. Moscow Math. Soc.* 19 (1968), pp. 197–231.
- [138] F. Ginelli et al. “Characterizing Dynamics with Covariant Lyapunov Vectors”. *Physical Review Letters* 99.13 (2007), p. 130601. DOI: [10.1103/PhysRevLett.99.130601](https://doi.org/10.1103/PhysRevLett.99.130601).
- [139] Isaac Goldhirsch, Pierre-Louis Sulem, and Steven A. Orszag. “Stability and Lyapunov stability of dynamical systems: A differential approach and a numerical method”. *Physica D: Nonlinear Phenomena* 27.3 (1987), pp. 311–337. DOI: [10.1016/0167-2789\(87\)90034-0](https://doi.org/10.1016/0167-2789(87)90034-0).
- [140] Robert Foot, Mark J. Thomson, and R.R. Volkas. “Large neutrino asymmetries from neutrino oscillations”. *Phys.Rev.* D53 (1996), pp. 5349–5353. DOI: [10.1103/PhysRevD.53.R5349](https://doi.org/10.1103/PhysRevD.53.R5349). arXiv: [hep-ph/9509327](https://arxiv.org/abs/hep-ph/9509327) [hep-ph].
- [141] Xiang-Dong Shi. “Chaotic amplification of neutrino chemical potentials by neutrino oscillations in big bang nucleosynthesis”. *Phys.Rev.* D54 (1996), pp. 2753–2760. DOI: [10.1103/PhysRevD.54.2753](https://doi.org/10.1103/PhysRevD.54.2753). arXiv: [astro-ph/9602135](https://arxiv.org/abs/astro-ph/9602135) [astro-ph].
- [142] Kari Enqvist, Kimmo Kainulainen, and Antti Sorri. “On chaoticity of the amplification of the neutrino asymmetry in the early universe”. *Phys.Lett.* B464 (1999), pp. 199–205. DOI: [10.1016/S0370-2693\(99\)00966-1](https://doi.org/10.1016/S0370-2693(99)00966-1). arXiv: [hep-ph/9906452](https://arxiv.org/abs/hep-ph/9906452) [hep-ph].
- [143] A. D. Dolgov, S. H. Hansen, S. Pastor, and D. V. Semikoz. “Neutrino oscillations in the early universe: How large lepton asymmetry can be generated?” *Astropart. Phys.* 14 (2000), pp. 79–90. DOI: [10.1016/S0927-6505\(00\)00111-0](https://doi.org/10.1016/S0927-6505(00)00111-0). arXiv: [hep-ph/9910444](https://arxiv.org/abs/hep-ph/9910444) [hep-ph].
- [144] Antti Sorri. “Physical origin of ‘chaoticity’ of neutrino asymmetry”. *Phys. Lett.* B477 (2000), pp. 201–207. DOI: [10.1016/S0370-2693\(00\)00203-3](https://doi.org/10.1016/S0370-2693(00)00203-3). arXiv: [hep-ph/9911366](https://arxiv.org/abs/hep-ph/9911366) [hep-ph].
- [145] Robert Buras. “Sterile neutrinos in big bang nucleosynthesis”. MA thesis. Munich, Tech. U., 1999. arXiv: [hep-ph/0002086](https://arxiv.org/abs/hep-ph/0002086) [hep-ph].
- [146] P. Di Bari, Robert Foot, R.R. Volkas, and Y.Y.Y. Wong. “Comment on ‘Neutrino oscillations in the early universe: How can large lepton asymmetry be generated?’” *Astropart.Phys.* 15 (2001), pp. 391–412. DOI: [10.1016/S0927-6505\(00\)00160-2](https://doi.org/10.1016/S0927-6505(00)00160-2). arXiv: [hep-ph/0008245](https://arxiv.org/abs/hep-ph/0008245) [hep-ph].
- [147] Poul-Erik N. Braad and Steen Hannestad. “On the Chaoticity of active-sterile neutrino oscillations in the early universe” (2000). arXiv: [hep-ph/0012194](https://arxiv.org/abs/hep-ph/0012194) [hep-ph].

