

Dirac solitons in general relativity:

many-fermion states, singular solutions &
the inclusion of a Higgs mechanism

Peter Edward Duncan Leith



University of
St Andrews

This thesis is submitted in partial fulfilment for the degree of

Doctor of Philosophy (PhD)

at the University of St Andrews

September 2022

Candidate's declaration

I, Peter Edward Duncan Leith, do hereby certify that this thesis, submitted for the degree of PhD, which is approximately 60,000 words in length, has been written by me, and that it is the record of work carried out by me, or principally by myself in collaboration with others as acknowledged, and that it has not been submitted in any previous application for any degree. I confirm that any appendices included in my thesis contain only material permitted by the 'Assessment of Postgraduate Research Students' policy.

I was admitted as a research student at the University of St Andrews in September 2018.

I received funding from an organisation or institution and have acknowledged the funder(s) in the full text of my thesis.

Date 17/11/22

Signature of candidate

Supervisor's declaration

I hereby certify that the candidate has fulfilled the conditions of the Resolution and Regulations appropriate for the degree of PhD in the University of St Andrews and that the candidate is qualified to submit this thesis in application for that degree. I confirm that any appendices included in the thesis contain only material permitted by the 'Assessment of Postgraduate Research Students' policy.

Date 17.xi.22.

Signature of supervisor

Permission for publication

In submitting this thesis to the University of St Andrews we understand that we are giving permission for it to be made available for use in accordance with the regulations of the University Library for the time being in force, subject to any copyright vested in the work not being affected thereby. We also understand, unless exempt by an award of an embargo as requested below, that the title and the abstract will be published, and that a copy of the work may be made and supplied to any bona fide library or research worker, that this thesis will be electronically accessible for personal or research use and that the library has the right to migrate this thesis into new electronic forms as required to ensure continued access to the thesis.

I, Peter Edward Duncan Leith, confirm that my thesis does not contain any third-party material that requires copyright clearance.

The following is an agreed request by candidate and supervisor regarding the publication of this thesis:

Printed copy

No embargo on print copy.

Electronic copy

No embargo on electronic copy.

Date 17/11/22

Signature of candidate

Date 17.xi.22

Signature of supervisor

Underpinning Research Data or Digital Outputs

Candidate's declaration

I, Peter Edward Duncan Leith, hereby certify that no requirements to deposit original research data or digital outputs apply to this thesis and that, where appropriate, secondary data used have been referenced in the full text of my thesis.

Date 17/11/22

Signature of candidate

Abstract

The study of gravitationally localised quantum states, in which quantum particles are bound together by their mutual gravitational interaction, has been a topic of considerable research for over 50 years. Stemming from John Wheeler’s initial concept of an electromagnetic ‘geon’ [1], focus quickly converged on scalar fields with the introduction of the objects today referred to as ‘boson stars’ [2]. It was not until more recently, however, that the fermionic sector was properly addressed by Finster, Smoller & Yau [3], who successfully constructed the first numerical solutions to the coupled Einstein–Dirac system. The resulting ‘particle-like’ objects, comprising pairs of neutral fermions, have become known as ‘Dirac solitons’ or ‘Dirac stars’, and have been the focus of significantly less study than their bosonic counterparts.

This thesis aims to expand the knowledge of Dirac solitons in a number of ways. First, we conduct a detailed study of the many-fermion system (up to a total of 90 fermions), and interpret the structure of the resulting solutions in terms of a ‘self-trapping’ effect. We also find somewhat unexpected results for the behaviour of the excited states of these many-fermion solitons. Second, we present particle-like solutions to the minimally-coupled Einstein–Dirac–Higgs system, and show that, in the presence of strong coupling, a mass-scale separation can occur, in which the total mass of the constituent fermions far outweighs the gravitational mass of the state. Finally, we introduce new singular solutions to the Einstein–Dirac system, including the first normalised analytic solution. The properties of these are somewhat unusual, but we show that many are related (and indeed can be smoothly connected) to the non-singular Einstein–Dirac states.

Acknowledgements

I would like to thank my trio of supervisors, Dr Chris Hooley, Prof. Keith Horne and Prof. David Dritschel, for providing the inspiration behind this project and for the many productive discussions in our weekly meetings. I would also like to thank my parents and my sister for their continual support over the years.

This work was jointly funded by a St Leonards scholarship from the University of St Andrews and by the UKRI via an EPSRC grant (Grant No. EP/R513337/1).

Contents

Abstract	i
Acknowledgements	ii
1 Introduction & Background	1
1.1 Outline	1
1.2 Quantum gravity	3
1.3 Gravitationally localised quantum states	5
1.3.1 Geons, boson stars & related objects	5
1.3.2 Einstein–Dirac solitons	6
1.4 General relativity	8
1.4.1 Notation & Conventions	8
1.4.2 Spacetime metrics & curvature	9
1.4.3 The Einstein equations	10
1.4.4 The Schwarzschild metric	11
1.4.5 Mass in general relativity	12
1.5 Spinors in curved spacetime	13
1.5.1 The Dirac equation in flat spacetime	13
1.5.2 The vierbein formalism	14
1.5.3 The Dirac equation in curved spacetime	15
2 Einstein–Dirac solitons	17
2.1 Equations of motion	17
2.1.1 The two-fermion system	17
2.1.2 The many-fermion system	20
2.2 Generating particle-like solutions	22
2.2.1 Boundary conditions	22
2.2.2 The rescaling method	23
2.3 Useful quantities	25
2.3.1 Fermion densities	25
2.3.2 Gravitational mass	26

2.3.3	Curvature invariants	27
2.4	Two-fermion states: a review	28
2.4.1	Individual states	28
2.4.2	Families of states	30
2.5	Structure of Dirac solitons	33
2.5.1	The power-law solution	33
2.5.2	Zonal structure	33
2.5.3	Infinite-redshift states	35
3	The fermion self-trapping effect	37
3.1	Many-fermion states	37
3.2	Optical geometry	41
3.2.1	Optical embedding diagrams	41
3.2.2	Geodesic motion	45
3.3	Self-trapping interpretation	47
3.4	Families of many-fermion states	50
3.4.1	Spiralling behaviour	50
3.4.2	Calculating the fermion energy	53
3.5	Discussion	56
4	Many-fermion excited states	59
4.1	Numerical method	59
4.1.1	Review of the two-fermion case	59
4.1.2	The many-fermion system	61
4.2	1st excited states ($n = 2$)	63
4.2.1	Varying N_f	63
4.2.2	The $N_f = 12$ system: a closer look	65
4.2.3	The $N_f = 20$ system	68
4.2.4	The $N_f = 38$ system	70
4.3	Higher excited states	72
4.3.1	Small fermion number	73
4.3.2	The $N_f = 38$ system: $n = 4$ states	74
4.3.3	Varying N_f : A summary	76
4.4	Discussion	78
5	Einstein–Dirac–Higgs solitons	82
5.1	The Einstein–Dirac–Higgs system	83
5.1.1	Setup	83
5.1.2	Equations of motion	85
5.1.3	Particle-like solutions	86
5.2	Einstein–Dirac–Higgs solitons	89
5.2.1	Numerical method	89

5.2.2	Individual states	91
5.2.3	Families of states	92
5.3	Mass-scale separation	95
5.3.1	Strong coupling	96
5.3.2	Large- r analysis	98
5.3.3	Discussion	103
5.4	Other results	105
5.4.1	Zonal structure & infinite-redshift states	105
5.4.2	Many-fermion states	108
5.5	Discussion	109
6	Singular solutions	112
6.1	Small- r analysis	112
6.1.1	Case [A]: $\{\alpha, \beta, A, T\} \sim \{r^{\kappa/2}, r^{\kappa/2+1}, r^0, r^0\}$ (FSY states)	114
6.1.2	Case [B]: $\{\alpha, \beta, A, T\} \sim \{r^{\kappa/2}, r^{3\kappa/2+1}, r^0, r^0\}$	115
6.1.3	Case [C]: $\{\alpha, \beta, A, T\} \sim \{r^1, r^1, r^0, r^{-1}\}$ ($z = \infty$ states)	116
6.1.4	Case [D]: $\{\alpha, \beta, A, T\} \sim \{r^\Omega, r^1, r^0, r^{-\Omega}\}$	117
6.1.5	Case [E]: $\{\alpha, \beta, A, T\} \sim \{r^0, r^0, r^{-2}, r^{-1}\}$	119
6.1.6	Case [F]: $\{\alpha, \beta, A, T\} \sim \{r^{\kappa/2}, r^{-\kappa/2}, r^0, r^1\}$	122
6.1.7	Case [G]: $\{\alpha, \beta, A, T\} \sim \{r^0, r^{-1}, r^0, r^2\}$	122
6.1.8	Case [H]: $\{\alpha, \beta, A, T\} \sim \{r^{-\kappa/2+1}, r^{-\kappa/2}, r^0, r^\kappa\}$	125
6.1.9	Case [I]: $\{\alpha, \beta, A, T\} \sim \{r^0, r^0, r^{-1}, r^0\}$	126
6.1.10	Case [J]: $\{\alpha, \beta, A, T\} \sim \{r^{3/2}, r^0, r^{-1}, r^{1/2}\}$	128
6.1.11	Case [K]: $\{\alpha, \beta, A, T\} \sim \{r^0, r^{3/2}, r^{-1}, r^{1/2}\}$	131
6.1.12	Case [L]: $\{\alpha, \beta, A, T\} \sim \{r^0, r^0, r^{-1}, r^{1/2}\}$	131
6.2	$\omega = 0$, $A = 1$ analytic solution	133
6.2.1	The two-fermion solution	134
6.2.2	The many-fermion solution	138
6.2.3	Interpretation	139
6.3	Discussion	140
7	Summary & Discussion	142
A	Low-redshift relations	144
B	Data for Figures	147
	Bibliography	150

Chapter 1

Introduction & Background

The issue of reconciling quantum mechanics with general relativity is one of the outstanding questions in modern theoretical physics. Given the current abundance and diversity of candidate theories aimed at addressing this ‘quantum gravity’ problem, combined with an associated lack of testable experimental hypotheses, it is clear that some form of direction is required. Historically speaking, progress in physics is often made by analysing the behaviour of approximate theories, in which the more troublesome aspects of the full theory are purposely avoided, allowing the analysis of otherwise unattainable phenomena. Applying this to the problem of quantum gravity, one obvious simplification is to treat the gravitational field as a purely classical object, but it can also be advantageous to approximate the matter sector, in order to avoid the calculational issues associated with quantum field theory.

This latter approach is employed within the semi-classical Einstein–Dirac formalism, where the interaction between a quantum wavefunction and a classical gravitational field is considered. The associated particle-like states, referred to here as ‘Dirac solitons’, provide an intriguing toy system in which the full dynamics of general relativity can be explored, while still retaining a quantum aspect. Although the direct application of these objects to the real world is somewhat debatable, their analysis may nonetheless prove useful in indicating potential directions towards a full theory of quantum gravity.

1.1 Outline

This thesis is arranged as follows. The subsequent sections in this first chapter introduce the wider context within which Dirac solitons are situated, starting with a discussion of approaches to quantum gravity, followed by a non-technical review of gravitationally localised quantum systems, including geons, boson stars and Dirac solitons. Thereafter follows a brief summary of the necessary physics background required throughout the

thesis, namely an understanding of general relativity and the treatment of spin in curved spacetime.

Chapter 2 contains a more technical introduction to Dirac solitons, starting with a brief derivation of the equations of motion first obtained by Finster, Smoller & Yau [3]. We then detail the method used to generate numerical solutions to these equations, and present a summary of the main known results for the two-fermion system. We also review in detail the results concerning the internal zonal structure of Dirac solitons discovered by Bakucz Canário *et al.* [4], as this will be relevant to much of the material in the subsequent chapters.

Chapter 3 is the first of four research chapters, and concerns the discovery of a fermion self-trapping effect. Results from this chapter have been published in ref. [5]. We begin by reviewing the generalisation of the two-fermion system to that of many fermions, before presenting numerical results illustrating how the behaviour of the system changes with increasing fermion number. We then introduce the concept of optical geometry as a convenient means for visualising the spacetime of Dirac solitons, and interpret the resulting structures in terms of a fermion self-trapping effect. Thereafter follows a discussion as to whether this effect can explain the characteristic spiralling behaviour of the soliton families.

Chapter 4 is a natural continuation of the preceding chapter, dealing with the behaviour of excited states in the many-fermion Einstein–Dirac system. This begins with a discussion of the challenges to obtaining numerical solutions, due to the onset of multivaluedness that occurs, followed by results illustrating the implications of this multiplicity on the soliton families. The observed effects are then explained by analysing the internal structure of individual states with high fermion number. Results from this chapter have been published in ref. [6].

Chapter 5 concerns the inclusion of a minimally-coupled Higgs field to the Einstein–Dirac system. We first review the Higgs mechanism and its application in this case, before describing how numerical solutions to the Einstein–Dirac–Higgs system can be generated. The properties of these are analysed and the appearance of a mass-scale separation is illustrated, in which the ADM mass becomes parametrically smaller than the total mass of the constituent fermions. We discuss possible explanations for this behaviour, including those drawn from results of a large- r analysis, before concluding with a brief presentation of infinite-redshift and many-fermion solutions. A summary of the results from this chapter can be found in ref. [7].

In Chapter 6, we return to the bare Einstein–Dirac system, and present a selection of new solutions, all of which exhibit a naked singularity at their centre. We show that these are obtained by varying the small- r asymptotic expansions used when generating solutions, and provide a detailed classification of the possible forms. We go on to analyse

the often unusual properties of these singular solutions and illustrate their connection to the original Einstein–Dirac states. We finish by presenting a new analytic solution to the system, which, although singular, can be correctly normalised, with the resulting states exhibiting zero total gravitational mass.

We conclude with a brief summary and discussion in Chapter 7.

1.2 Quantum gravity

Individually, the theories of general relativity and quantum mechanics represent two of the major triumphs of 20th century physics, but taken together they appear fundamentally incompatible. Both theories have proved highly successful within their own regimes, with quantum mechanics (through its modern guise as quantum field theory) underpinning the Standard Model of particle physics, and general relativity vital in our understanding of modern astrophysics and cosmology. It is the general consensus, however, that a complete understanding of gravity extends beyond that provided by general relativity, and that inevitably the theory must be altered at the quantum scale. It is hoped that, if such a theory of ‘quantum gravity’ can be found, it would not only provide the required bridge between quantum mechanics and general relativity, but also eliminate the somewhat disquieting presence of singularities within the latter.

Naive attempts at ‘quantising’ the classical gravitational field, in analogy with the other fundamental forces of nature, fail, however, due to the inherent non-renormalisability of the resulting theory [8]. It is therefore generally considered that an entirely new approach is required, but despite intensive research over the past 70 years or so, no fully satisfactory theory of quantum gravity has yet been formulated. Examples of candidates that aim to construct such a theory from first principles include string theory [9], quantum loop gravity [10] and causal set theory [11], while alternative approaches include that of asymptotic safety [12], in which the problem of renormalisability is addressed by introducing an ultraviolet fixed point to the field theory.

Common to all these theories is the assumption that Einstein’s general relativity is the correct theory of gravity at large (non-quantum) scales. The discoveries of dark matter and dark energy, however, have introduced an element of doubt to this position, and as a consequence a variety of modified gravity theories have been suggested. Although the majority of these are viewed as irrelevant to the quantum nature of gravity, one example of a theory that is not is conformal gravity, the quantum version of which is claimed to be fully renormalisable [13]. In addition, there are some physicists, perhaps most notably Roger Penrose, who argue that the quantisation of the gravitational field is in fact unnecessary, and that instead it is quantum mechanics that should be modified to accommodate the principles of general relativity [14]. Full details of how to achieve this, however, have yet

to be formulated.

It should be clear from the above discussion that there is thus no clear consensus on how to approach the problem of quantum gravity, let alone a theory with which detailed calculations can be performed. How then are we to proceed if we wish to study a particular system where the effects of quantum mechanics and general relativity are both important (black holes for instance)? The answer is usually found in the approach known as ‘semiclassical gravity’, in which the gravitational field is treated purely as a classical entity. This is often justified by arguing that any quantum aspects of gravity should only have a significant effect below distances approaching the Planck length, defined as:

$$l_{\text{P}} = \sqrt{\frac{\hbar G}{c^3}} \approx 1.616 \times 10^{-35} \text{ m}, \quad (1.1)$$

or in situations where the spacetime curvature approaches or exceeds the order of the Planck curvature, defined as:

$$\mathcal{R}_{\text{P}} = \frac{c^3}{\hbar G} \approx 3.829 \times 10^{69} \text{ m}^{-2}. \quad (1.2)$$

Thus if we restrict ourselves to systems which do not approach these scales, we can safely assume that quantum gravity corrections are negligible, and hence treat the gravitational field classically.

In the majority of cases, however, this approximation alone proves insufficient to make much progress. This is ultimately due to one essential feature of general relativity — the back-reaction of the matter involved on the spacetime itself. As John Wheeler concisely put it:

“*Spacetime tells matter how to move; matter tells spacetime how to curve*”. [15]

This proves problematic since in order to construct a quantum field theory, a fixed background upon which to do so is required. To include back-reaction, however, the matter distribution and spacetime structure must be determined simultaneously, a calculation that is beyond the current capabilities of quantum field theory. One solution to this quandary is to make a further approximation, valid when the effects of back-reaction are negligible for the system in question. In such cases, the spacetime is treated as a fixed background, allowing a quantum field theory to be constructed and solved perturbatively. Often, it is then possible for any (small) neglected effects of back-reaction to be subsequently included by an iterative process, for example. This approach, known under the blanket term of ‘quantum fields in curved spacetime’ has proved highly successful, most notably perhaps in the prediction of Hawking radiation from black holes [16], and its application to cosmological inflation [17].

What should be the approach, however, when confronted with quantum systems in

which the effects of back-reaction cannot be assumed negligible? This is the case, for example, for gravitationally localised quantum states, the properties and indeed very existence of which rely heavily on the presence of back-reaction. For such systems, one solution is to abandon quantum field theory entirely, and instead treat matter as a simple quantum wavefunction. Coupling general relativity to what is then effectively a classical field (with added properties) is fairly straightforward, allowing the full non-linear dynamics of general relativity to be restored. It is within this general framework that the objects under consideration in this thesis, Dirac solitons, are situated.

1.3 Gravitationally localised quantum states

1.3.1 Geons, boson stars & related objects

Before discussing Dirac solitons themselves, we first briefly review the history and development of the related objects that preceded them. The first of these is the ‘electromagnetic geon’, a singularity-free, particle-like object existing within the coupled Einstein–Maxwell system. Initially introduced in 1955 by Wheeler [1], these entities can be thought of as standing electromagnetic waves held together by their own self-gravity. A similar object, consisting purely of gravitational waves (dubbed a ‘gravitational geon’), was subsequently discovered by Brill & Hartle [18], this being an approximate solution to the vacuum Einstein equations. It was later shown that in fact these two systems were governed by the same set of leading-order equations [19]. Unfortunately, interest in both the electromagnetic and gravitational geon was relatively short-lived, since it was discovered that the solutions were most likely unstable [20], although more recently stable gravitational geons have been argued to exist in anti-de Sitter spacetimes [21].

Geon-like objects for bosonic scalar fields were first constructed independently in 1968 by Feinblum & McKinley [22] and Kaup [2], these being localised solutions to the coupled Einstein–Klein–Gordon equations. Initially referred to as ‘Klein–Gordon geons’, these are better known today as ‘boson stars’. Unlike the electromagnetic and gravitational geons, however, a stable branch of boson star solutions is known to exist [23–25], and thus significant interest has been generated in these objects. In their usual formulation, boson stars can be interpreted as macroscopic quantum states somewhat akin to Bose–Einstein condensates, consisting of large numbers of bosons bound together gravitationally but prevented from collapse by the effects of the uncertainty principle.