Bibliography

- [148] Kevork N. Abazajian and Prateek Agrawal. “Chaos, Determinacy and Fractals in Active-Sterile Neutrino Oscillations in the Early Universe”. *JCAP* 0810 (2008), p. 006. DOI: [10.1088/1475-7516/2008/10/006](https://doi.org/10.1088/1475-7516/2008/10/006). arXiv: [0807.0456](https://arxiv.org/abs/0807.0456) [[hep-ph](#)].
- [149] Xiang-Dong Shi and George M. Fuller. “Leptonic domains in the early universe and their implications”. *Phys. Rev. Lett.* 83 (1999), pp. 3120–3123. DOI: [10.1103/PhysRevLett.83.3120](https://doi.org/10.1103/PhysRevLett.83.3120). arXiv: [astro-ph/9904041](https://arxiv.org/abs/astro-ph/9904041) [[astro-ph](#)].
- [150] P. Di Bari and Robert Foot. “On the sign of the neutrino asymmetry induced by active-sterile neutrino oscillations in the early universe”. *Phys. Rev.* D61 (2000), p. 105012. DOI: [10.1103/PhysRevD.61.105012](https://doi.org/10.1103/PhysRevD.61.105012). arXiv: [hep-ph/9912215](https://arxiv.org/abs/hep-ph/9912215).
- [151] X. Shi, D.N. Schramm, and B.D. Fields. “Constraints on neutrino oscillations from big bang nucleosynthesis”. *Phys.Rev.* D48 (1993), pp. 2563–2572. DOI: [10.1103/PhysRevD.48.2563](https://doi.org/10.1103/PhysRevD.48.2563). arXiv: [astro-ph/9307027](https://arxiv.org/abs/astro-ph/9307027) [[astro-ph](#)].
- [152] Xiaoye S. Li. “An Overview of SuperLU: Algorithms, Implementation, and User Interface”. *ACM Trans. Math. Softw.* 31.3 (2005), pp. 302–325.
- [153] James W. Demmel, John R. Gilbert, and Xiaoye S. Li. “An Asynchronous Parallel Supernodal Algorithm for Sparse Gaussian Elimination”. *SIAM J. Matrix Analysis and Applications* 20.4 (1999), pp. 915–952.
- [154] E. Hairer, S.P. Nørsett, and G. Wanner. *Solving Ordinary Differential Equations: Stiff and differential-algebraic problems*. Springer series in computational mathematics. Springer-Verlag, 1993. ISBN: 9783540604525.
- [155] P. Di Bari. “Amplification of isocurvature perturbations induced by active sterile neutrino oscillations”. *Phys. Lett.* B482 (2000), pp. 150–160. DOI: [10.1016/S0370-2693\(00\)00487-1](https://doi.org/10.1016/S0370-2693(00)00487-1). arXiv: [hep-ph/9911214](https://arxiv.org/abs/hep-ph/9911214) [[hep-ph](#)].
- [156] Kari Enqvist, Kimmo Kainulainen, and Antti Sorri. “Creation of large spatial fluctuations in neutrino asymmetry by neutrino oscillations”. *JHEP* 04 (2001), p. 012. DOI: [10.1088/1126-6708/2001/04/012](https://doi.org/10.1088/1126-6708/2001/04/012). arXiv: [hep-ph/0012291](https://arxiv.org/abs/hep-ph/0012291) [[hep-ph](#)].
- [157] John F. Cherry et al. “Neutrino scattering and flavor transformation in supernovae”. *Phys. Rev. Lett.* 108 (2012), p. 261104. DOI: [10.1103/PhysRevLett.108.261104](https://doi.org/10.1103/PhysRevLett.108.261104). arXiv: [1203.1607](https://arxiv.org/abs/1203.1607) [[hep-ph](#)].
- [158] Alessandro Mirizzi, Gianpiero Mangano, and Ninetta Saviano. “Self-induced flavor instabilities of a dense neutrino stream in a two-dimensional model” (2015). arXiv: [1503.03485](https://arxiv.org/abs/1503.03485) [[hep-ph](#)].