We shall not attempt to provide a comprehensive overview of the subsequent literature on boson stars (excellent reviews of which can be found in refs. [26] and [27]), and instead highlight only a few aspects of interest. Boson stars share many similarities with compact astrophysical objects, for example the presence of a maximum gravitational mass analogous to the Chandrasekhar limit. One finds that this maximum mass is inversely proportional to

the boson mass itself, thus suggesting that boson stars consisting of ultra-light bosons may be able to form stellar-sized objects. This feature has generated significant astrophysical interest, with boson stars having been proposed as candidates for dark matter in galactic halos [28, 29], and as black hole mimickers [30, 31]. Such applications of course rely on the existence of a fundamental scalar matter particle (the axion for example [32]), thus invoking theories beyond the Standard Model.

The dynamics of boson stars is also a subject worth reviewing. As mentioned, a stable branch of states is known to exist, but what is the fate of solutions on the unstable branch? The first attempt at addressing this question was by Seidel & Suen [33], who performed a numerical evolution of the Klein–Gordon system in the context of spherically symmetric perturbations. The results of this analysis indicated that unstable boson stars can have two distinct fates, dependent on the type of initial perturbation — collapse to a black hole, or migration to a solution on the stable branch with an associated expulsion of scalar field to infinity. Later work by Guzmán [34, 35], however, reports instead collapse to a black hole or total dispersal to infinity, depending on the binding energy of the initial state. Although these results may appear contradictory, it is likely that the ultimate fate of unstable boson stars is highly sensitive to the type of initial perturbation.

Finally, we note that gravitationally localised states have also been constructed for massive spin-1 bosons, these being referred to as ‘Proca stars’ [36]. There have also been various studies of geon-like objects in which both bosons and fermions are present (so-called ‘fermion-boson’ stars) [37, 38], although in these cases the fermionic component is treated using a fluid approximation.

1.3.2 Einstein–Dirac solitons

We now review the literature relating to gravitationally localised states consisting of Dirac fermions. These have been known under various guises but today are commonly referred to as either ‘Dirac stars’ or ‘Dirac solitons’. It is the latter that we shall utilise here. Constructing such objects is significantly more challenging than for the bosonic case, mainly due to the complexities associated with formulating the Dirac equation in curved spacetime (see section 1.5), but also since the Pauli exclusion principle must be considered. Dirac solitons have therefore been the subject of significantly less research than their bosonic counterparts.

Early analysis by Ruffini & Bonazzola [39] and Lee & Pang [40] avoided many of these complications by relying on elements of approximation, motivated by considering Fermi–Dirac statistics and equations of state. As such their results are applicable only for systems comprising large numbers of fermions, and at sufficiently low density. It was not until 1999 that (numerical) localised solutions to the full Einstein–Dirac system were finally constructed by Finster, Smoller & Yau [3]. The resulting particle-like states are

spherically symmetric, by virtue of comprising pairs of fermions arranged in singlet states, and have the desirable property of being free from singularities. This initial analysis also established the presence of a stable branch of solutions (at least with respect to spherically symmetric perturbations), in analogy with that found for boson stars.

The same authors subsequently extended their formalism to include charged fermions [41], discovering similar stable, particle-like solutions in the coupled Einstein–Dirac–Maxwell system. Interestingly, bound states are shown to exist even when the overall interaction is repulsive in the Newtonian limit. The addition of the electroweak interaction has also been studied [42] by considering the Einstein–Dirac–Yang–Mills system, and again stable localised solutions (in this case single-particle states) have been found. Detailed analysis involving black holes within the Einstein–Dirac system and its extensions has also been performed [43–47], with the results in all cases proving that the Dirac fields must ultimately disappear within the black hole horizon. This is perhaps surprising since one might expect states with non-zero angular momentum to form stable orbits around the black hole; these analyses show that this cannot be the case. In addition, rigorous proofs of existence of Dirac solitons can be found in refs. [48–50], and a comprehensive analysis of their Newtonian limit (solutions of the Newton–Schrödinger system) in refs. [51] and [52].

More recently, particle-like solutions have been generated for the (conformally-coupled) Einstein–Dirac–Higgs and Einstein–Dirac–Maxwell–Higgs systems [53]. We shall provide a more thorough analysis of the former (within the minimally-coupled framework) in Chapter 5. The internal structure of Dirac solitons has also recently been studied by Bakucz Canário *et al.* [4], who derive an analytic solution to the massless Einstein–Dirac system, and illustrate its relevance to an observed zonal structure. A detailed review of this is presented in section 2.5. We note that this analytic solution (among others) was independently found by Blázquez-Salcedo & Knoll [54], whose work in addition generalises aspects of the Einstein–Dirac system to d dimensions (see also ref. [55]).

All of the above studies analyse Dirac solitons in a spherically symmetric context, a convenience that provides much simplification to the equations of motion. Spinning, axisymmetric solutions, however, have recently been presented by Herdeiro *et al.* [56], representing states comprising only a single fermion. Included in their work is a comparison between the spinning Dirac, boson and Proca cases, with an equivalent analysis concerning the spherical states presented in ref. [57]. The results of these comparative studies indicate that the three cases are qualitatively similar, but an important distinction is highlighted, in that a proper treatment of the Dirac problem demands the normalisation of the fermion wavefunction (a requirement unnecessary in the bosonic cases). Consequently, while in bosonic systems the boson mass is a free parameter and the boson number can vary considerably, in the Dirac system the fermion number is fixed (set equal to the number of spinor fields) and it is the fermion mass that varies. Thus the macroscopic stellar-sized

quantum states that exist in the context of boson stars, comprising a large number of light particles, cannot be as easily obtained in the Dirac case.

It is of course possible to study ‘classical’ Dirac solitons, in which the normalisation condition is left explicitly unsatisfied. Such an approach is employed in refs. [58–60], where the effects of coupling various fields to the Einstein–Dirac system are investigated, and also by Daka *et al.* [61], who consider the dynamical evolution of Dirac stars. This latter study utilises a framework similar to that employed by Seidel & Suen to analyse the bosonic case [33], and finds similar results (black hole formation and migration to a stable branch). It is unclear, however, whether this behaviour would persist if the normalisation condition were to be reinstated.

Finally, in relation to dynamics, it is important to note that, although these fermionic objects are often referred to as ‘solitons’, their solitonic nature has yet to be established, i.e. it is not known whether they collide in a non-diffractive manner. Head-on collisions of boson stars, however, have been performed, with solitonic behaviour reported for sufficiently large initial momenta [62, 63]. More relevant, perhaps, is the case of ℓ -boson stars [64], which are much more similar in structure to Dirac solitons. An analysis of the collisions between these objects, however, reports no indication of solitonic behaviour [65]. Regardless of whether Dirac solitons can truly be considered solitons, we shall continue to refer to them as such for the purposes of this thesis.

1.4 General relativity

We now proceed to the more technical aspects of this thesis, starting with a summary of some relevant concepts in Einstein’s theory of general relativity. Undoubtedly one of the finest achievements of 20th century physics, the theory has remained virtually unchanged since its formulation over 100 years ago, and still constitutes our best current understanding of the gravitational force. We shall not attempt to provide a thorough review of general relativity, as this can be found in any good textbook on the subject (Wald [66] for example, which uses many of the same conventions as we do), but instead focus on those areas most relevant to the study of Dirac solitons.

1.4.1 Notation & Conventions

We begin by stating the notation and conventions used within this thesis. Throughout we work in units where $\hbar = c = 1$, and additionally set the value of Newton’s constant G to unity when generating numerical solutions, although factors of G are retained in all written equations. The mostly-positive metric signature $(-, +, +, +)$ is used, noting that this is the opposite choice to that of Finster, Smoller & Yau in their analysis of Dirac solitons. Tensors with (lowercase) Greek indices run over both space and time, with

the Einstein summation convention assumed for repeated indices. On occasion, we shall use Latin indices to indicate sums that run over purely spatial dimensions. The explicit components of a tensor are labelled relative to the spatial co-ordinate system in question. For example, in spherical polar co-ordinates (r, θ, ϕ) the components of a space-time vector A^μ would be written as

$$A^\mu \equiv (A^t, A^r, A^\theta, A^\phi)^T. \quad (1.3)$$

We note of course that θ and ϕ are themselves Greek letters, but to avoid confusion we shall refrain from using these to label generic indices. Other conventions, involving for instance the definitions of the Riemann and Ricci tensors, are detailed in the appropriate sections that follow.

1.4.2 Spacetime metrics & curvature

The geometry of spacetime is that of a (3+1) dimensional pseudo-Riemannian manifold that is everywhere locally Minkowskian (i.e. flat) and endowed with a metric $g_{\mu\nu}$, which encodes the notion of distance via the infinitesimal line element:

$$ds^2 = g_{\mu\nu} dx^\mu dx^\nu. \quad (1.4)$$

Ultimately the metric is the fundamental object that is used to distinguish between different spacetimes. In flat space, for example, the metric reduces to $\eta_{\mu\nu} = \text{diag}(-1, 1, 1, 1)$, referred to as the Minkowski metric, and the line element can be written simply as $ds^2 = -dt^2 + d\vec{x}^2$.

The flat-space concept of a derivative is generalised to curved spacetime via the covariant derivative ∇_μ , which acts on a vector in the following manner:

$$\nabla_\nu A^\mu = \partial_\nu A^\mu + \Gamma_{\nu\sigma}^\mu A^\sigma. \quad (1.5)$$

There is a straightforward generalisation of this to tensors of arbitrary rank. Here, $\Gamma_{\nu\sigma}^\mu$ are the Christoffel symbols, which can be expressed in terms of the metric:

$$\Gamma_{\nu\sigma}^\mu = \frac{1}{2} g^{\mu\tau} (\partial_\nu g_{\sigma\tau} + \partial_\sigma g_{\tau\nu} - \partial_\tau g_{\nu\sigma}). \quad (1.6)$$

This relation uniquely determines the covariant derivative above (which should then strictly be referred to as the Levi-Civita connection), and ensures that the covariant derivative of the metric itself vanishes, i.e. $\nabla_\mu g_{\nu\sigma} = 0$.

The curvature of spacetime can be quantified using a variety of measures. Examples include the Riemann curvature tensor, $R^\mu{}_{\nu\sigma\tau}$, the Ricci curvature tensor, $R_{\mu\nu}$, and the Ricci scalar, R , which can be defined sequentially as follows:

$$R^\mu{}_{\nu\sigma\tau} = \partial_\sigma \Gamma_{\nu\tau}^\mu - \partial_\tau \Gamma_{\nu\sigma}^\mu + \Gamma_{\sigma\lambda}^\mu \Gamma_{\nu\tau}^\lambda - \Gamma_{\tau\lambda}^\mu \Gamma_{\nu\sigma}^\lambda; \quad (1.7)$$

$$R_{\mu\nu} = R^{\sigma}{}_{\mu\sigma\nu}; \quad (1.8)$$

$$R = R^{\mu}{}_{\mu}. \quad (1.9)$$

The Ricci scalar is the simplest curvature invariant that can be constructed in general relativity, and is often used as a convenient measure to assess the curvature of a particular spacetime and the presence of singularities within it (infinite curvature providing a decent, although not conclusive, indicator of an essential singularity). There exist spacetimes, however, in which the Ricci scalar will evaluate to zero despite the presence of non-zero (Riemann) curvature or indeed singularities. In many such cases, the Kretschmann scalar, defined as:

$$K = R_{\mu\nu\sigma\tau}R^{\mu\nu\sigma\tau}, \quad (1.10)$$

provides a more robust curvature measure. This quantity will prove particularly useful when assessing the singular solutions presented in Chapter 6.

1.4.3 The Einstein equations

Thus far we have dealt purely with the geometry of spacetime, but how do we incorporate matter into this picture? By analogy with special relativity, we can describe matter via an energy-momentum tensor, $T^{\mu\nu}$, which is both symmetric and has zero divergence with respect to the covariant derivative:

$$\nabla_{\mu}T^{\mu\nu} = 0. \quad (1.11)$$

This latter property ensures the local conservation of energy. The task is then to relate the energy-momentum tensor to some geometric quantity, in order to obtain an equation that governs precisely how the distribution of matter affects the curvature of spacetime. This quantity is required to be similarly divergenceless, leading Einstein to postulate the following expression, referred to as the Einstein equations:

$$R_{\mu\nu} - \frac{1}{2}Rg_{\mu\nu} = 8\pi GT_{\mu\nu}. \quad (1.12)$$

The terms on the left-hand side can be combined to form the Einstein tensor $G_{\mu\nu}$, and a cosmological constant term, $\Lambda g_{\mu\nu}$, can be optionally included, although this is not of relevance at the scales considered here. It is important to note that the Einstein equations are a postulate of general relativity, i.e. they are not derivable from fundamental principles. Indeed it is possible to construct alternative statements that nonetheless are consistent with current experimental constraints. Such alternatives are of particular interest with regard to the study of modified gravity theories that aim to address the dark matter and dark energy problems.

When studying field theory in the context of general relativity, it is convenient to utilise the Lagrangian formulation to allow the Einstein equations to be derived from a

variational principle. The starting point for this is the action:

$$S = \int \left(\frac{R}{16\pi G} + \mathcal{L}_m \right) \sqrt{-g} d^4x, \quad (1.13)$$

where g is the determinant of the metric, necessary to obtain the correct form for the volume element, and \mathcal{L}_m is the Lagrangian density for whichever matter fields are present. The first part of this expression is referred to as the Einstein-Hilbert action, chosen such that the variation of (1.13) with respect to the metric tensor produces the Einstein equations, with the energy-momentum tensor defined as:

$$T_{\mu\nu} = \frac{-2}{\sqrt{-g}} \frac{\delta}{\delta g^{\mu\nu}} (\sqrt{-g} \mathcal{L}_m). \quad (1.14)$$

1.4.4 The Schwarzschild metric

Although the Einstein equations can be written fairly concisely, they in fact consist of a set of (at most ten) coupled, non-linear differential equations, and hence constructing exact analytic solutions proves challenging. One obvious simplification is to consider spacetimes with well-defined symmetries, the existence of which can be expressed mathematically by the presence of one or more Killing vectors, ξ^μ , for which the following holds:

$$\nabla_\mu \xi_\nu + \nabla_\nu \xi_\mu = 0. \quad (1.15)$$

A spacetime is considered stationary if it possesses a timelike Killing vector, and considered static if this Killing vector is hypersurface-orthogonal, i.e. if the spacetime consists of a foliation of spacelike hypersurfaces defined by $t = \text{const}$.

The simplest non-trivial solution to the Einstein equations is the Schwarzschild metric, which represents the spacetime around a static, spherically symmetric point mass. In this case, the timelike Killing vector is δ_t^μ and that associated with spherical symmetry is δ_ϕ^μ . The Schwarzschild solution can be written in spherical co-ordinates as:

$$ds^2 = - \left(1 - \frac{2GM}{r} \right) dt^2 + \left(1 - \frac{2GM}{r} \right)^{-1} dr^2 + r^2 d\Omega^2, \quad (1.16)$$

where $d\Omega^2 = d\theta^2 + r^2 \sin^2 \theta d\phi^2$, and M is the total gravitational mass of the point particle (see section 1.4.5). Note that the metric contains an essential singularity at $r = 0$ and a black hole horizon at $r = 2GM$. This solution is of relevance to the study of Dirac solitons, in that the notion of ‘particle-like’ is understood to imply that the spacetime approaches the Schwarzschild form at a sufficiently large distance from the fermion source.

A generalisation of the Schwarzschild metric to that of a static, spherical object of

constant density and radius R is given by the interior Schwarzschild solution:

$$ds^2 = -\frac{1}{4} \left[3 \left(1 - \frac{2GM}{R} \right)^{1/2} - \left(1 - \frac{r^2}{a^2} \right)^{1/2} \right]^2 dt^2 + \frac{dr^2}{1 - (r/a)^2} + r^2 d\Omega^2, \quad (1.17)$$

where $a = (R^3/2GM)^{1/2}$. This is valid only within the object ($r \leq R$), and matches on to the exterior Schwarzschild metric (1.16) for $r > R$. Note that this spacetime contains no central singularity and no black hole horizon provided $R > 2GM$.

1.4.5 Mass in general relativity

The concepts of mass and energy in general relativity are not in general simple to define, in essence due to the lack of a consistent, frame-independent definition for the energy of the gravitational field (see chapter 11 of Wald [66] for a detailed discussion). For spacetimes that are asymptotically flat, however, some progress can be made. For example, in such cases the total Arnowitt–Deser–Misner (ADM) mass can be obtained by considering the deviation of the metric from that of Minkowski space as $r \rightarrow \infty$. Applying this to the Schwarzschild metric (1.16), one finds that the quantity M is in fact the total ADM mass of the system.

This definition, however, is only valid strictly at spatial infinity. How then can one quantify the concept of mass/energy contained within an arbitrary region? For spacetimes that are static and spherically symmetric, in addition to being asymptotically flat, an approximate answer is provided by the Komar mass. This quantity is calculated by considering the gravitational force (applied from spatial infinity) required to hold stationary a test particle, and takes the following surface integral form [66]:

$$M_K = -\frac{1}{4\pi} \int_S \frac{1}{\chi} n^\mu \xi^\nu (\nabla_\nu \xi_\mu) dS, \quad (1.18)$$

where ξ^μ is a timelike Killing vector normalised such that $\chi = \sqrt{-\xi_\mu \xi^\mu} \rightarrow 1$ as $r \rightarrow \infty$, and n^μ is the outward-pointing unit normal to the 2-sphere S . Using Stokes's theorem and applying the Einstein equations, this can be converted into a volume integral form:

$$M_K = 2 \int_V \left(T_{\mu\nu} - \frac{1}{2} T^\sigma{}_\sigma g_{\mu\nu} \right) n^\mu \xi^\nu dV. \quad (1.19)$$

We shall make use of these expressions when analysing the interior mass distributions of Dirac solitons. Note that the Komar mass evaluated at infinity coincides with the ADM mass, as can be easily verified for the Schwarzschild metric (1.16).

1.5 Spinors in curved spacetime

We now discuss the second background topic required for a technical understanding of Dirac solitons, namely the treatment of fermions within the framework of general relativity. Again we shall only give a brief summary of the required concepts. For a more detailed discussion see for example refs. [67] and [68].

1.5.1 The Dirac equation in flat spacetime

We begin with a review of the situation in flat (Minkowski) spacetime. Unlike spinless bosonic particles, which can be adequately described by a scalar field, the treatment of particles with non-zero spin necessitates the introduction of objects endowed with additional structure. In the case of massive spin- $\frac{1}{2}$ fermions, the required objects are four-component Dirac spinors $\Psi(x)$, which are contained within the fundamental representation of the Lorentz group. For a Dirac fermion of mass m , the spinor wavefunction obeys the (flat-space) Dirac equation:

$$(i\bar{\gamma}^\mu \partial_\mu - m) \Psi = 0, \quad (1.20)$$

where $\bar{\gamma}^\mu$ are the flat-space Dirac gamma matrices, defined by their anti-commutation relations:

$$\{\bar{\gamma}^\mu, \bar{\gamma}^\nu\} \equiv \bar{\gamma}^\mu \bar{\gamma}^\nu + \bar{\gamma}^\nu \bar{\gamma}^\mu = -2\eta^{\mu\nu}, \quad (1.21)$$

where $\eta_{\mu\nu}$ is the Minkowski metric. In the Dirac representation, these 4x4 matrices take the following explicit forms:

$$\bar{\gamma}^0 = \begin{pmatrix} \mathbf{1} & 0 \\ 0 & -\mathbf{1} \end{pmatrix}; \quad \bar{\gamma}^i = \begin{pmatrix} 0 & \sigma^i \\ -\sigma^i & 0 \end{pmatrix}, \quad (1.22)$$

where $i \in \{1, 2, 3\}$ and σ^i are the standard Pauli matrices. Note that we are choosing to label the flat-space gamma matrices with a additional bar in order to distinguish them from their curved-space counterparts.

Note also that, although the spinor Ψ is written without any explicit spacetime indices, this does not imply that it transforms as a scalar under Lorentz boosts $\Lambda^\mu{}_\nu$. Instead its explicit transformation can be written as:

$$\Psi \rightarrow \exp\left(-\frac{i}{2}\omega_{\mu\nu}\sigma^{\mu\nu}\right) \Psi \equiv S\Psi, \quad (1.23)$$

where $\sigma^{\mu\nu} = -\frac{i}{4}[\bar{\gamma}^\mu, \bar{\gamma}^\nu]$ and $\omega_{\mu\nu}$ are the generators of the infinitesimal Lorentz transform:

$$\Lambda^\mu{}_\nu \simeq \delta^\mu{}_\nu + \omega^\mu{}_\nu. \quad (1.24)$$

Using this transformation property, it is then straightforward to show that the Dirac equation is indeed Lorentz invariant, as required.

It is also possible to define an action from which the Dirac equation can be obtained via a variational principle. This can be written as:

$$S_D = \int \bar{\Psi} (i\bar{\gamma}^\mu \partial_\mu - m) \Psi d^4x, \quad (1.25)$$

where the adjoint spinor $\bar{\Psi} \equiv \Psi^\dagger \bar{\gamma}^0$. Varying this action with respect to $\bar{\Psi}$ then results in the Dirac equation, the conjugate of which is obtained by varying with respect to the spinor itself.

1.5.2 The vierbein formalism

How should one generalise the above formalism to curved spacetime, and thus general relativity? The first obvious step is to define curved-space Dirac gamma matrices that obey the generalised anti-commutation relations:

$$\{\gamma^\mu, \gamma^\nu\} = -2g^{\mu\nu}. \quad (1.26)$$

One could then imagine obtaining the curved-space Dirac equation by simply replacing the flat-space gamma matrices with those above and promoting the partial derivative to a covariant derivative. This naive approach fails, however, since the resulting equation will not be invariant under general co-ordinate transforms, ultimately due to the Lorentz transformation properties of the Dirac spinor. One way to proceed is to utilise the principle of equivalence, i.e. that at each point in the spacetime manifold, one can construct a flat, locally inertial reference frame. Since we know how to deal with spinors in flat spacetime, all that remains is to relate quantities written in the local inertial co-ordinate basis to those in the general non-inertial basis. The vierbein (or tetrad) formalism achieves precisely this.

More formally, if we construct a set of local inertial co-ordinates X_p^a at a point p in the space-time manifold, then the vierbein e^a_μ at that point is defined as:

$$e^a_\mu(p) = \left(\frac{\partial X_p^a}{\partial x^\mu} \right)_{x=p}. \quad (1.27)$$

This object should be interpreted not as a tensor, since it does not transform as one, but instead as a set of four vector fields, with the first (Latin) index referring to the local inertial co-ordinates and the second (Greek) index to the non-inertial co-ordinates. We now have a means by which we can transform between quantities written in the two basis systems, e.g. for a vector field:

$$A^a = e^a_\mu A^\mu; \quad A^\mu = e_a^\mu A^a. \quad (1.28)$$

The generalisation to tensors of arbitrary rank follows straightforwardly. In particular, we

can relate the metric tensor to the Minkowski metric via the following expression:

$$g_{\mu\nu} = e^a{}_{\mu} e^b{}_{\nu} \eta_{ab}. \quad (1.29)$$

Note that Latin indices are raised and lowered by the Minkowski metric, and Greek indices by the full metric.

1.5.3 The Dirac equation in curved spacetime

How do we now apply the vierbein formalism to the Dirac equation? First, we can relate the curved-space gamma matrices to their flat-space counterparts:

$$\gamma^{\mu} = e_a{}^{\mu} \bar{\gamma}^a. \quad (1.30)$$

One can easily check that the anti-commutation relations (1.21) follow as a consequence of this definition. We then must address the problem of the derivative term. To this end, we write the curved-space Dirac equation as:

$$(i\gamma^{\mu} D_{\mu} - m)\Psi = 0, \quad (1.31)$$

where we have introduced a corrected derivative $D_{\mu} = \partial_{\mu} + \Gamma_{\mu}$. The explicit form of the correction term Γ_{μ} can be determined by requiring that (1.31) be invariant under general co-ordinate transforms. For this we follow the derivation found in ref. [67]. The vierbein formalism allows us to work solely in a local inertial frame, and thus the condition of covariance requires $e_a{}^{\mu} D_{\mu} \Psi \rightarrow S \Lambda_a{}^b e_b{}^{\mu} D_{\mu} \Psi$ under a (position-dependent) Lorentz transform $\Lambda^a{}_b$, with S defined as in (1.23). This implies that Γ_{μ} must transform in the following manner:

$$\Gamma_{\mu} \rightarrow S \Gamma_{\mu} S^{-1} - (\partial_{\mu} S) S^{-1}. \quad (1.32)$$

For an infinitesimal Lorentz transform of the form (1.24), S and S^{-1} can be written as:

$$S \simeq 1 - \frac{i}{2} \omega_{ab} \sigma^{ab}; \quad S^{-1} \simeq 1 + \frac{i}{2} \omega_{ab} \sigma^{ab}, \quad (1.33)$$

noting that ω_{ab} is now position-dependent. Substituting these into (1.32) then gives:

$$\Gamma_{\mu} \rightarrow \Gamma_{\mu} - \frac{i}{2} \omega_{ab} [\sigma^{ab}, \Gamma_{\mu}] + \frac{i}{2} (\partial_{\mu} \omega_{ab}) \sigma^{ab}. \quad (1.34)$$

This defines the Lorentz transformation property required to ensure that the Dirac equation remains invariant under general co-ordinate transforms. To determine the explicit form of Γ_{μ} , note first that the vierbein transforms like a vector in the inertial frame:

$$e^a{}_{\mu} \rightarrow \Lambda^a{}_b e^b{}_{\mu} \simeq e^a{}_{\mu} + \omega^a{}_b e^b{}_{\mu}. \quad (1.35)$$

One can then show that the quantity $e_b{}^\nu \partial_\mu e_{a\nu}$ transforms as follows:

$$e_b{}^\nu \partial_\mu e_{a\nu} \rightarrow e_b{}^\nu \partial_\mu e_{a\nu} + \omega_b{}^c e_c{}^\nu \partial_\mu e_{a\nu} + \omega_a{}^c e_b{}^\nu \partial_\mu e_{c\nu} + \partial_\mu \omega_{ab}. \quad (1.36)$$

Contracting this with σ^{ab} and multiplying by $i/2$, one finds that this is precisely the transformation required for Γ_μ given in (1.34). Thus the appropriate derivative that acts on a spinor field in curved spacetime is:

$$D = \partial_\mu + \frac{i}{2} \sigma^{ab} e_a{}^\nu \partial_\mu e_{b\nu}. \quad (1.37)$$

Chapter 2

Einstein–Dirac solitons

We now review in detail the Einstein–Dirac system and the particle-like states that exist within it. This comprises a derivation of the equations of motion, details on the numerical method by which these can be solved, and a summary of the main results found for the two-fermion system studied by Finster, Smoller & Yau [3]. In addition, we shall review the more recent analysis of Bakucz Canário *et al.* [4] concerning the internal structure of Dirac solitons, as well as calculating a variety of quantities used throughout the thesis.

2.1 Equations of motion

2.1.1 The two-fermion system

We begin by outlining the derivation of the equations of motion valid for the two-fermion Einstein–Dirac system following ref. [3], although we note that expressions may differ due to the difference in sign conventions. The starting point is the Einstein–Dirac action, defined as:

$$S_{\text{ED}} = \int \left(\frac{R}{16\pi G} + \bar{\Psi} (\not{D} - m) \Psi \right) \sqrt{-g} d^4x, \quad (2.1)$$

where we have made use of the Feynman slash notation to write $\not{D} \equiv i\gamma^\mu D_\mu$. Variation of this action with respect to the adjoint spinor $\bar{\Psi}$ produces the curved-space Dirac equation (1.31), while variation with respect to the metric results in the Einstein equations (1.12). In this case, the matter Lagrangian density from which the energy-momentum tensor can be obtained is simply:

$$\mathcal{L}_m = \bar{\Psi} (\not{D} - m) \Psi. \quad (2.2)$$

Note that together the Einstein and Dirac equations form a fully coupled system, since the Dirac operator \not{D} depends on the curved-space gamma matrices (and the vierbein), which in turn depend on the metric.

To solve such a coupled system is evidently not straightforward, and thus to make

progress we simplify the problem by seeking static, spherically symmetric solutions. In spherical polar co-ordinates (t, r, θ, ϕ) , the most general form for the spacetime metric can be written as:

$$ds^2 = -\frac{1}{T(r)^2}dt^2 + \frac{1}{A(r)}dr^2 + r^2d\Omega^2, \quad (2.3)$$

where the metric functions $T(r)$ and $A(r)$ depend only on the radial co-ordinate. Note that the condition $A(r) > 0$ is required in order to preserve the metric signature. This choice of metric implies that a vierbein may be constructed in which the only non-zero components are $e^t_t = T$, $e^r_r = \sqrt{A}$ and $e^\theta_\theta = e^\phi_\phi = 1$. Thus the curved-space gamma matrices can be written as:

$$\gamma^0 = T \begin{pmatrix} \mathbf{1} & 0 \\ 0 & -\mathbf{1} \end{pmatrix}; \quad \gamma^r = \sqrt{A} \begin{pmatrix} 0 & \sigma^r \\ -\sigma^r & 0 \end{pmatrix}; \quad (2.4)$$

$$\gamma^\theta = \frac{1}{r} \begin{pmatrix} 0 & \sigma^\theta \\ -\sigma^\theta & 0 \end{pmatrix}; \quad \gamma^\phi = \frac{1}{r} \begin{pmatrix} 0 & \sigma^\phi \\ -\sigma^\phi & 0 \end{pmatrix}, \quad (2.5)$$

where σ^r , σ^θ and σ^ϕ are the Pauli matrices in spherical co-ordinates, defined in terms of their usual Cartesian counterparts by:

$$\sigma^r = \sigma^1 \sin \theta \cos \phi + \sigma^2 \sin \theta \sin \phi + \sigma^3 \cos \theta; \quad (2.6)$$

$$\sigma^\theta = \sigma^1 \cos \theta \cos \phi + \sigma^2 \cos \theta \sin \phi - \sigma^3 \sin \theta; \quad (2.7)$$

$$\sigma^\phi = \frac{1}{\sin \theta} (-\sigma^1 \sin \phi + \sigma^2 \cos \phi). \quad (2.8)$$

Imposing spherical symmetry on the fermionic sector, however, is less straightforward, since in general the presence of spin will identify a preferred direction in the spacetime. The simplest solution to this problem is to consider the case of two fermions arranged in a singlet state, where, despite each individual fermion having non-zero spin, the total angular momentum of the overall state is identically zero. This is mathematically encoded by introducing the following ansatz for the spinor wavefunction:

$$\Psi_a = \frac{\sqrt{T(r)}}{r} \begin{pmatrix} \alpha(r)e_a \\ -i\sigma^r\beta(r)e_a \end{pmatrix} e^{-i\omega t}, \quad (2.9)$$

where $a \in \{1, 2\}$ labels each fermion, and e_a are two-component basis vectors with $e_1 = (1, 0)^T$ and $e_2 = (0, 1)^T$. Both fermions are chosen to have the same radial dependence, controlled by the (real-valued) functions $\alpha(r)$ and $\beta(r)$, and in addition the same mass m and energy ω , with the latter controlling the harmonic time-dependence as standard for a stationary state. We shall discuss the derivation of this ansatz in more detail when considering the many-fermion system in section 2.1.2. If required, the total wavefunction of the singlet state can be constructed using the Hartree–Fock formalism:

$$\Psi = \Psi_1 \wedge \Psi_2, \quad (2.10)$$

although this proves unnecessary in deriving the equations of motion.

The next task is to determine an explicit form for the Dirac operator \mathcal{D} . Although it is possible to use the vierbein formalism and the expressions given in section 1.5, it is more convenient to employ the approach formulated in ref. [69], using which the following form for the Dirac operator can be obtained:

$$\mathcal{D} = i\gamma^\mu \partial_\mu + \frac{i}{2} \nabla_\mu \gamma^\mu \quad (2.11)$$

$$= i\gamma^t \frac{\partial}{\partial t} + i\gamma^r \left(\frac{\partial}{\partial r} + \frac{1}{r} \left(1 - \frac{1}{\sqrt{A}} \right) - \frac{T'}{2T} \right) + i\gamma^\theta \frac{\partial}{\partial \theta} + i\gamma^\phi \frac{\partial}{\partial \phi}, \quad (2.12)$$

where in the second line the covariant derivatives of the gamma matrices have been explicitly evaluated, and the prime denotes a radial derivative. Using this, along with the spinor ansatz (2.9), it is straightforward to evaluate the Dirac equation (1.31), which ultimately reduces to the following two independent equations:

$$\sqrt{A} \alpha' = +\frac{\alpha}{r} - (\omega T + m)\beta; \quad (2.13)$$

$$\sqrt{A} \beta' = -\frac{\beta}{r} + (\omega T - m)\alpha. \quad (2.14)$$

Turning to the Einstein equations, the energy-momentum tensor $T_{\mu\nu}$ can be calculated using (1.14). It can be shown that the variation of the second term in (2.11) vanishes for a singlet state, leaving only the first term to consider. This results in the following expression for the mixed energy-momentum tensor:

$$T^\mu{}_\nu = - \sum_{a=1}^2 \Re \{ \bar{\Psi}_a (i\gamma^\mu \partial_\nu) \Psi_a \}, \quad (2.15)$$

the only non-zero components of which are:

$$T^t{}_t = -\frac{2\omega T^2}{r^2} (\alpha^2 + \beta^2); \quad (2.16)$$

$$T^r{}_r = \frac{2T\sqrt{A}}{r^2} (\alpha\beta' - \beta\alpha'); \quad (2.17)$$

$$T^\theta{}_\theta = T^\phi{}_\phi = \frac{2T}{r^3} \alpha\beta. \quad (2.18)$$

The Einstein tensor $G_{\mu\nu}$ can be evaluated by using the expressions given in section 1.4.2, from which we find the non-zero components to be:

$$G^t{}_t = \frac{1}{r^2} (-1 + A + rA'); \quad (2.19)$$

$$G^r{}_r = \frac{1}{r^2} \left(-1 + A - \frac{2rAT'}{T} \right); \quad (2.20)$$

$$G^\theta{}_\theta = G^\phi{}_\phi = \frac{A'}{2r} - \frac{A'T'}{2T} + \frac{2A(T')^2}{T^2} - \frac{AT'}{rT} - \frac{AT''}{T}. \quad (2.21)$$

In principle, four Einstein equations can therefore be constructed, although it is clear that spherical symmetry will guarantee that the $\theta\theta$ and $\phi\phi$ equations are equal, as is evidenced above. In addition, since the Einstein equations are divergenceless ($\nabla_\mu T^\mu{}_\nu = \nabla_\mu G^\mu{}_\nu = 0$), one of the remaining three can always be written in terms of the other two. This leaves only two independent Einstein equations, which we choose for convenience to be:

$$-1 + A + rA' = -16\pi G\omega T^2 (\alpha^2 + \beta^2); \quad (2.22)$$

$$-1 + A - 2rA\frac{T'}{T} = 16\pi GT\sqrt{A} (\alpha\beta' - \alpha'\beta). \quad (2.23)$$

These, along with the Dirac equations (2.13) and (2.14) form a set of four coupled differential equations for the four unknown fields α , β , A and T . Solving this system fully defines the wavefunction and spacetime for a pair of neutral, gravitationally-interacting fermions, within the context of the Einstein–Dirac formalism.

2.1.2 The many-fermion system

Although initially valid only for systems of two fermions, the above derivation was subsequently extended to include states consisting of an arbitrarily large (even) number of fermions [43]. To achieve this, the fermions must be arranged in a filled shell, somewhat akin to an atomic orbital, in order to guarantee that the total angular momentum remains zero and the overall state is still spherically symmetric. More formally, a collection of fermions, each with angular momentum $j \in \{\frac{1}{2}, \frac{3}{2}, \frac{5}{2}, \dots\}$, will constitute a filled shell if there exists exactly one fermion for every possible projection of the angular momentum in the z -direction (the value of which is denoted k). Elementary quantum mechanics dictates that k can take only quantised values within the set $\{-j, -j+1, \dots, j-1, j\}$, and thus the total number of fermions in a filled shell is $N_f = 2j+1$, a quantity that can take any positive even value. Labelling the wavefunction of each individual fermion by its quantum numbers $\{j, k\}$, the total wavefunction of the state can as before be constructed via the Hartree-Fock product:

$$\Psi = \Psi_{j,k=-j} \wedge \Psi_{j,k=-j+1} \wedge \dots \wedge \Psi_{j,k=+j-1} \wedge \Psi_{j,k=j}. \quad (2.24)$$

This overall state is guaranteed to have zero total angular momentum, and is therefore spherically symmetric.

Encoding the notion of a filled shell of fermions into the Einstein–Dirac system requires an amendment to the spinor ansatz (2.9). The precise form of this can be obtained by requiring each spinor wavefunction Ψ_{jk} to be an eigenstate of the total angular momentum operator J , with eigenvalue $\sqrt{j(j+1)}$, and the z -projected angular momentum operator J_z , with eigenvalue k . In addition, to include fermions with both positive and negative parity, each wavefunction must be an eigenstate of the parity operator P , with eigenvalue ± 1 as appropriate. This calculation is similar to that implemented for the Hydrogen atom

(see [70] for example), and the resulting spinor ansatz is

$$\Psi_{jk}^{\pm}(t, r, \theta, \phi) = \frac{\sqrt{T(r)}}{r} \begin{pmatrix} \chi_{j\mp\frac{1}{2}}^k \alpha(r) \\ i\chi_{j\pm\frac{1}{2}}^k \beta(r) \end{pmatrix} e^{-i\omega t}, \quad (2.25)$$

where \pm indicates positive or negative parity (noting that all fermions in the filled shell are required to have the same parity). As in the two-fermion case, there is a harmonic time dependence, while the radial dependence is controlled by the real-valued fermion fields α and β . The angular part of the wavefunction differs for each fermion, and is determined by the two-component functions χ :

$$\chi_{j-\frac{1}{2}}^k(\theta, \phi) = \sqrt{\frac{j+k}{2j}} Y_{j-\frac{1}{2}}^{k-\frac{1}{2}}(\theta, \phi) \begin{pmatrix} 1 \\ 0 \end{pmatrix} + \sqrt{\frac{j-k}{2j}} Y_{j-\frac{1}{2}}^{k+\frac{1}{2}}(\theta, \phi) \begin{pmatrix} 0 \\ 1 \end{pmatrix}; \quad (2.26)$$

$$\chi_{j+\frac{1}{2}}^k(\theta, \phi) = \sqrt{\frac{j+1-k}{2j+2}} Y_{j+\frac{1}{2}}^{k-\frac{1}{2}}(\theta, \phi) \begin{pmatrix} 1 \\ 0 \end{pmatrix} - \sqrt{\frac{j+1+k}{2j+2}} Y_{j+\frac{1}{2}}^{k+\frac{1}{2}}(\theta, \phi) \begin{pmatrix} 0 \\ 1 \end{pmatrix}, \quad (2.27)$$

where $Y_m^l(\theta, \phi)$ are the usual spherical harmonics. A detailed calculation extending this ansatz to d dimensions can be found in ref. [54]. Note that the factor of \sqrt{T}/r appearing in (2.25) cannot be derived by such a calculation, and is added purely in order to simplify the equations of motion. Note also that the two-fermion ansatz (2.9) can be recovered by considering the case of $j = \frac{1}{2}$.

Using this amended ansatz, the derivation of the many-fermion equations of motion follows in a similar manner to the two-fermion case, although significantly more care is required when combining the angular parts of the spinors (see ref. [43] for details). As before, each fermion obeys an identical Dirac equation and has an equal contribution to the energy-momentum tensor, and thus it transpires that the equations of motion need only a slight modification:

$$\sqrt{A} \alpha' = +\frac{\kappa\alpha}{2r} - (\omega T + m)\beta; \quad (2.28)$$

$$\sqrt{A} \beta' = -\frac{\kappa\beta}{2r} + (\omega T - m)\alpha; \quad (2.29)$$

$$-1 + A + rA' = -8\pi G|\kappa|\omega T^2 (\alpha^2 + \beta^2); \quad (2.30)$$

$$-1 + A - 2rA\frac{T'}{T} = 8\pi G|\kappa|T\sqrt{A} (\alpha\beta' - \alpha'\beta), \quad (2.31)$$

where $\kappa = \pm(2j+1)$, the sign of which indicates positive or negative parity. This set of equations, henceforth referred to as the ‘FSY equations’, fully describes the behaviour of a static, self-gravitating filled shell consisting of $N_f = |\kappa|$ fermions of equal mass m and energy ω . Despite being derived over 20 years ago, there exist no published results presenting particle-like solutions of this system, although it is remarked in ref. [43] that such states exist (within the context of charged fermions) and exhibit the same qualitative

behaviour as those in the two-fermion system. Here, in Chapters 3 and 4, we shall provide a thorough analysis of these many-fermion states, and demonstrate that this is in fact untrue; for sufficiently large numbers of fermions, the behaviour of the states differs significantly from the two-fermion case.