- [159] Alessandro Mirizzi. “Breaking the symmetries of the bulb model in two-dimensional self-induced supernova neutrino flavor conversions” (2015). arXiv: [1506.06805](https://arxiv.org/abs/1506.06805) [[hep-ph](#)].
- [160] Alessandro Mirizzi. “Multi-azimuthal-angle effects in self-induced supernova neutrino flavor conversions without axial symmetry”. *Phys.Rev.* D88 (2013), p. 073004. arXiv: [1308.1402](https://arxiv.org/abs/1308.1402) [[hep-ph](#)].
- [161] Sovan Chakraborty and Alessandro Mirizzi. “Multi-azimuthal-angle instability for different supernova neutrino fluxes”. *Phys. Rev.* D90.3 (2014), p. 033004. DOI: [10.1103/PhysRevD.90.033004](https://doi.org/10.1103/PhysRevD.90.033004). arXiv: [1308.5255](https://arxiv.org/abs/1308.5255) [[hep-ph](#)].

- [162] Huaiyu Duan. “Flavor Oscillation Modes In Dense Neutrino Media”. *Phys.Rev.* D88 (2013), p. 125008. DOI: [10.1103/PhysRevD.88.125008](https://doi.org/10.1103/PhysRevD.88.125008). arXiv: [1309.7377](https://arxiv.org/abs/1309.7377) [[hep-ph](#)].
- [163] Andreu Esteban-Pretel et al. “Decoherence in supernova neutrino transformations suppressed by deleptonization”. *Phys.Rev.* D76 (2007), p. 125018. DOI: [10.1103/PhysRevD.76.125018](https://doi.org/10.1103/PhysRevD.76.125018). arXiv: [0706.2498](https://arxiv.org/abs/0706.2498) [[astro-ph](#)].
- [164] Gianluigi L. Fogli, Eligio Lisi, Antonio Marrone, and Alessandro Mirizzi. “Collective neutrino flavor transitions in supernovae and the role of trajectory averaging”. *JCAP* 0712 (2007), p. 010. DOI: [10.1088/1475-7516/2007/12/010](https://doi.org/10.1088/1475-7516/2007/12/010). arXiv: [0707.1998](https://arxiv.org/abs/0707.1998) [[hep-ph](#)].
- [165] Giancarlo Benettin, Luigi Galgani, Antonio Giorgilli, and Jean-Marie Strelcyn. “Lyapunov Characteristic Exponents for smooth dynamical systems and for hamiltonian systems; A method for computing all of them. Part 2: Numerical application”. *Meccanica* 15 (1 1980), pp. 21–30.
- [166] Gianpiero Mangano, Alessandro Mirizzi, and Ninetta Saviano. “Damping the neutrino flavor pendulum by breaking homogeneity”. *Phys. Rev.* D89.7 (2014), p. 073017. DOI: [10.1103/PhysRevD.89.073017](https://doi.org/10.1103/PhysRevD.89.073017). arXiv: [1403.1892](https://arxiv.org/abs/1403.1892) [[hep-ph](#)].
- [167] Irene Tamborra et al. “Self-sustained asymmetry of lepton-number emission: A new phenomenon during the supernova shock-accretion phase in three dimensions”. *Astrophys. J.* 792.2 (2014), p. 96. DOI: [10.1088/0004-637X/792/2/96](https://doi.org/10.1088/0004-637X/792/2/96). arXiv: [1402.5418](https://arxiv.org/abs/1402.5418) [[astro-ph.SR](#)].
- [168] Mikhail Danilov. “Sensitivity of DANSS detector to short range neutrino oscillations” (2014). arXiv: [1412.0817](https://arxiv.org/abs/1412.0817) [[physics.ins-det](#)].
- [169] A. P. Serebrov et al. “NEUTRINO4 experiment: preparations for search for sterile neutrino at 100 MW Reactor SM-3 at 6-12 Meters” (2012). arXiv: [1205.2955](https://arxiv.org/abs/1205.2955) [[hep-ph](#)].