Finally, one might be concerned as to the physical relevance of an isolated filled shell of high-angular-momentum fermions. Certainly a more realistic situation, and one more in keeping with the true spirit of the Pauli exclusion principle, would be to consider a series of nested filled shells, the first of which contains two fermions, with each subsequent shell adding a further two. Although such a construction can indeed be treated within a spherically symmetric framework, the resulting set of equations is significantly more challenging to solve, since each shell will obey its own distinct Dirac equation with distinct energy ω . Thus, in order to study systems containing large numbers of fermions (the cumulative gravitational impact of which is of interest), the only computationally-feasible choice is to work with the admittedly somewhat unrealistic concept of an isolated filled shell.

2.2 Generating particle-like solutions

We now address the issue of constructing localised, particle-like solutions to the FSY equations (2.28)–(2.31), again following the method outlined in ref. [3]. These are the states that we refer to as Einstein–Dirac solitons.

2.2.1 Boundary conditions

To solve the equations of motion we are of course required to specify boundary conditions for the fields α , β , A and T . Since Dirac solitons are localised objects, the spacetime in which they are contained must be asymptotically flat, i.e.

$$\lim_{r \rightarrow \infty} T(r) = 1; \quad \lim_{r \rightarrow \infty} A(r) = 1. \quad (2.32)$$

In addition, the notion of particle-like implies that the fermion fields α and β must decay sufficiently quickly at large r in order for the state to be normalisable. This property is to be understood in relation to the norm of a state, which is defined as:

$$(\Psi|\Psi) = \int_{\mathcal{H}} \bar{\Psi} \gamma^\mu \Psi n_\mu \sqrt{-h} \, d^3x, \quad (2.33)$$

where \mathcal{H} is a space-like hypersurface with future-directed normal n_μ , and h is the determinant of the spatial part of the metric (2.3). This is simply the conserved charge associated with the usual conserved Dirac current $\bar{\Psi} \gamma^\mu \Psi$, and thus can be interpreted as defining the particle number. A quantum treatment of fermions strictly requires the normalisation of the spinor wavefunction, i.e. $(\Psi_{jk}|\Psi_{jk}) = 1$, such that each spinor field Ψ_{jk} is associated

with precisely one fermion. Using the ansatz (2.25) and taking $n_\mu = (T^{-1}, 0, 0, 0)$, this normalisation condition can be written explicitly as:

$$4\pi \int_0^\infty \frac{T}{\sqrt{A}} (\alpha^2 + \beta^2) dr = 1, \quad (2.34)$$

the integrand of which can be interpreted as the probability of finding a fermion in an infinitesimal region dr .

Both the conditions of asymptotic flatness and normalisation can be deduced from the physical characteristics of a particle-like quantum state. In contrast, the remaining boundary conditions, which are specified at $r = 0$, are obtained solely by ensuring that the state contains no central singularity, a property that is certainly desirable but not necessarily required. This regularity is realised by the following small- r asymptotic expansion (valid only for positive κ):

$$\alpha(r) = \alpha_1 r^{\kappa/2} + \mathcal{O}(r^{\kappa/2+2}); \quad (2.35)$$

$$\beta(r) = \frac{1}{\kappa+1} \alpha_1 (\omega T_0 - m) r^{\kappa/2+1} + \mathcal{O}(r^{\kappa/2+3}); \quad (2.36)$$

$$A(r) = 1 - \frac{8\pi G \kappa}{\kappa+1} \omega T_0 \alpha_1^2 r^\kappa + \mathcal{O}(r^{\kappa+2}); \quad (2.37)$$

$$T(r) = T_0 - \frac{8\pi G}{2\kappa+2} (2\omega T_0 - m) T_0^2 \alpha_1^2 r^\kappa + \mathcal{O}(r^{\kappa+2}), \quad (2.38)$$

where α_1 and T_0 are unconstrained parameters. We note that this is by no means the unique small- r expansion for the equations of motion; indeed it is not even the unique non-singular expansion. It is, however, the unique non-singular expansion from which normalisable solutions can be constructed. We shall discuss this further in Chapter 6.

As mentioned, the above asymptotic expansion is valid only for positive values of κ . For negative parity states, the expansion is altered such that β becomes the dominant field at small r . Note, however, that the FSY equations are symmetric under the transformation $\{\kappa \rightarrow -\kappa, m \rightarrow -m, \alpha \rightarrow \beta, \beta \rightarrow -\alpha\}$. One can therefore obtain negative parity solutions by solving the system for negative values of m (with the expansion above) and subsequently interchanging the fermion fields. Moreover, the change in relative sign of the fermion fields has no effect on the properties of the solution, and thus anti-fermion states with positive parity are equivalent to fermion states with negative parity. Similarly, anti-fermion states with negative parity are equivalent to fermion states with positive parity. Throughout the majority of this thesis we shall consider only positive parity, positive mass states, with the exception being Chapter 6.

2.2.2 The rescaling method

For a specific choice of κ , the FSY equations (2.28)–(2.31) seemingly contain two free parameters — the fermion mass m and the fermion energy ω . The small- r expansion

(2.35)–(2.38) adds a further two unspecified parameters to the system in α_1 and T_0 . The values of three of these, however, will be constrained by satisfying the conditions of normalisation and asymptotic flatness. Hence one would expect the particle-like solutions of the FSY equations to constitute a one-parameter family of states, which indeed turns out to be the case.

From a numerical point of view, imposing asymptotic flatness and normalisation is not straightforward, however. Fortunately, a rescaling procedure exists through which these conditions can be converted into constraints applicable at $r = 0$, greatly simplifying the process of obtaining solutions. To understand this, first note that normalisable solutions to the FSY equations automatically satisfy the condition $A \rightarrow 1$ as $r \rightarrow \infty$, but imply only the weaker condition $T \rightarrow \text{const}$. To remedy this latter discrepancy, a value for T_0 must be chosen such that this constant is equal to one. This can be achieved by first solving the system with an arbitrary specified value of T_0 , after which the solution can be subsequently rescaled to ensure $T \rightarrow 1$ as required. A similar rescaling trick can be used to convert a solution with an arbitrary value of the normalisation integral to one which is correctly normalised. More formally, we can seek normalisable solutions for the ‘unscaled’ fields $\{\tilde{\alpha}, \tilde{\beta}, \tilde{A}, \tilde{T}\}$ that satisfy $\tilde{T}_0 = \tilde{m} = 1$, along with the weaker boundary conditions:

$$\lim_{r \rightarrow \infty} \tilde{T}(r) = \tau; \quad (2.39)$$

$$\int_0^\infty \frac{\tilde{T}}{\sqrt{\tilde{A}}} (\tilde{\alpha}^2 + \tilde{\beta}^2) dr = \chi^2. \quad (2.40)$$

Once obtained, these unscaled solutions can then be normalised and made asymptotically flat by rescaling the fields and parameters as follows:

$$\begin{aligned} \alpha(r) &= \sqrt{\frac{\tau}{\chi}} \tilde{\alpha}(\chi r); & A(r) &= \tilde{A}(\chi r); & \omega &= \chi \tau \tilde{\omega}; \\ \beta(r) &= \sqrt{\frac{\tau}{\chi}} \tilde{\beta}(\chi r); & T(r) &= \frac{1}{\tau} \tilde{T}(\chi r); & m &= \chi \tilde{m}. \end{aligned} \quad (2.41)$$

We note of course that the unscaled solutions (denoted by an additional tilde) do not represent physically acceptable states, and are merely a convenient tool to simplify the numerical solving procedure. On occasion, however, it can prove enlightening to analyse the behaviour of states in terms of their unscaled parameters, as we shall demonstrate in Chapter 4.

The above rescaling procedure assigns values to two of the four previously unspecified parameters in the unscaled system (\tilde{T}_0 and \tilde{m}), leaving only $\tilde{\alpha}_1$ and $\tilde{\omega}$ to be considered. It turns out that $\tilde{\omega}$ must be tuned in order to obtain normalisable solutions (i.e. those in which the fermion fields decay at large r), while $\tilde{\alpha}_1$ is a truly free parameter that can take any positive value, and thus parametrises the families of particle-like states. Note, however, that the physical (rescaled) quantity α_1 cannot be used in a similar manner. Instead it is

T_0 that acquires the role of a free parameter, and hence the equivalent physical quantity that should be used to identify states is the central redshift z , defined as:

$$z = T_0 - 1, \quad (2.42)$$

which can also take any positive value. This can be used as a measure of the central compactness of a state, or alternatively how relativistic a state is, with $z \approx 1$ marking the boundary between non-relativistic (low-redshift) and relativistic (high-redshift) states. In the Einstein–Dirac system, z and $\tilde{\alpha}_1$ are guaranteed to be in one-to-one correspondence, and thus they can be used interchangeably. This does not apply, however, to the Einstein–Dirac–Higgs system studied in Chapter 5.

As a final remark, we note that the classical analysis of Dirac solitons does not require the normalisation of the spinor wavefunction, and therefore in such a treatment the fermion mass becomes an additional free parameter in the system. The number of fermions is then determined by evaluating the norm $(\Psi|\Psi)$, which in the classical case is no longer required to equal one. Hence the possibility of obtaining highly compact, macroscopic states that contain large numbers of light fermions can be realised. In contrast, in the quantum treatment, the fermion mass is not a free parameter, and by dimensional analysis is required to be of the order of the Planck mass. Thus quantum Einstein–Dirac solitons should be thought of as microscopic, Planck-scale states that contain a relatively small number of ultra-high-mass fermions.

2.3 Useful quantities

Before presenting numerical examples of Dirac solitons, we first derive expressions for a variety of quantities that will prove useful in studying their properties.

2.3.1 Fermion densities

The first of these is the fermion number density $n_f(r)$, defined as the number of fermions per unit volume, the form of which can be deduced from the normalisation condition $(\Psi_{jk}|\Psi_{jk}) = 1$. Using the expression for the norm in (2.33) and evaluating the integrand, we obtain the following expression for the total number of fermions $N_f = |\kappa|$:

$$N_f = N_f (\Psi_{jk}|\Psi_{jk}) = N_f \int_{\mathcal{H}} \frac{T}{r^2} (\alpha^2 + \beta^2) \sqrt{-h} \, d^3x. \quad (2.43)$$

Identifying the spatial volume element $dV = \sqrt{-h} \, d^3x$, the fermion number density can therefore be written as:

$$n_f(r) = N_f \frac{T}{r^2} (\alpha^2 + \beta^2). \quad (2.44)$$

One might then be tempted to define the fermion energy density as simply $\omega \times n_f$, but this turns out to be incorrect. Instead, the energy density $\rho_f(r)$ is obtained from the tt component of the energy-momentum tensor, and takes the form:

$$\rho_f(r) = \omega N_f \frac{T^2}{r^2} (\alpha^2 + \beta^2). \quad (2.45)$$

Comparing this with (2.44) implies that each fermion has a ‘local’ energy ωT (when viewed by an observer at spatial infinity). Since the system is spherically symmetric, it will often be more instructive to consider radial fermion densities (i.e. densities per radius as opposed to per volume). We shall define these quantities as $r^2 n_f(r)$ and $r^2 \rho_f(r)$ respectively, noting that we have dropped the factor of 4π that should strictly be included.

2.3.2 Gravitational mass

Since Dirac solitons are localised solutions to the static, spherically symmetric FSY equations, the spacetime outside the fermion source asymptotes towards the usual Schwarzschild form (1.16). Hence the total ADM mass M of a state can be obtained by simply comparing this to the form of the metric (2.3) at large r , i.e.

$$M = \lim_{r \rightarrow \infty} \left(\frac{r}{2G} (1 - A(r)) \right). \quad (2.46)$$

To analyse the gravitational mass distribution within the fermion source, we turn to the Komar mass (see section 1.4.5). Using the surface integral form (1.18), and taking the surface S as a 2-sphere of radius r , the Komar mass evaluates to:

$$\begin{aligned} M_K(r) &= -\frac{1}{4\pi} r^2 \int_0^\pi \sin \theta \, d\theta \int_0^{2\pi} d\phi \frac{1}{V} n^\mu \xi^\nu (\nabla_\nu \xi_\mu) \\ &= -r^2 \sqrt{A} \frac{T'}{T^2}, \end{aligned} \quad (2.47)$$

where we have taken the outward pointing normal to be $n^\mu = (0, \sqrt{A}, 0, 0)^T$ and the time-like Killing vector to be $\xi^\mu = (1, 0, 0, 0)^T$. This quantity can be interpreted as the gravitational mass contained within a sphere of radius r . An alternative expression can be obtained by using the volume integral form (1.19), with V taken to be a spherical volume of radius r with future-directed normal $n^\mu = (T, 0, 0, 0)^T$. This evaluates to:

$$\begin{aligned} M_K(r) &= 8\pi \int_0^r \left(T^t_t - \frac{1}{2} T^\mu_\mu \right) g_{tt} \frac{T}{\sqrt{A}} s^2 \, ds \\ &= 4\pi N_f \int_0^r (2\omega T(\alpha^2 + \beta^2) - m(\alpha^2 - \beta^2)) \frac{1}{\sqrt{A}} \, ds, \end{aligned} \quad (2.48)$$

where we have used the explicit forms for the components of the energy-momentum tensor. Note that this integral form is valid only if the spacetime does not contain a singularity, since Stokes’s theorem is used in its derivation. Also note that the total Komar mass

$M_K(\infty)$ coincides with the ADM mass M , as can be easily verified by substituting the Schwarzschild metric into (2.47).

2.3.3 Curvature invariants

A variety of measures that can be used to assess the curvature of a spacetime were introduced in section 1.4.2. Here, we shall only present explicit expressions for the two scalar curvature quantities (the Ricci scalar and the Kretschmann scalar), since these prove useful in detecting the presence of singularities. Starting with the Ricci scalar, an expression in terms of the metric fields A and T can be obtained by performing a lengthy calculation involving the sequence Christoffel symbols \rightarrow Riemann tensor \rightarrow Ricci tensor \rightarrow Ricci scalar. This results in the following:

$$R = \frac{2}{r^2} - \frac{2A}{r^2} - \frac{2A'}{r} + \frac{4AT'}{rT} + \frac{A'T'}{T} - \frac{4A(T')^2}{T^2} + \frac{2AT''}{T}. \quad (2.49)$$

Alternatively, the Ricci scalar can be written in terms of the energy-momentum tensor by taking the trace of the Einstein equations (1.12):

$$R = -8\pi GT^\mu{}_\mu \quad (2.50)$$

$$= 8\pi G \frac{mN_f}{r^2} T(\alpha^2 - \beta^2). \quad (2.51)$$

This is a much more convenient expression from a numerical point of view, since it does not involve any derivatives. Note, however, that this evaluates to zero when $m = 0$, and thus the Ricci scalar cannot provide a reliable measure of curvature in such cases. We therefore turn to the Kretschmann scalar, for which the following expression in terms of the metric fields can be obtained by using (1.10):

$$K = \frac{4}{r^4}(1 - 2A + A^2) + \frac{2}{r^2} \left((A')^2 + \frac{4A^2(T')^2}{T^2} \right) + \frac{1}{T^2} (A'T' + 2AT'')^2 - \frac{8}{T^3} (AA'(T')^3 + 2(AT')^2T'') + \frac{16}{T^4} (A^2(T')^4). \quad (2.52)$$

We can then use the form for the Ricci scalar (2.49), along with the FSY equations themselves, to eliminate the derivative terms, and thus obtain the following lengthier, but more numeric-friendly expression:

$$K = \frac{4}{r^4} \left(3\hat{A}^2 + \Delta^2 + 3\hat{A}\Xi + \Delta\Xi + 2\hat{A}\Omega + \Omega\Delta + 2\Xi\Omega + 6\Xi^2 + 3\Omega^2 \right), \quad (2.53)$$

where

$$\hat{A} = A - 1; \quad (2.54)$$

$$\Delta = -8\pi GN_f \omega T^2 (\alpha^2 + \beta^2); \quad (2.55)$$

$$\Xi = 8\pi G \kappa N_f \frac{T\alpha\beta}{r}; \quad (2.56)$$

$$\Omega = 8\pi GN_f m T(\alpha^2 - \beta^2). \quad (2.57)$$

We shall make particular use of these curvature invariants when analysing the properties of our singular solutions in Chapter 6.

2.4 Two-fermion states: a review

We now present a brief summary outlining the behaviour of the two-fermion system, including examples of individual numerical solutions as well as the spiralling nature exhibited by the families of states. Results for this system were first presented by Finster, Smoller & Yau in ref. [3], and much of what follows is contained therein, although here we extend the analysis to much higher excited states.

2.4.1 Individual states

The procedure for generating particle-like solutions of the FSY equations (2.28)–(2.31) can be summarised as follows. In terms of numerics, we find MATHEMATICA’s differential equation solver, NDSolve, to be more than adequate, and utilise its built-in explicit Runge-Kutta method. An accuracy of 8 digits is usually sufficient for the two-fermion system, although to obtain ultra-high-redshift states this needs to be increased. We initialise the numerical solver at a small, but non-zero, radius, using the small- r asymptotic expansion (2.35)–(2.38), and initially seek unscaled solutions for which $\tilde{m} = \tilde{T}_0 = 1$. For each choice of the parameter $\tilde{\alpha}_1$ (or equivalently z), a discrete set of particle-like states is found to exist, with a distinct ground state and an infinite sequence of excited states, each with an increasingly higher (unscaled) fermion energy and an increasing number of nodes (zeros) in the fermion fields α and β . To obtain these states, we implement a 1-parameter shooting method (a simple binary chop) in $\tilde{\omega}$, and search for normalisable solutions that contain the appropriate number of fermion field nodes. Once acquired, these unscaled solutions can be converted into normalised, physical states by using the rescaling technique outlined in section 2.2.2.

Examples of three such states are shown in Fig. 2.1, where for each we plot the radial profiles of the fermion fields α and β , the metric fields A and T , and the radial fermion energy density $r^2 n_f$ (with n_f defined as in (2.44)). All three states have the same central redshift $z = 2.00$, but differ in the number of fermion field nodes n . Note that we are choosing to label states by the *total* number of nodes in α and β ; thus the $n = 1$ state contains a single node in β and none in α , while the $n = 8$ state contains four in each. This numbering system encompasses both positive and negative parity states, with even values of n denoting the former and odd values the latter. We therefore refer to the $n = 0$ state as an even-parity ground state, the $n = 1$ state as an odd-parity ground state, and the $n = 8$ state as a fourth even-parity excited state. Since this terminology can be somewhat

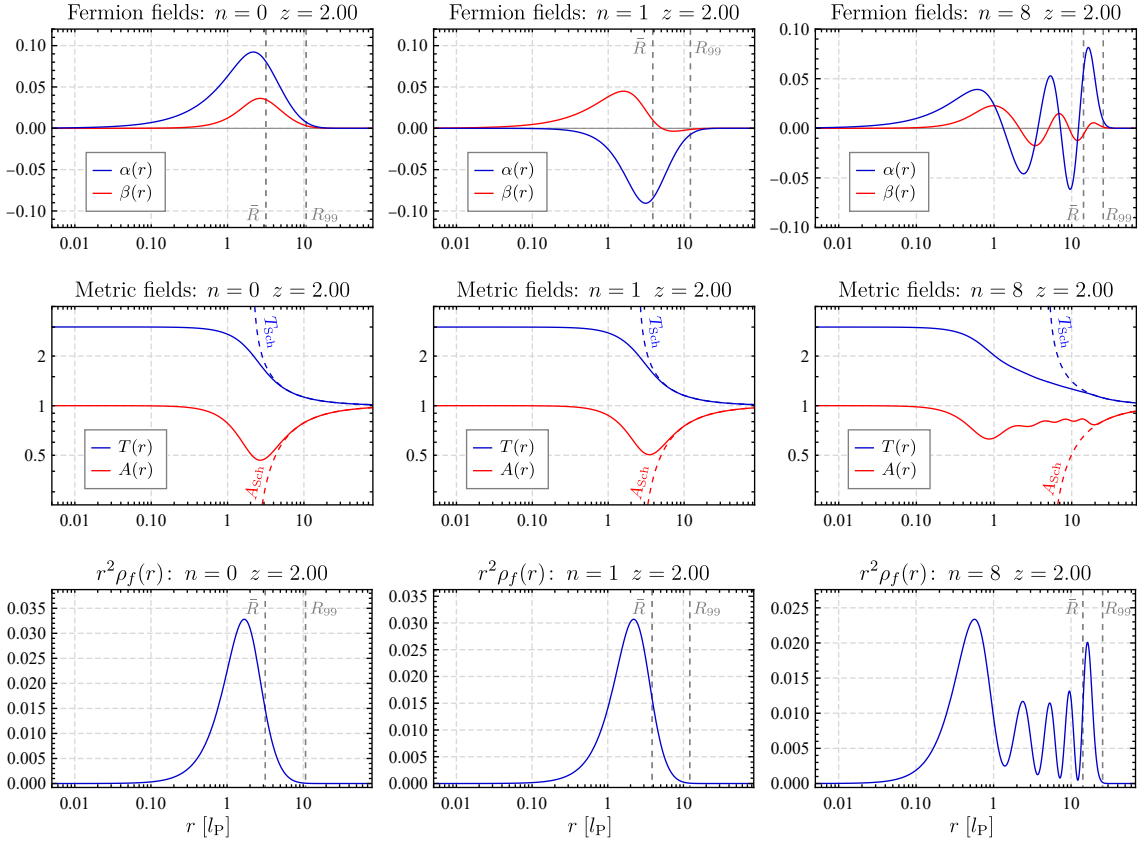


Figure 2.1: Plots showing the radial profiles of the fermion fields, metric fields and radial fermion energy density for an $n = 0$ state (**left**), an $n = 1$ state (**middle**) and an $n = 8$ state (**right**). All three states have the same central redshift $z = 2.00$. Included as dashed lines alongside the metric fields are the corresponding Schwarzschild forms A_{Sch} and T_{Sch} . Radial distances are measured in units of the Planck length $l_{\text{P}} = \sqrt{\hbar G/c^3}$, and two measures of radial extent, \bar{R} and R_{99} , are shown. Numerical values of quantities such as the fermion mass and energy for these states can be found in Appendix B.

confusing, we shall attempt always to specify the value of n when discussing excited-state solutions, although for the majority of this thesis (with the exception of Chapter 4) we are concerned solely with $n = 0$ even-parity ground states.

With regard to the spacetime structure of the states shown in Fig. 2.1, all three exhibit similar metric field profiles, with T decreasing monotonically from a central maximum, while A drops from an initial value of 1 before ultimately rising again. Both fields latch on to their appropriate Schwarzschild forms, $T_{\text{Sch}}^{-2} = A_{\text{Sch}} = 1 - 2GM/r$, at large r and asymptote towards 1 as required. These features are observed to be common to all (non-singular) particle-like solutions of the FSY equations. The presence of additional fermion field oscillations in the $n = 8$ state results in corresponding oscillations in A . These are of relatively small-amplitude, however, indicating that the spatial distortion caused by excited states is less pronounced, although it is spread over a larger radial extent.

An indication of the distribution of the fermion source can be obtained by considering

the radial fermion energy density profiles. These are almost identical for the $n = 0$ and $n = 1$ states, despite the additional fermion node in the latter, while the effects of the fermion field oscillations are clear to see for the $n = 8$ state. Note that the larger the peak in the radial energy density, the larger the amount of local spatial distortion as measured by the metric field A . Included on the fermion plots are also two measures of soliton radius, R_{99} and \bar{R} . The former is defined as the smallest radius which contains 99.9% of the total ADM mass of the state, while the latter is defined as:

$$\bar{R} = \left(4\pi \int_0^\infty r^2 \frac{T}{\sqrt{A}} (\alpha^2 + \beta^2) dr \right)^{1/2}, \quad (2.58)$$

i.e. the rms radius of the state weighted by the radial fermion probability density. These two quantities fulfil different purposes. Whereas R_{99} can be used as a basic estimate of the total radial extent of the fermion source, \bar{R} is a more subtle measure that takes into account its interior distribution. We shall utilise both these notions of radius throughout this thesis. In relation to the solutions in Fig. 2.1, note that the $n = 8$ state has a much larger radius (by either measure) than the $n = 0$ and $n = 1$ states, in order to accommodate the additional oscillations in the fermion fields. Note also that, although the sizes of these states are only of the order of tens of Planck lengths, this is purely a consequence of studying fermions which interact only gravitationally. In more realistic models, where for example the fermions are charged, one would expect states to be significantly larger in extent.

2.4.2 Families of states

Taken together, individual solutions with a common value of n form a continuous family of states parametrised by the central redshift z (or equivalently the unscaled parameter $\tilde{\alpha}_1$). Within these families, it is found that quantities such as the fermion mass and energy oscillate towards constant values as the redshift increases, resulting in a spiralling behaviour when these are plotted against each other.




Figure 2.2: The behaviour of the family of $n = 0$ ground-state solutions, showing the fermion binding energy (**left**), the soliton radius (**middle**), and the total binding energy (**right**) as a function of fermion mass. The insets show zooms of the spiralling parts of the curves, and the total binding energy is separated into stable and unstable branches. All masses and energies are measured in units of the Planck mass $m_P = \sqrt{\hbar c/G}$.

The realisation of this for the $n = 0$ ground state is shown in Fig. 2.2, where we plot the fermion binding energy $m - \omega$, the soliton radius R_{99} and the total binding energy of the state E_b as a function of the fermion mass. This latter quantity is defined as

$$E_b = M - N_f m, \quad (2.59)$$

i.e. the difference between the ADM mass of the state and the mass of N_f (in this case two) isolated fermions. Note that, with this definition, a state is considered bound if it has a negative value of E_b . There are a number of important features to extract from these plots. First, there exists a maximum fermion mass, $m_{\max} \approx 0.607 m_{\text{P}}$, beyond which ground-state solutions do not exist, and similarly an accompanying maximum ADM mass, $M_{\max} \approx 1.168 m_{\text{P}}$. It is understood that these limits owe their existence to the fact that the repulsive effects of the uncertainty principle cannot prevent the gravitational collapse of states above a certain mass threshold (in analogy with e.g. the Chandrasekhar limit).

Both the fermion binding energy and the mass-radius relation for the $n = 0$ family exhibit a spiralling behaviour, as mentioned, where the central redshift increases monotonically from zero (at $m = 0$) to an infinite value at the centre of the spirals. Low-redshift, non-relativistic states have a large radial extent, since they consist of (relatively) light fermions, and can be well-approximated by solutions to the Newton–Schrödinger system (see ref. [51] for details). We derive some of the low-redshift scaling relations that exist between quantities in Appendix A. A consequence of this spiralling behaviour is that, in the relativistic regime, there can exist multiple ground-state solutions with the same value of the fermion mass. This is perhaps surprising, since one might expect there to be a unique equilibrium configuration for a pair of gravitationally-bound fermions of mass m , and this is indeed the case in the non-relativistic regime. This behaviour appears to be a consequence of the non-linear nature of general relativity, as we shall discuss further when considering the fermion self-trapping effect in Chapter 3.

The total binding energy plot in Fig. 2.2 also includes information concerning the stability of solutions. It was shown in ref. [3], using Conley index theory, that states with redshifts below that of the maximum-mass state are stable to spherically-symmetric perturbations, while those with redshifts above are unstable. The unstable branch thus includes solutions in which the total binding energy is negative and one might naively expect to be stable. This overall behaviour is similar to that of many compact astrophysical objects, in which the maximum mass similarly corresponds to a stable-to-unstable transition point. It is important to remember, however, that all states on the $n = 0$ curve (and indeed higher n curves) are static, i.e. they are equilibrium configurations in which the gravitational attraction between the fermions is precisely balanced by the repulsive effects of the uncertainty principle. Thus, left unperturbed, these states will continue to exist ad infinitum, with the fermion wavefunction oscillating at a constant frequency ω . What distinguishes the stable and unstable branches is the behaviour of states under

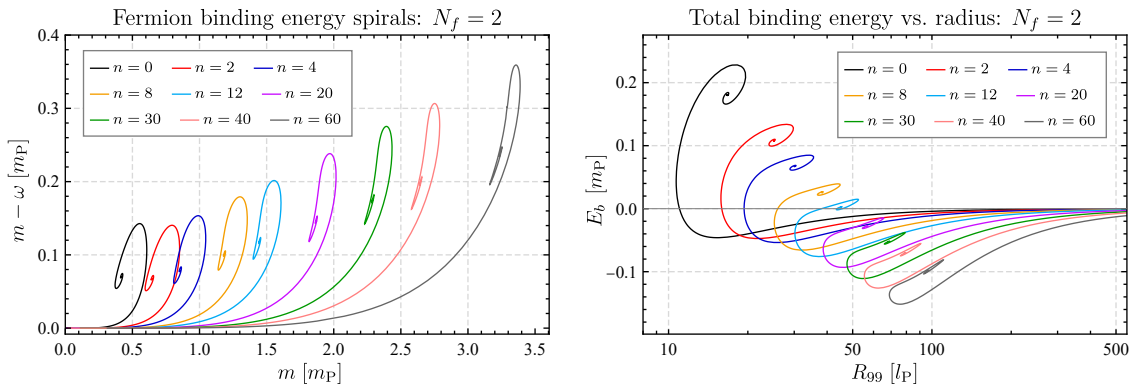


Figure 2.3: The behaviour of (even-parity) excited-state families of solutions, showing the fermion binding energy as a function of fermion mass (**left**) and the total binding energy as a function of soliton radius (**right**) for families with the values of n indicated.

small perturbations; those on the stable branch will ultimately return to their original form, while those on the unstable branch will not (tending to collapse, explode, or decay to stable states).

Similar spiralling families exist for states with higher values of n , a selection of which are shown in Fig. 2.3, where we plot the fermion binding energy as a function of fermion mass, and the total binding energy as a function of radial extent. Note that we are restricting our analysis here to even-parity states. The curves are qualitatively similar for all values of n , although the higher excited-state spirals become increasingly distorted. The maximum values of both m and ω increase with increasing n , and indeed an extrapolation indicates that states can exist with arbitrary large values of fermion mass, ADM mass and fermion energy, provided they are sufficiently excited. Interestingly, the total binding energy shows a general decrease with increasing n , to such an extent that, for $n \geq 14$, all states become bound. This does not of course imply stability; indeed we would expect the maximum-mass state for each family to again represent a stable-unstable transition point. It does, however, suggest that highly excited states should remain localised under perturbations, since delocalisation would be energetically unfavourable.

We end this section with a remark on terminology. The word ‘excited’ has been used here to designate states with higher numbers of fermion nodes, but this does not necessarily imply higher fermion energy as one would expect from a non-relativistic perspective. If we were to consider the set of states at a constant fermion mass, for example, we might find that (for large fermion masses) an ‘excited’ state in fact has the lowest possible energy, or indeed (due to the spiralling nature of families) that more than one ‘ground-state’ solution exists. Comparing states instead at constant redshift somewhat mitigates this inconsistency, although only for systems with small numbers of fermions. Thus the term ‘excited’ here should solely be regarded as referring to the number of fermion nodes present in a particular state. Interestingly, however, the notion of ‘excited’ does apply as

expected in the unscaled system, where each subsequent excited state does indeed have a larger fermion energy than the last. This feature will prove particularly useful when multivaluedness becomes an issue in Chapter 4.

2.5 Structure of Dirac solitons

We now review the more recent analysis by Bakucz Canário *et al.* [4] concerning the internal structure of Einstein–Dirac solitons, which will prove relevant to much of the material in Chapters 3 and 4.

2.5.1 The power-law solution

This begins with the discovery of an analytic, singular solution to the FSY equations (2.28)–(2.31), valid when the fermions are massless, in which all fields exhibit a simple power-law dependence. We note that this solution (extended to d dimensions) was also independently found by Blázquez Salcedo & Knoll [54], who interpret it as a ‘light-like singularity’. The metric and fermion fields for this solution take the following forms:

$$\begin{aligned}\alpha(r) &= \sqrt{\frac{\omega}{12\pi GN_f^2 \kappa_-}} r; & A(r) &= \frac{1}{3}; \\ \beta(r) &= \sqrt{\frac{\omega}{12\pi GN_f^2 \kappa_+}} r; & T(r) &= \frac{1}{\omega} \sqrt{\frac{N_f^2}{4} - \frac{1}{3}} r^{-1},\end{aligned}\quad (2.60)$$

where $\kappa_{\pm} = \kappa/2 \pm \sqrt{1/3}$. Note that this solution cannot be normalised, since the integrand in the normalisation integral (2.34) diverges as $r \rightarrow \infty$, and thus it is neither localised nor should be considered a proper quantum state. The presence of a central singularity can be inferred from the Kretschmann scalar (2.53), which evaluates to:

$$K = \frac{8}{3} r^{-4} \quad (\text{for } \kappa > 0); \quad K = \frac{56}{9} r^{-4} \quad (\text{for } \kappa < 0). \quad (2.61)$$

This clearly diverges at small r and thus the spacetime is singular at $r = 0$. Note that the Ricci scalar (2.51) evaluates to zero everywhere, since the solution is massless, and therefore cannot be used for this purpose.

2.5.2 Zonal structure

Given that the above power-law solution is valid only when $m = 0$, its relationship to the massive particle-like states of the Einstein–Dirac system is not immediately obvious. It was shown in ref. [4], however, that, for relativistic states ($z \gtrsim 1$), a zone exists in which the mass terms in the equations of motion are vastly subdominant, and thus the wavefunction in that region behaves as if the fermions were massless. This region is referred to as the




Figure 2.4: (Left): Illustration of the zonal structure of Einstein–Dirac solitons, showing the radial profiles of the metric and fermion fields for a highly-relativistic $n = 4$ state. The separation into four distinct zones (core, power-law zone, wave zone and evanescent zone) can clearly be seen. Note that the fields have been rescaled such that they appear constant within the power-law zone. (Right): Behaviour of the power-law oscillations, indicating how the radii of extrema in the metric field A vary as the redshift of the $n = 0$ state is varied. Also included are the boundaries of the power-law zone (which we define as the radii at which $A = 2/3$) and the total radial extent of the fermion source.

‘power-law’ zone, and is one of four distinct zones identified in the internal structure of relativistic Dirac solitons.

This zonal structure is illustrated in the left-hand panel of Fig. 2.4, where we plot the radial profiles of the fermion and metric fields (scaled by appropriate powers of r) for a highly-relativistic $n = 4$ state. The four distinct zones can be described as follows. The innermost is the ‘core’, in which the fields follow the small- r asymptotic expansion (2.35)–(2.38). Thereafter follows the power-law zone, in which the fermion mass is negligible compared to the local fermion energy ωT , and the fields do indeed approximately follow the massless power-law solution (2.60). As we move further outwards in radius, ωT continues to decrease and the fermion mass term becomes increasingly relevant, ultimately signalling the end of the power-law zone. For excited states ($n \geq 2$), a ‘wave zone’ is then entered, which contains the fermion field oscillations that define the value of n . At even larger radii, the local fermion energy drops to a value comparable to that of the fermion mass, at which point the state transitions to the ‘evanescent zone’, where the fermion fields decay exponentially and the metric fields are approximately Schwarzschild. For ground-state solutions, the wave zone ceases to exist and there is therefore a direct transition from the power-law zone (or core if the state is non-relativistic) to the evanescent zone.

Notice that the fields in the power-law zone do not follow the power-law dependence precisely, and, in particular, damped, small-amplitude oscillations are in evidence near the inner boundary. An analysis of these, for the metric field A , is shown in the right-hand panel of Fig. 2.4, where the radii of the extrema in A are plotted as a function of central redshift. Note that the power-law zone only appears at $z \approx 1$, i.e. when the states become relativistic, and its inner boundary shrinks towards $r = 0$ as the redshift increases,

while its outer boundary tends to a constant radius. For non-relativistic states, a single minimum in A exists, with subsequent pairs of extrema gradually appearing within the power-law zone as the redshift increases. One can see that these oscillations are evenly spaced in $\log(r)$ and shrink towards $r = 0$ along with the inner power-law zone boundary. Overall, highly-relativistic states therefore contain an extended power-law zone with a large number of ‘power-law oscillations’. It is noted in ref. [4] that the frequency at which new extrema appear in the power-law zone coincides with that of the oscillations in m , ω etc. that ultimately source the spiralling structures shown earlier in Fig. 2.2. This suggests that there is a fundamental connection between the power-law oscillations and the existence of a spiralling behaviour. We shall discuss this point further in Chapter 3.

2.5.3 Infinite-redshift states

As noted above, the inner boundary of the power-law zone shrinks towards $r = 0$ as the central redshift is increased, while the outer boundary remains at a finite radius. This implies that, in the limit of infinite redshift, the core zone should no longer exist, and the power-law zone should extend all the way outwards from $r = 0$. In turn, this suggests that the infinite-redshift state (the state that lies at the centre of the spirals) can be obtained by simply replacing the small- r expansions (2.35)–(2.38) with the power-law expressions (2.60) when generating solutions. This indeed proves to be the case, allowing such states to be constructed numerically in a straightforward manner.

The infinite-redshift $n = 0$ state is shown in Fig. 2.5, where we plot the fermion fields, the metric fields and the radial fermion energy density, as before. The power-law dependence is evident at small r , up until the point at which the fermion mass term becomes significant, and the state transitions to an evanescent zone. Note that the radial fermion energy density approaches a constant value as $r \rightarrow 0$, in contrast to the finite-redshift solutions shown in Fig. 2.1. This implies that the volume density $\rho_f(r)$ must diverge as $r \rightarrow 0$, and hence a central singularity is present, a fact that could be predicted by considering the power-law form of the Kretschmann scalar. Note, however, that since the radial probability density is finite at $r = 0$, the state can still be correctly normalised,




Figure 2.5: The infinite-redshift $n = 0$ state, showing the radial profiles of the fermion fields, metric fields and radial fermion energy density. This state is located at the centre of the $n = 0$ spiral shown in Fig. 2.2, as evidenced by the parameter values given in Appendix B.

despite the existence of the singularity. As a final remark, one might wonder why the infinite-redshift solution does not appear to exhibit power-law oscillations, unlike the finite redshift states. The answer to this can be construed by considering again the right-hand panel of Fig. 2.4. As the central redshift increases, each new pair of extrema appears at a slightly smaller radius than the last, and thus, in the infinite-redshift limit, all power-law oscillations will be located at precisely $r = 0$.

This concludes our brief summary of the two-fermion system. For further details, we point the reader to the aforementioned refs. [3] and [4]. In the following two chapters, we shall extend this analysis to many-fermion Dirac solitons, giving a physical interpretation for the power-law oscillations as well as showing that the distinction between the power-law zone and the wave zone ceases to exist in the context of many-fermion excited states.

Chapter 3

The fermion self-trapping effect

We now begin the first research chapter of this thesis, which concerns the discovery of a so-called fermion ‘self-trapping’ effect, where the fermion wavefunction becomes trapped within highly distorted regions of spacetime, giving rise to a multi-shell-like structure. We shall begin by presenting ground-state, particle-like solutions to the many-fermion Einstein–Dirac system, in which this phenomenon is most prominent, before introducing the concept of optical geometry, which acts as a useful tool for visualising the spacetime of Dirac solitons. We then discuss the self-trapping interpretation itself and its impact on the spiralling behaviour of solution families at large fermion number.