- [170] J. Ashenfelter et al. “PROSPECT - A Precision Reactor Oscillation and Spectrum Experiment at Short Baselines”. *Community Summer Study 2013: Snowmass on the Mississippi (CSS2013) Minneapolis, MN, USA, July 29-August 6, 2013*. 2013. arXiv: [1309.7647](https://arxiv.org/abs/1309.7647) [[physics.ins-det](#)].
- [171] G. Bellini et al. “SOX: Short distance neutrino Oscillations with BoreXino”. *JHEP* 1308 (2013), p. 038. DOI: [10.1007/JHEP08\(2013\)038](https://doi.org/10.1007/JHEP08(2013)038). arXiv: [1304.7721](https://arxiv.org/abs/1304.7721) [[physics.ins-det](#)].
- [172] A. Gando et al. “CeLAND: search for a 4th light neutrino state with a 3 PBq ^{144}Ce - ^{144}Pr electron antineutrino generator in KamLAND” (2013). arXiv: [1312.0896](https://arxiv.org/abs/1312.0896) [[physics.ins-det](#)].
- [173] N. Agafonova et al. “Search for $\nu_\mu \rightarrow \nu_e$ oscillations with the OPERA experiment in the CNGS beam”. *JHEP* 1307 (2013), p. 004. DOI: [10.1007/JHEP07\(2013\)004](https://doi.org/10.1007/JHEP07(2013)004), [10.1007/JHEP07\(2013\)085](https://doi.org/10.1007/JHEP07(2013)085). arXiv: [1303.3953](https://arxiv.org/abs/1303.3953) [[hep-ex](#)].
- [174] Lisa Goodenough and Dan Hooper. “Possible Evidence For Dark Matter Annihilation In The Inner Milky Way From The Fermi Gamma Ray Space Telescope” (2009). arXiv: [0910.2998](https://arxiv.org/abs/0910.2998) [[hep-ph](#)].

Bibliography

- [175] Oscar Adriani et al. “An anomalous positron abundance in cosmic rays with energies 1.5-100 GeV”. *Nature* 458 (2009), pp. 607–609. DOI: [10.1038/nature07942](https://doi.org/10.1038/nature07942). arXiv: [0810.4995](https://arxiv.org/abs/0810.4995) [[astro-ph](#)].
- [176] M. Aguilar et al. “Electron and Positron Fluxes in Primary Cosmic Rays Measured with the Alpha Magnetic Spectrometer on the International Space Station”. *Phys. Rev. Lett.* 113 (2014), p. 121102. DOI: [10.1103/PhysRevLett.113.121102](https://doi.org/10.1103/PhysRevLett.113.121102).
- [177] Kfir Blum, Boaz Katz, and Eli Waxman. “AMS-02 Results Support the Secondary Origin of Cosmic Ray Positrons”. *Phys.Rev.Lett.* 111.21 (2013), p. 211101. DOI: [10.1103/PhysRevLett.111.211101](https://doi.org/10.1103/PhysRevLett.111.211101). arXiv: [1305.1324](https://arxiv.org/abs/1305.1324) [[astro-ph.HE](#)].
- [178] Timur Delahaye, Kumiko Kotera, and Joseph Silk. “What could we learn from a sharply falling positron fraction?” *Astrophys. J.* 794.2 (2014), p. 168. DOI: [10.1088/0004-637X/794/2/168](https://doi.org/10.1088/0004-637X/794/2/168). arXiv: [1404.7546](https://arxiv.org/abs/1404.7546) [[astro-ph.HE](#)].
- [179] N. Aghanim et al. “Planck 2015 results. XI. CMB power spectra, likelihoods, and robustness of parameters”. *Submitted to: Astron. Astrophys.* (2015). arXiv: [1507.02704](https://arxiv.org/abs/1507.02704) [[astro-ph.CO](#)].
- [180] Roland Diehl et al. “Radioactive Al-26 and massive stars in the galaxy”. *Nature* 439 (2006), pp. 45–47. DOI: [10.1038/nature04364](https://doi.org/10.1038/nature04364). arXiv: [astro-ph/0601015](https://arxiv.org/abs/astro-ph/0601015) [[astro-ph](#)].