3.1 Many-fermion states

The procedure for generating localised solutions to the many-fermion Einstein–Dirac system is almost identical to that employed in the two-fermion case, outlined in section 2.2. We again initially search for unscaled solutions of the FSY equations (2.28)–(2.31), by tuning the value of the unscaled fermion energy $\tilde{\omega}$, before converting these into physical states via the rescaling method. Given a system of N_f fermions, we find that, for each value of $\tilde{\alpha}_1$ (or equivalently z), there exists a unique $n = 0$ ground state, with lowest unscaled energy $\tilde{\omega}_0$, as in the two-fermion system. This uniqueness does not extend, however, to excited states ($n \geq 2$), where instead we find that multiple solutions can exist that have identical values for both n and $\tilde{\alpha}_1$. In this chapter, we shall restrict our analysis to $n = 0$ ground states (excited states are addressed in Chapter 4), and thus the required values of $\tilde{\omega}$ can be obtained via a simple 1-parameter shooting procedure, as was the case in the two-fermion system. There is a slight complication, however — as the fermion number is increased, the precision to which $\tilde{\omega}_0$ must be specified, in order to ensure that the numerical solver reaches a radius at which the fermion fields have decayed sufficiently, also increases. Although this does not require an alteration to the basic numerical method, solution generation becomes significantly more lengthy and challenging due to the complications that

arise from working with numerics at high precision and accuracy. Fortunately, MATHEMATICA is particularly well-suited to this purpose, allowing us to obtain states containing up to a maximum of $N_f = 90$ fermions (requiring 50-digit precision), which is more than sufficient to demonstrate the fermion self-trapping effect discussed here. Note that this limit is solely a consequence of our computational equipment, and could in principle be extended by using a more powerful system.

The observed behaviour of states with large numbers of fermions is markedly different from the two-fermion case, as is illustrated in Fig. 3.1, which shows three individual $N_f = 90$ states at three different redshift values. Recall that, in the many-fermion Einstein–Dirac system, the fermions are arranged in a filled shell with each having an angular momentum $j = \frac{1}{2}(N_f - 1)$. Thus one would expect the $N_f = 90$ fermion wavefunctions to exhibit a single peak, centred around a large radius, consistent with the notion of a single shell of high-angular-momentum fermions. This is precisely what occurs in the non-relativistic regime, as can be seen for the $z = 0.092$ state shown in Fig. 3.1. Here, both the fermion fields and the radial fermion energy density peak at a radius of approximately $480 l_P$, and

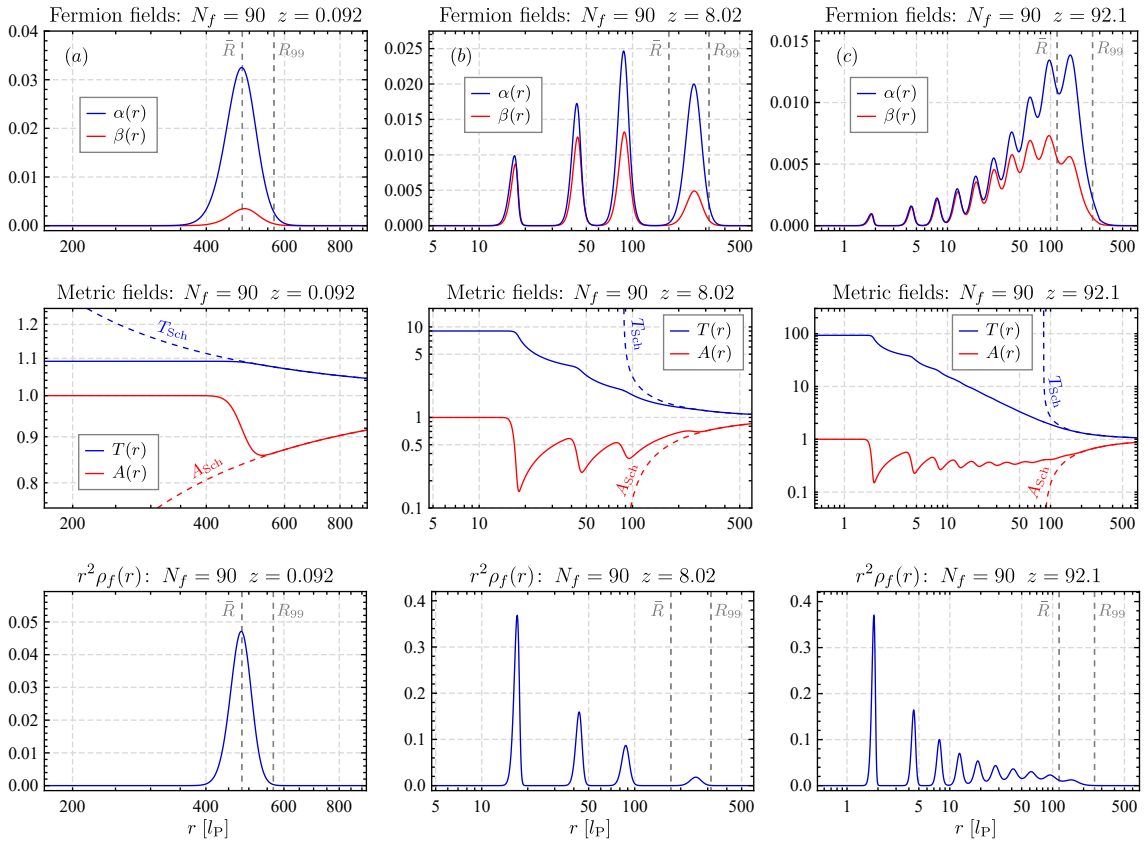


Figure 3.1: Examples of three individual particle-like states for the $N_f = 90$ fermion system, showing the radial profiles of the fermion fields, metric fields and radial fermion energy density for states with redshift values of $z = 0.092$ (left), $z = 8.02$ (middle) and $z = 92.1$ (right). For non-relativistic solutions, the fermion wavefunction exhibits a single peak, while for relativistic states it breaks into a series of multiple shells.

are relatively delocalised, resulting in a moderate spatial distortion as measured by the metric field A .

Once the system becomes relativistic, however, the situation alters significantly. Considering first the mid-redshift $z = 8.02$ state, we see that the fermion wavefunction (and energy density) appears to have split into four distinct peaks, with the inner three located at a much smaller radius than one would expect high-angular-momentum fermions to occupy. The spatial distortion caused by these peaks is significant, as can be inferred from the structure of the metric field A , due to the highly concentrated radial energy density contained within them. In-between the peaks, there exist regions in which the probability of finding a fermion is virtually zero, an observation supported by the fact that the metric fields there are approximately Schwarzschild. This situation is amplified as the redshift is increased further, with the $z = 92.1$ state exhibiting a total of 10 fermion peaks, although those at larger radius are not as spatially separated.

The overall behaviour of relativistic states therefore appears more akin to what one would expect from a multiple-shell model, where fermions of differing angular momenta occupy shells at differing radii. But this is categorically not the situation here — each fermion has precisely the same (total) angular momentum, and furthermore the wavefunction of each fermion has an identical radial dependence. Thus the observed splitting into multiple shells cannot be attributed to angular momentum effects, and instead must have an alternative origin.

Before discussing this origin in more detail, we first summarise the behaviour of the system as the number of fermions is varied. This is illustrated in Fig. 3.2, where we plot three states, each with an identical redshift value of $z = 92.1$, for three different values of N_f . There are a number of features to note. First, all three states exhibit oscillations in the fermion and metric fields, with the amplitude of these increasing as the fermion number is increased. For the $N_f = 20$ and $N_f = 40$ states, the first oscillation is of large enough amplitude to cause the inner peak to become separated from the remainder of the wavefunction, whereas this does not occur for the $N_f = 6$ case. Furthermore, the number of oscillations also increases with the fermion number, at least when comparing states at constant redshift. Note that this is not particularly insightful, however, since the number of oscillations at constant N_f itself increases with increasing redshift (as evidenced in Fig. 3.1). This last observation makes it difficult to analyse precisely how the structure of states changes as we vary the fermion number, since we would ideally compare solutions that have an equal number of oscillations. Such a set of states cannot be obtained by evaluating at constant redshift (or $\tilde{\alpha}_1$), however, nor does there appear to be a simple scaling that relates the redshift to the number of oscillations. The analysis we can perform is therefore limited in extent, but we are aided by the observation that, once a peak fully forms, its height and width (in $\log(r)$) do not change significantly as the redshift is further increased. Using this feature, we find that, for sufficiently large N_f ,

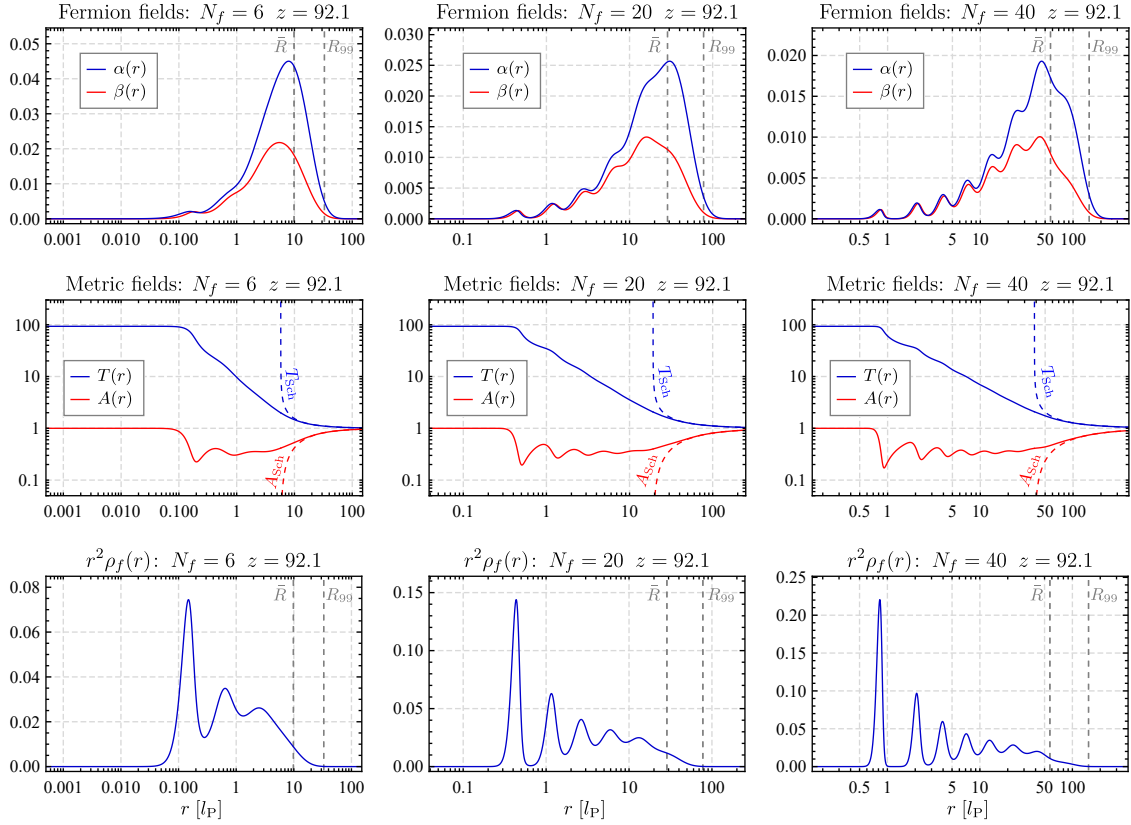


Figure 3.2: Illustration of the change in structure of many-fermion Dirac solitons as the number of fermions is varied, showing the fermion fields, metric fields and radial fermion energy densities for states with fermion numbers of $N_f = 6$ (left), $N_f = 20$ (middle), and $N_f = 40$ (right). All three states have the same central redshift $z = 92.1$. Note in particular the increased separation between the inner fermion peaks as the fermion number is increased.

and for peaks that are sufficiently well-separated, the height of each peak in $r^2 \rho_f$ scales approximately as $N_f^{2/3}$, while the width (in $\log(r)$) scales as $N_f^{-2/3}$. Thus the peaks in the energy density become progressively taller and narrower as the fermion number increases. Intriguingly, this implies that, in the limit of infinitely many fermions, the energy density should split into a series of delta functions, suggesting that some analytic progress may be possible in the limit of $N_f \rightarrow \infty$.

As a final remark, we note that the above oscillatory behaviour is in fact present at all values of N_f , even for the two-fermion system studied previously, in which it manifests itself as the small-amplitude power-law oscillations discussed in section 2.5. We thus conclude that, as the fermion number increases, these oscillations increase in amplitude, ultimately sourcing the extreme wavefunction splitting effects seen at large N_f . As further evidence for this, we observe that all peaks in the fermion energy density (regardless of the value of N_f) are indeed located within the power-law zone, and in addition are evenly spaced in $\log(r)$. The exceptions to this are the final fermion peaks, which are located within the evanescent zone, and any inner peaks that are highly separated from the remainder of the wavefunction, for which the $\log(r)$ spacing does not appear to hold.

3.2 Optical geometry

How exactly are we to interpret the observed splitting of the fermion wavefunction into a multiple-shell-like structure? To answer this question, we must analyse more thoroughly the spacetimes associated with these highly-relativistic, many-fermion Dirac solitons. The tool that will prove most useful in this regard is the ‘optical geometry’, a concept initially introduced by Dowker & Kennedy [71] and Gibbons & Perry [72], and one which has found a variety of applications within general relativity. It provides a means through which the structure of a particular spacetime can be analysed from the perspective of a null particle, allowing features such as the presence of ‘photon spheres’ (circular null geodesics) to become obvious.

3.2.1 Optical embedding diagrams

This visualisation is achieved by constructing a so-called optical geometry embedding diagram, the basic procedure for which is as follows. Given a static, spherically symmetric spacetime with line element ds^2 , one first introduces the ‘optical metric’ $d\tilde{s}^2$ by rescaling ds^2 such that the prefactor in front of the tt component is unity, i.e. we define:

$$d\tilde{s}^2 = \frac{1}{g_{tt}} ds^2. \quad (3.1)$$

One can then implement the standard procedure for embedding the spatial part of a spacetime within a Euclidean cylindrical co-ordinate system (h, ρ, φ) , for which the line element de^2 takes the form:

$$de^2 = dh^2 + d\rho^2 + \rho^2 d\varphi^2. \quad (3.2)$$

The aim is to construct a surface $h = f(\rho)$ (the optical embedding diagram), within this co-ordinate system, that is isometric to the spatial part of our optical metric. This can be achieved by rewriting $d\tilde{s}^2$ in the following form:

$$d\tilde{s}^2 = -dt^2 + \left(1 + \left(\frac{df}{d\rho} \right)^2 \right) d\rho^2 + \rho^2 d\varphi^2, \quad (3.3)$$

where the two angular co-ordinates $\{\theta, \phi\}$ have been combined into a single co-ordinate φ by evaluating at $\theta = \pi/2$, possible due to the redundancy that exists in spherically symmetric systems. An expression for the surface $h = f(\rho)$ can then be inferred by a straightforward comparison of terms. We note that a standard, non-optical embedding diagram can be obtained in a similar manner, by simply embedding the spatial part of the original metric ds^2 as above, without the initial rescaling.

Examples of optical embedding diagrams for various well-known spacetime metrics can be found in refs. [73–77]. Of particular relevance here is the study of compact as-

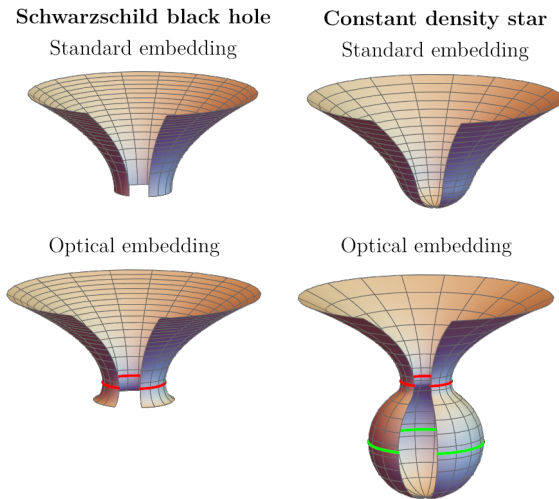


Figure 3.3: Standard and optical embedding diagrams for the usual Schwarzschild black-hole metric (1.16) (**left**), and a constant density star described by the interior Schwarzschild metric (1.17) (**right**). Both solutions are generated with an ADM mass of $M = 1$, while the interior solution represents an ultra-compact star of radius $R = 2.3$. Note that the black-hole Schwarzschild metric is only embeddable for $r > 2M$ (in the standard case) and $r > 9M/4$ (in the optical case). For the optical embedding diagrams, orbits corresponding to circular null geodesics (photon spheres) are highlighted, with those that are stable shown in green and unstable in red.

trophysical objects, such as those described by the interior Schwarzschild solution (1.17). An analysis of the optical geometry of this metric was first undertaken by Abramowicz *et al.* [78], the results of which we reproduce and summarise in Fig. 3.3. This shows both the standard and optical embedding diagrams for the usual Schwarzschild black-hole metric (1.16), alongside the corresponding diagrams for an ultra-compact, constant-density Schwarzschild star with ADM mass $M = 1$ and radius $R = 2.3$. Considering first the standard embedding diagrams, the black hole metric can be embedded (using the aforementioned technique) only outside the horizon ($r > 2M$), whereas the spacetime of the constant density star is everywhere embeddable. Other than this, these diagrams are relatively featureless, and there is little useful information to gain from them.

By contrast, the optical geometry embedding diagrams exhibit a significantly more interesting structure. As mentioned, these provide a visualisation of the spacetime from the point of view of a null particle, since the trajectories of null geodesics are confined to the resulting embedded surfaces. For the Schwarzschild black hole solution, we observe the appearance of a ‘neck’ within the optical geometry (shown in red) at $r = 3M$, which indicates the location of an unstable circular null geodesic (photon sphere). For the constant density star, this neck is joined by a ‘bulge’ (shown in green), where a stable photon sphere resides, creating an overall ‘bottleneck’ structure. We note that this type of behaviour is only apparent for stars that are ultra-compact ($R < 3M$), and as such the unstable photon sphere is always located outside the star, whereas the stable orbit resides within.

One particularly interesting property of ultra-compact stars is their ability to trap gravitational waves, a feature first discovered by Chandrasekhar & Ferrari [79]. This was further explored by Abramowicz *et al.* [78], using the optical geometry viewpoint, who argue that the presence of a pair of photon spheres is crucial to this trapping phenomenon, ultimately due to the reversal in centrifugal force that accompanies each circular null geodesic. The effect of this reversal is such that null particles are always attracted

towards the locations of stable circular null geodesics, while being repelled from the locations of their unstable counterparts. Thus gravitational waves, or indeed any disturbances in a similarly massless field, will become trapped within the ‘bottle-neck’ structure illustrated in Fig. 3.3. For further details concerning this centrifugal force reversal, and its often counter-intuitive consequences, see refs. [80–83]. Of relevance here is simply the observation that the presence of a bottleneck-type structure permits the trapping of null particles.

We now proceed to the task of applying this optical geometry framework to our many-fermion Einstein–Dirac solitons. To this end, we shall first derive an expression for the optical embedding surface $f(\rho)$, valid for a static, spherically symmetric spacetime, using the general metric form given in (2.3), before subsequently applying it to our numerical solutions. We begin by defining the optical metric using (3.1):

$$d\tilde{s}^2 = T^2 ds^2 = -dt^2 + \frac{T^2}{A} dr^2 + r^2 T^2 d\Omega^2. \quad (3.4)$$

Evaluating this at $\theta = \pi/2$, and identifying the remaining angular co-ordinate with φ , we can immediately write the cylindrical radial co-ordinate as $\rho = rT$. Comparing the form of (3.4) with (3.3), we then obtain the following condition:

$$df^2 + (Tdr + rdT)^2 = \frac{T^2}{A} dr^2, \quad (3.5)$$

which can be integrated to find an explicit expression for $f(r)$:

$$f(r) = \int_0^r T(u) \sqrt{\frac{1}{A(u)} - \left(1 + \frac{uT'(u)}{T(u)}\right)^2} du. \quad (3.6)$$

This, along with $\rho = rT$, fully defines the optical embedding diagram for a general spherically symmetric, static spacetime.

Using these expressions, we can straightforwardly construct optical embedding diagrams for our many-fermion Dirac solitons. Examples of these, for the 90 fermion system, are shown in Fig. 3.4, for the three redshift cases presented earlier in Fig. 3.1, alongside the infinite-redshift state. For comparison, we have also included the standard embedding diagrams, for which $\rho = r$ and

$$f(r) = \int_0^r \sqrt{\frac{1}{A(u)} - 1} du. \quad (3.7)$$

Consider first the finite-redshift states. For the non-relativistic $z = 0.092$ state, the optical geometry exhibits a saucer-like appearance, but as the redshift increases, the base of this extends into a tubular structure, along which a series of necks and bulges gradually form. This region of spacetime corresponds to the power-law zone (the size of which increases

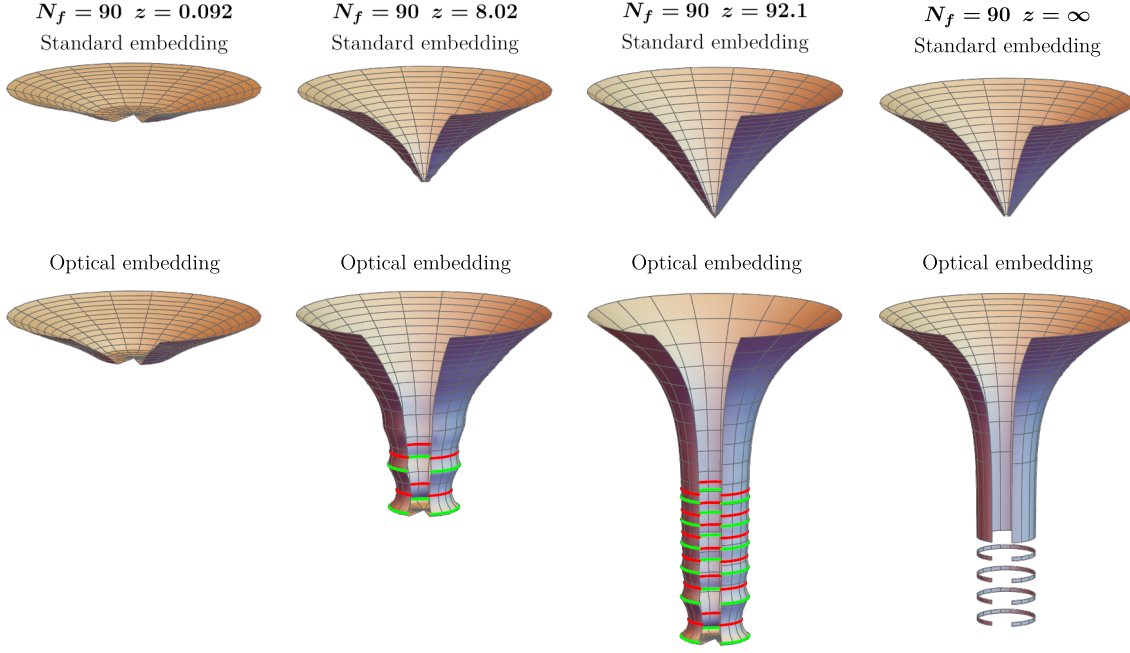


Figure 3.4: Standard and optical embedding diagrams for $N_f = 90$ Einstein–Dirac solitons, with the values of redshift indicated. The left-most three correspond to the same three states shown in Fig. 3.1, while the right-most is the infinite-redshift state. Since it contains a central singularity, the spacetime of the infinite-redshift state cannot be embedded at $r = 0$, with this point being located at $h = -\infty$ in the optical embedding diagram. Trajectories of stable circular geodesics are indicated in green, with the corresponding unstable orbits in red.

with redshift), with the necks and bulges a direct consequence of the power-law oscillations. Located at the base of each optical embedding diagram is the core region, which is almost entirely flat for large values of N_f . This is of course absent in the infinite-redshift case, where the power-law zone extends all the way to $r = 0$. Due to the central singularity, the point $r = 0$ in the infinite-redshift state cannot be embedded, and indeed we observe that the power-law zone in the optical geometry is infinite in extent, with $r = 0$ located at $h = -\infty$. The power-law oscillations are also compressed to $h = -\infty$, and thus the optical geometry in the power-law zone is purely cylindrical (noting that the width $\rho = rT$ is a constant for the power-law solution). At large radii (large h), this cylindrical structure ultimately opens up into a parabolic Schwarzschild form, once the evanescent zone has been reached, as is the case for all four solutions.

The behaviour of the optical embedding diagrams can therefore be summarised as follows. For low-redshift, non-relativistic states, no power-law zone exists and hence the flat core broadens directly into an exterior Schwarzschild geometry. As the redshift increases beyond $z = 1$, the appearance of a power-law zone gives rise to a tubular structure containing a series of stable and unstable photon spheres, with both the height of this region and the number of photon spheres increasing with redshift. Ultimately this becomes infinite in extent, exhibiting a purely cylindrical geometry at infinite redshift. Although we have presented here results only for the 90 fermion system, we note that this general

behaviour is replicated for all values of N_f , although the bottleneck structures become increasingly pronounced as the number of fermions increases.

3.2.2 Geodesic motion

The appearance of photon spheres in the optical geometries presented above suggests that, in analogy with the ultra-compact stars studied in ref. [78], the spacetimes of relativistic Dirac solitons should have the ability to trap null particles. This can be explored further by analysing the trajectories of null geodesics. To do so, we consider the Lagrangian $\mathcal{L} = (ds/d\lambda)^2$, i.e.

$$\mathcal{L} = -\frac{1}{T(r)^2}\dot{t}^2 + \frac{1}{A(r)}\dot{r}^2 + r^2\dot{\theta}^2 + r^2\sin^2\theta\dot{\phi}^2, \quad (3.8)$$

where a dot represents a derivative with respect to the affine parameter λ . Without loss of generality, we can restrict our analysis to orbits in the equatorial plane $\theta = \pi/2$. Then, noting that \mathcal{L} has no explicit dependence on either t or ϕ , we can extract the following conservation equations:

$$0 = \frac{\partial\mathcal{L}}{\partial t} = \frac{d}{d\lambda} \left(\frac{\partial\mathcal{L}}{\partial\dot{t}} \right) = \frac{d}{d\lambda} \left(-\frac{2\dot{t}}{T^2} \right) \equiv -2\frac{dE}{d\lambda}; \quad (3.9)$$

$$0 = \frac{\partial\mathcal{L}}{\partial\phi} = \frac{d}{d\lambda} \left(\frac{\partial\mathcal{L}}{\partial\dot{\phi}} \right) = \frac{d}{d\lambda} \left(2r^2\dot{\phi} \right) \equiv 2\frac{dL}{d\lambda}, \quad (3.10)$$

where we have defined the total energy $E = \dot{t}/T^2$ and the angular momentum $L = r^2\dot{\phi}$, which are constants of the geodesic motion. For null geodesics, we have the condition $\mathcal{L} = 0$, which gives rise to the following energy equation:

$$E^2 = \frac{1}{AT^2}\dot{r}^2 + \frac{L^2}{r^2T^2}. \quad (3.11)$$

Applying a co-ordinate transform such that $du^2 = dr^2/(AT^2)$, we can then write this in the form of kinetic energy + potential energy:

$$\frac{1}{2}E^2 = \frac{1}{2}\dot{u}^2 + \frac{L^2}{2r(u)^2T(u)^2}. \quad (3.12)$$

Hence null particles are subject to an effective potential of the form $V_{\text{eff}} = L^2/(2r^2T^2)$. Circular orbits will occur at radii where $V'_{\text{eff}}(u) = 0$, which results in the condition:

$$0 = \frac{dV_{\text{eff}}}{du} = \frac{dr}{du} \frac{dV_{\text{eff}}}{dr} = -L^2T\sqrt{A} \left(\frac{1}{r^3T^2} + \frac{T'}{r^2T^3} \right); \quad (3.13)$$

$$\implies T + rT' = (rT)' = 0. \quad (3.14)$$

Thus circular null geodesics (photon spheres) do indeed exist at the extrema of the function $\rho = rT$, in agreement with the earlier optical geometry analysis.

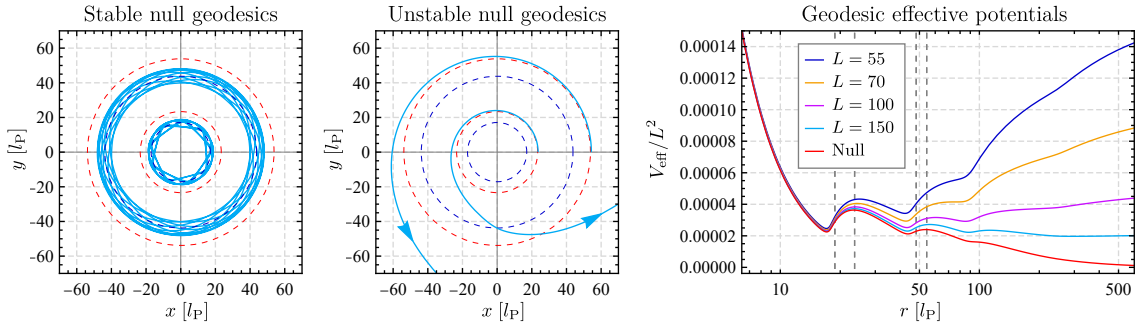


Figure 3.5: Plots showing (left) trajectories of stable, trapped null geodesics within the spacetime of an $N_f = 90$ Einstein–Dirac soliton with redshift $z = 8.02$, and (middle) examples of unstable null geodesics within the same spacetime that escape to infinity. Stable and unstable photon spheres are depicted by dark blue and red dashed lines, respectively. Also shown (right) are the effective potentials experienced by both null and massive geodesics, the latter for the values of angular momentum L indicated. The vertical dashed lines are the radii at which the four geodesics shown left are initialised.

The trajectories for a few (non-circular) null geodesics are illustrated in the two left-hand panels of Fig. 3.5, for the $\kappa = 90$, $z = 8.02$ redshift state shown in previous figures. These are obtained by solving the energy equation (3.11) with zero initial radial velocity and taking $L = 1$ (noting that varying the value of L simply equates to a rescaling of the affine parameter). The two geodesics shown in the left-hand plot are initialised at radii close to those of the two stable photon spheres, and we observe that these indeed equate to stable, trapped orbits, which oscillate between a minimum and maximum radius. In contrast, the geodesics shown in the middle plot are started just outside the unstable photon spheres, and consequently their trajectories escape to infinity. This behaviour can also be predicted by analysing the effective potential V_{eff} to which null geodesics are subject. This is included in the right-hand panel of Fig. 3.5, where the dashed lines indicate the radii at which the aforementioned geodesics are initialised. From this we conclude that trapped orbits can only occur within the potential wells that surround the stable photon spheres, which become progressively shallower as we move outwards in radius.

Figure 3.5 also contains information concerning the behaviour of massive, time-like geodesics. These obey the condition $\mathcal{L} = -1$, from which we obtain the slightly altered effective potential:

$$V_{\text{eff}} = \frac{L^2}{2r^2T^2} + \frac{1}{2T^2}, \quad (3.15)$$

where L is now the angular momentum per unit mass. Massive circular orbits will occur at radii where $V'_{\text{eff}} = 0$, which evaluates to:

$$(rT)' = \frac{T}{1 + L^2/r^2}. \quad (3.16)$$

For geodesics that are highly-relativistic (those with large values of L), this will approximate the equivalent null condition $(rT)' = 0$, as will also be the case at small r . This

behaviour can be observed in the right-hand plot of Fig. 3.5, where we plot the effective potential for massive geodesics with various angular momenta. These closely match the form of the null potential at small r , but gradually diverge at larger radius, with those at lowest angular momentum showing the greatest divergence. An additional circular orbit (in comparison to the null case) also appears for intermediate values of L , located at a larger radius, around which massive particles may become trapped but null particles will not. Finally, we note that the right hand side of (3.16) is strictly positive, and hence massive circular orbits can only exist in regions in which the derivative of $\rho = rT$ is similarly positive. By considering the optical embedding diagrams in Fig. 3.4, we conclude that stable massive circular geodesics must occur at radii slightly inside the corresponding stable photon spheres, whereas unstable massive circular geodesics must occur slightly outside their null counterparts.

3.3 Self-trapping interpretation

We are now finally in a position to explain the wavefunction splitting observed for relativistic many-fermion Dirac solitons, described earlier in section 3.1. The geodesic and optical geometry analysis above has established that the spacetimes of such objects contain regions in which null and (sufficiently relativistic) massive particles can become trapped, due to the presence of photon spheres. One might ask, however, how any of this relates to the behaviour of the fermion wavefunction, since the fermions comprising Dirac solitons are typically of the order of the Planck mass, and hence one would expect their behaviour to be far removed from that of massless particles.

The key to understanding this is to note that, as mentioned, photon spheres occur at relatively small radii, within the power-law zone of solutions, where recall the fermion mass term in the equations of motion is negligible compared to the fermion energy ($m \ll \omega T$). In some sense, therefore, the fermion wavefunction can be considered ‘relativistic’ within the power-law zone. This can be understood by employing a semi-classical, WKB-type argument, in which it should be possible to approximate the form of the fermion wavefunction by considering the orbits of classical particles. Certainly, from a purely classical perspective, high-angular-momentum fermions that orbit at small radii must indeed be highly relativistic, and hence we would expect the structure of the wavefunction to reflect this. Thus we conclude that the fermion wavefunction within the power-law zone should be subject to the same trapping mechanisms as highly-relativistic classical fermions.

This trapping effect is illustrated in Fig. 3.6, where in the left-hand panel we have overlaid the radial co-ordinate in the optical geometry, $\rho = rT$, on top of the radial fermion energy density $r^2\rho_f$, for the highest redshift, 90-fermion state shown previously. From this, it is clear to see that the peaks in the fermion density line up almost precisely

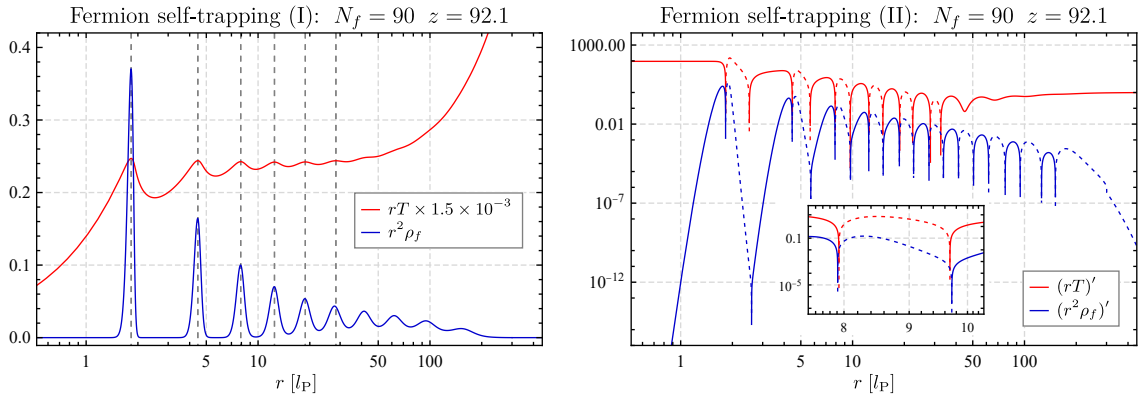


Figure 3.6: Illustration of the fermion-self trapping effect for a highly-relativistic, 90-fermion Einstein–Dirac soliton. Shown (**left**) is the radial profile of the radial fermion energy density $r^2\rho_f$, alongside the optical geometry radial co-ordinate $\rho = rT$. The dashed lines indicate the radii at which stable photon spheres occur. The derivatives of these functions are also shown (**right**), with the dashed portions of the curves corresponding to regions in which the derivatives are negative.

with the positions of the stable photon spheres in the optical geometry (the maxima in rT), at least for the first five oscillations. This strongly implies that indeed the fermion wavefunction has become trapped within the regions surrounding stable photon spheres, as suggested by the argument above. In addition, the unstable photon spheres (located at the minima of rT) produce an overall repulsive effect on the wavefunction, in common with the behaviour of classical null geodesics shown earlier in Fig. 3.5. Combined, these two factors can account for the observed splitting of the fermion wavefunction into a series of distinct peaks. This trapping effect can also explain the relative heights and widths of the energy density peaks, since more pronounced bulges in the optical geometry, which equate to more highly distorted regions of spacetime, result in deeper and more highly-confined potential wells. Since the depth of the bottleneck structures in the optical geometry decreases as we move outwards in radius, this then implies that the inner photon spheres should produce a stronger trapping effect, resulting in a larger proportion of fermions confined to a smaller region of spacetime, precisely as observed. Indeed, the trapping effect of the first two photon spheres is so strong that it causes the first two peaks in the fermion wavefunction to become almost entirely spatially separated from the remainder. This is not the case for subsequent peaks, however, since the associated potential wells are relatively shallow.

The precise correspondence between the fermion energy density peaks and the optical geometry can be seen more clearly in the right-hand panel of Fig. 3.6, where we plot together the radial derivatives of rT and $r^2\rho_f$. As can be evidenced, there is an excellent agreement between the extrema of these functions for the first five oscillations, even extending to the minima in the energy density. As we move outwards in radius, however, the correspondence becomes progressively worse, and indeed at large radii there exist oscillations in the fermion energy density that are not accompanied by an associated photon sphere. This can be understood by noting that the fermion wavefunction becomes progres-

sively less relativistic as the power-law zone is traversed, since ωT decreases monotonically with radius. Thus it is not surprising that the inner peaks of the energy density match the optical geometry more closely, since the effective potential experienced by highly-relativistic fermions resembles more closely that of null particles (as can be deduced from (3.16)). Towards the outer end of the power-law zone, as the wavefunction becomes less relativistic, the effective potential will begin to deviate significantly from the null case, to such an extent that in fact additional potential wells can appear, as shown in Fig. 3.5. This explains why the oscillations in the fermion energy density continue beyond those of the optical geometry, in agreement with the simple intuition that massive particles should be easier to trap than massless ones. Even at smaller radii, the agreement between the energy density and the optical geometry is not perfect, as can be seen from the inset shown in the right-hand panel of Fig. 3.6. We should not expect it to be, however, since the fermions are never strictly massless — what we would expect, however, is for the energy density peaks to occur slightly inside the stable photon spheres, and the minima to occur slightly outside their unstable counterparts, as argued in the previous section when analysing the effective potentials for massive geodesics. This indeed matches precisely what we observe.

Having introduced this trapping interpretation, it is important to reflect on the origins of the fermion wavefunction and the spacetime in which it has become trapped. In particular, it should be emphasised that here we have not first constructed a fixed spacetime (containing a series of photon spheres), and subsequently determined the form of the fermion wavefunction that responds to it — instead, both the metric and the fermion fields are solved for simultaneously. Thus the highly-distorted regions in which the fermion wavefunction becomes trapped are created solely by the energy density of the wavefunction itself, in a completely self-consistent way. It is therefore more appropriate to refer to this phenomenon as a ‘self-trapping’ effect. We also note that the discovery of this effect is heavily dependent on the accurate modelling of back-reaction, an ability only possible within semi-classical approaches of the type utilised here. Without the automatic inclusion of back-reaction, it would seem unlikely that structures such as those in evidence here could ever be obtained without prior knowledge of their existence.

Back-reaction is also crucial in understanding why the self-trapping effect is most prominent in systems containing large numbers of fermions. As mentioned, power-law oscillations are in fact present for Dirac solitons with any value of N_f , but it is only at large fermion number that these induce such extreme effects as wavefunction splitting. This is simply because a larger number of fermions produces a larger combined gravitational effect, and thus the back-reaction of the total wavefunction on the spacetime is greater, in turn creating more pronounced bottlenecks in the optical geometry. Hence for systems with small numbers of fermions, the spacetime distortion caused by the fermion energy density is insignificant, and thus the fermion wavefunction exhibits only small-amplitude oscillations, representing slight over and under-densities arising from the attractive (and

repulsive) effects of the photon spheres. From a purely mathematical perspective, this can be understood in terms of an increase in non-linearity of the FSY equations (2.28)–(2.31) as the fermion number is increased.

3.4 Families of many-fermion states

Following on from the above analysis of individual states, we now discuss the behaviour of the families of many-fermion Einstein–Dirac solitons, and demonstrate that the resulting spiralling structures are heavily influenced by the self-trapping effect. As before, our analysis is here restricted to $n = 0$ ground states, with excited states being addressed in Chapter 4.

3.4.1 Spiralling behaviour

As is the case for the two-fermion system, the set of many-fermion states with common N_f constitute a one-parameter family of solutions, parametrised by the central redshift z or equivalently the unscaled parameter $\tilde{\alpha}_1$. We similarly find that physical properties such as the fermion mass, fermion energy and ADM mass all oscillate towards their respective infinite-redshift values, and that as a consequence spiralling behaviour is observed. The visual appearance of these spiralling structures, however, differs significantly from the two-fermion case, as is illustrated for the $N_f = 90$ system in Fig. 3.7. In the upper panels we plot the fermion binding energy $m - \omega$ and the soliton radius R_{99} , both as a function of fermion mass. These can be compared with the equivalent two-fermion diagrams shown in Fig. 2.2. In addition, in the lower panels, we plot the fermion mass, fermion energy, ADM mass and soliton radius as a function of the central redshift. Overall, we find that the oscillations in redshift are of significantly larger amplitude than those seen in the two-fermion system, and as a result the spirals occupy a significantly larger area, to such an extent that the upper redshift limit of our numerics is still a considerable distance from the centre.

There are a number of features to note from these figures. Most importantly, there appear to be values of redshift at which all quantities exhibit a sudden, marked change in behaviour — m , ω and M all contain cusp-like points at the first few minima of their oscillatory periods, and these are accompanied by a discontinuous jump in the soliton radius R_{99} . The appearance of these can be understood by again invoking the self-trapping interpretation, as follows. At low values of redshift ($z < 1$), the states are non-relativistic, there is no power-law zone, and the fermion wavefunction contains only a single peak. In this regime, the curves therefore approximate the low-redshift relationships derived in Appendix A. As the redshift is increased beyond $z = 1$, however, the central compactness of the soliton increases to such an extent that a pair of photon spheres appears (at some critical value of redshift) at radii outside the bulk of the fermion source. As soon as

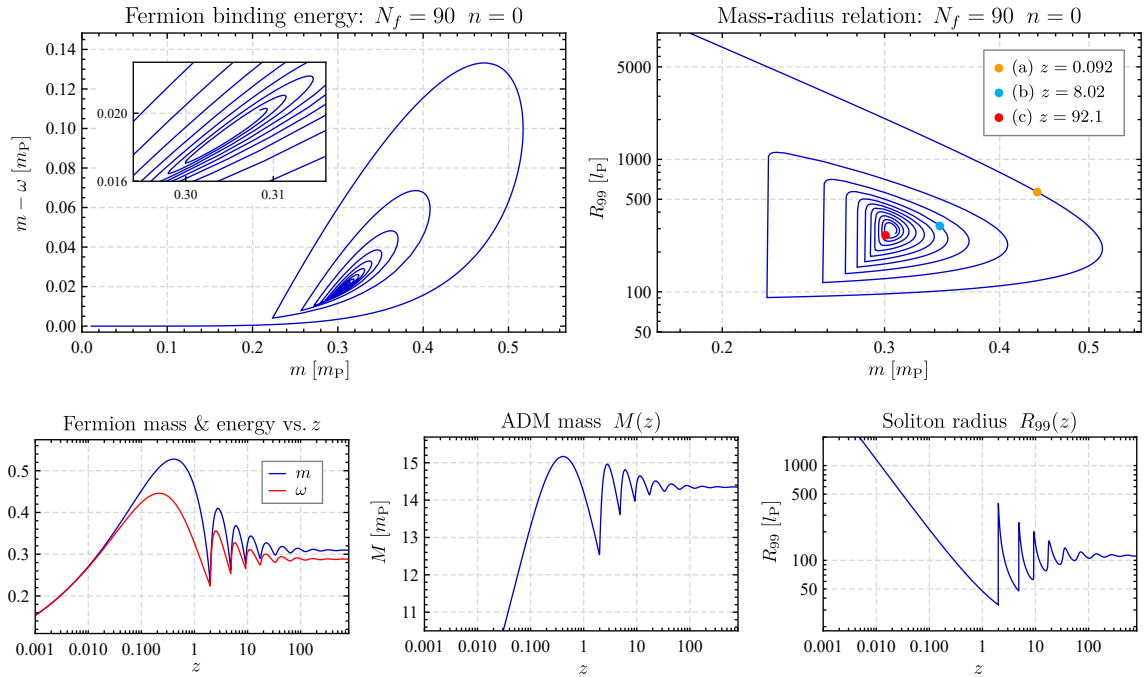


Figure 3.7: Plots summarising the behaviour of the family of ground-state 90-fermion Dirac solitons, showing (top) the spiralling nature of the fermion mass-binding energy and mass-radius relations. The locations of the three states presented in Fig. 3.1 are indicated on the latter curve. Also shown (bottom) are the the fermion mass m , fermion energy ω , ADM mass M and soliton radius R_{99} , all as a function of central redshift.

this occurs, the fermion wavefunction rearranges itself such that a small amount becomes trapped within the resulting bottleneck structure. The radius of the soliton therefore increases substantially (by a factor of over 100), to reflect this new doubly-peaked distribution of the fermion source. In addition, the redshift trajectories of the fermion mass, energy and ADM mass are also affected, since the subsequent evolution of the wavefunction is dominated by the gradual population of this new outer shell, as opposed to the continued development of the inner shell. As the redshift is increased further, a second pair of photon spheres materialise, causing this behaviour to repeat, and so on. Note that the jumps in radius become progressively smaller as the redshift increases, since each subsequent bottleneck is shallower than the last, resulting in a less confined outer trapping region. Ultimately, there will come a redshift beyond which new photon spheres begin to appear within the fermion source, as opposed to outside, and thus eventually the soliton radius becomes entirely continuous, and further cusps in m , ω and M no longer occur.

The evolution of the spiral curves as the fermion number is varied is illustrated in Fig. 3.8, where we plot the fermion binding energy and the mass-radius relations for the families of states with $N_f = 6$, 12 and 30, respectively. This clearly shows a gradual appearance of cusps and discontinuities, with the $N_f = 6$ curve appearing smooth, although still larger in extent than the $N_f = 2$ case, and the $N_f = 12$ curve exhibiting only a single jump in radius, before becoming continuous. Again this behaviour is related to

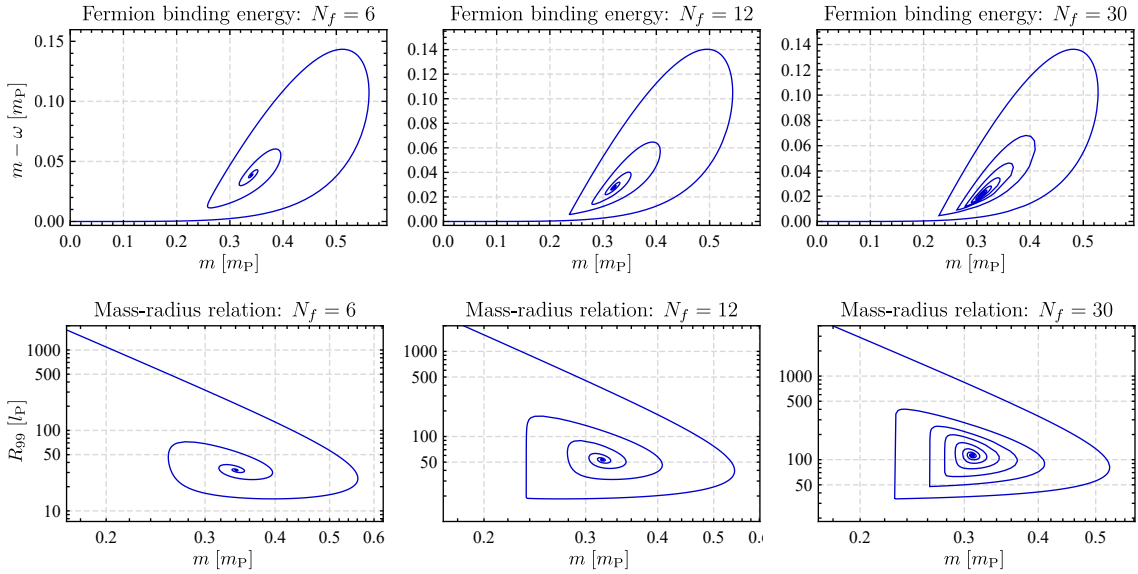


Figure 3.8: Plots showing the fermion binding energy (**top**) and mass-radius relations (**bottom**) for the families of ground states with N_f values of 6, 12 and 30.

the self-trapping effect. At small fermion number, all photon spheres materialise within the fermion source, and thus no discontinuities arise, whereas for intermediate numbers, it may be the case that only the first trapping region is spatially separated, and thus only the first oscillation is discontinuous. We note that these discontinuities are present partly due to employing R_{99} as our definition of soliton radius, i.e. the radius at which 99.9% of the ADM mass is first enclosed. With a more nuanced definition, such as \bar{R} , the radius becomes entirely continuous and instead cusp-like points appear, similar to those exhibited by the fermion mass and energy curves.

Given the importance of the fermion trapping to the visual appearance of these spiralling structures, it is perhaps tempting to conclude that the very existence of a spiralling behaviour is itself a direct consequence of the self-trapping effect. Indeed in ref. [4], it is asserted that, in the two-fermion system, the oscillations exhibited by physical observables are in fact caused by the oscillations of the fields themselves within the power-law zone. Since we have now provided an interpretation for these power-law oscillations, in the form of a fermion self-trapping effect, this does suggest a strong link to the spiralling behaviour. We explore this possibility further in Fig. 3.9, which shows the derivative of the fermion energy-redshift curves, $\omega'(z)$, for three different values of N_f . The redshift values at which new peaks in the radial fermion energy density first occur are then indicated by the dashed red lines. Considering first the $N_f = 90$ case, it can be clearly seen that, for $z > 1$, each oscillation in ω is indeed accompanied by the appearance of a new trapped region, in agreement with our earlier arguments. This is also the case for the first four oscillations in the $N_f = 30$ case, but once the curve become smooth, it appears that more than one oscillation completes before the appearance of each subsequent trapped region. This is substantiated by the $N_f = 6$ case (for which the curve is entirely smooth), where we find




Figure 3.9: An investigation of the relationship between the oscillatory nature of the fermion energy ω and the fermion self-trapping effect. These plots show the absolute value of the derivative of the fermion energy $d\omega/dz$ as a function of redshift, for the families of states with the N_f values indicated. The dashed red lines signify the values at which new trapped peaks in the radial fermion energy density first appear.

that new trapped regions appear approximately every one and a half oscillations in ω . How are we to interpret these results? On the one hand, they undoubtedly indicate a link between strong self-trapping effects and the properties of ω , but on the other, the periodicity of ω does not appear to match that of the self-trapping when the latter becomes weak, although they are indeed both evenly-spaced in redshift. We therefore tentatively conclude that the spiralling nature of Dirac solitons is not purely a consequence of the fermion self-trapping effect, although there is certainly a relation between the two.

3.4.2 Calculating the fermion energy

Having established a link between the physical properties of Dirac solitons and the fermion self-trapping effect, we now present a calculation through which an estimate of the fermion energy can be obtained by considering solely the form of the optical geometry. This is partly inspired by a similar calculation performed by Abramovich *et al.* [78], who successfully approximate the frequencies of gravitational wave modes that become trapped within the bottleneck structures associated with ultra-compact stars.

Our calculation relies on a semi-classical WKB-type argument, where we assume that the fermion wavefunction can be well approximated by considering the paths of classical particles, as discussed previously. In particular, we aim to relate the travel time of trapped null geodesics to the value of the fermion energy, noting that, for highly-relativistic fermions, this should differ only slightly from that of massive geodesics. As a first approximation, let us consider purely circular null geodesics (photon spheres), the equation of motion for which can be obtained by setting $ds^2 = dr^2 = 0$ in (2.3):

$$0 = -\frac{1}{T^2}dt^2 + r^2d\phi^2 \quad \implies \quad \frac{dt}{d\phi} = rT, \quad (3.17)$$

where we have confined the motion to the equatorial plane $\theta = \pi/2$, without loss of generality. Integrating over one complete orbit, we acquire a simple expression for the

travel time τ_c of a circular null geodesic:

$$\tau_c(r) = 2\pi r T(r). \quad (3.18)$$

Note that this is precisely equal to the distance travelled by circular paths in the optical geometry. For highly-relativistic states, multiple stable photon spheres can exist within the spacetime, and in those cases we should aim to obtain an averaged travel time, with an appropriate weighting applied to each path. In a true WKB analysis, this weighting is provided by the classical action, but here we choose instead to utilise the radial width of each trapping region, since this is observed to be roughly proportional to the amount of fermion energy that becomes trapped. Note that we are purposely avoiding using any explicit information from the fermionic sector, and thus we cannot use as a weighting the relative fermion density, for example. The mean travel time around the stable photon spheres within the optical geometry is then:

$$\tau_{\text{og,circ}} = \frac{\sum_n (r_n^+ - r_n^-) \tau_c(r_n)}{\sum_n (r_n^+ - r_n^-)}, \quad (3.19)$$

where r_n is the radius of the n^{th} stable photon sphere, and r_n^\pm are the outer boundaries of the n^{th} trapping region, calculated using the observed form of the null effective potential.

We now wish to relate this mean travel time to the energy of the fermion wavefunction. This can be achieved by considering, instead of a classical particle, a planar matter wave propagating purely in the ϕ -direction, with angular momentum j and energy ω_p , of the form $e^{i(j\phi - \omega_p t)}$. For this to reinforce constructively, the phase acquired after a single radial orbit must be set equal to that acquired after a temporal orbit τ , i.e. $\omega_p \tau = 2\pi j$. For Dirac solitons, $j = (N_f - 1)/2$, and hence a relation between travel time and energy can be extracted. This argument, however, arises from a non-relativistic, flat-space perspective, and we would not expect it to hold precisely in our circumstance. Given the lack of an obvious appropriate generalisation, we instead turn to the power-law solution (2.60), for which the following analytic relation between energy and travel time, $\tau = 2\pi r T$, exists:

$$\omega_{\text{pl}} = \frac{2\pi}{\tau} \sqrt{\frac{N_f^2}{4} - \frac{1}{3}} \equiv \frac{2\pi\xi}{\tau}. \quad (3.20)$$

We note that this approximates the form obtained from the argument above when the number of fermions is large. We are now in a position to express the fermion energy solely using information concerning the circular motion of null geodesics:

$$\omega_{\text{og,circ}} = \frac{2\pi\xi}{\tau_{\text{og,circ}}}. \quad (3.21)$$

A comparison between this expression and the true fermion energy is shown in Fig. 3.10, where we plot both as a function of redshift for the $N_f = 50$ fermion system, focusing on

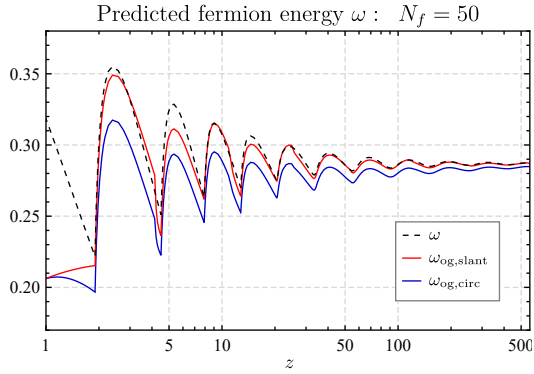


Figure 3.10: Predictions for the fermion energy, applied to the $N_f = 50$ system, obtained by analysing the travel time of null geodesics within the optical geometry. This shows the true fermion energy ω , as a function of redshift, alongside the approximations derived by first considering only circular orbits (blue) and subsequently including slanted orbits (red). Note that the curves do not agree at the low-redshift end of this plot, since this is prior to the formation of the first photon sphere.

the relativistic regime. From this, it can be seen that our approximation has successfully captured the desired oscillatory behaviour of ω , but the precise numerical values are consistently on the low side, suggesting that our estimate of the travel time is slightly too high.

With the aim of improving this approximation, we note that the above analysis equates to treating the fermion wavefunction as a series of delta functions located at the radii of stable photon spheres. This is clearly unrepresentative, since the self-trapping effect confines the wavefunction not to specific points but to regions surrounding them. To rectify this, we should therefore include geodesics that are non-circular, i.e. those that exhibit trapped ‘rosette’ orbits of the kind shown earlier in Fig. 3.5. We model these, somewhat crudely, as slanted orbits, for which the travel time is simply:

$$\tau_{\text{slant}}(r, r_c) = 2\pi \sqrt{r^2 T(r)^2 + (r - r_c)^2}, \quad (3.22)$$

where r_c is the radius of the associated photon sphere. Weighting each path equally, we can then calculate an average travel time for the null geodesics trapped around each bulge in the optical geometry:

$$\tau_{\text{bulge}}(r_c) = \frac{1}{(r^+ - r^-)} \int_{r^-}^{r^+} \tau_{\text{slant}}(r, r_c) dr. \quad (3.23)$$

As before, we then use the width of the trapping regions as the weighting factor to obtain a new estimate for the fermion energy, now including slanted, non-circular orbits:

$$\omega_{\text{og,slant}} = \frac{2\pi\xi \sum_n (r_n^+ - r_n^-)}{\sum_n (r_n^+ - r_n^-) \tau_{\text{bulge}}(r_n)}. \quad (3.24)$$

The application of this expression to the $N_f = 50$ system is also included in Fig. 3.10, resulting in a clear improvement compared to our earlier prediction. The numerical match to the true fermion energy is still not exact, but this is not surprising given the obvious shortcomings of our analysis. Indeed, since we have neglected the mass of the fermions, modelled non-circular geodesics only crudely, and implemented a somewhat ad-hoc weight-

ing system, the agreement is perhaps more accurate than one might have anticipated.

Regardless, it is clear that the principle behind the calculation is sound, i.e. the fermion energy can be obtained by analysing only the properties of the optical geometry, at least to a first approximation. We should not be overly surprised by this, since the fermion wavefunction and the spacetime metric are inherently coupled, in such a way that the spacetime should indeed respond precisely to the fermion energy density. What is perhaps more surprising, however, is the fact that states exist for which this calculation proves relatively simple, a characteristic that only arises due to the fermion self-trapping effect.

3.5 Discussion

Before closing this chapter, we briefly summarise the main results and provide a further discussion. Overall, we have shown that, at large fermion number, and sufficiently high redshift, the fermion wavefunction splits into a series of separated peaks, such that its appearance is more akin to a multiple-shell model. We have attributed this phenomenon to the appearance of a self-trapping effect, in which the fermion wavefunction becomes trapped within regions of spacetime surrounding stable photon spheres. This in turn affects the behaviour of physical observables, resulting in a marked difference in appearance of the binding energy and mass-radius spirals when compared with the two-fermion system. Although the self-trapping effect is indeed present for systems with any number of fermions, its consequences are most pronounced when the fermion number is large, due to the stronger back-reaction of the fermion energy on the spacetime metric.

One issue that we have not yet addressed is the precise origin of the self-trapping effect, i.e. why do the spacetimes of highly-relativistic Dirac solitons exhibit a series of photon spheres? The answer to this question is not entirely obvious. Certainly the appearance of a single photon sphere (or a single pair) can be understood by invoking the comparison with ultra-compact astrophysical objects. In these systems, a photon sphere is an inevitable consequence of the central compactness reaching a threshold value, due to the large spatial distortion caused by the object. The behaviour of Dirac solitons is clearly analogous, despite the obvious difference in scale — the central compression of the state is measured by the central redshift, and thus the appearance of a pair of photon spheres at some threshold value of z is entirely expected. The emergence of subsequent photon spheres, however, proves more difficult to explain. From an astrophysical perspective, objects with multiple photon spheres appear uncommon, and in fact there are proofs prohibiting such structures in static, vacuum spacetimes [84, 85], although recently hairy black holes containing three photon spheres have been discovered in the context of Einstein–Maxwell–scalar theory [86, 87]. Of more relevance here perhaps are the studies by Karlovini *et al.*

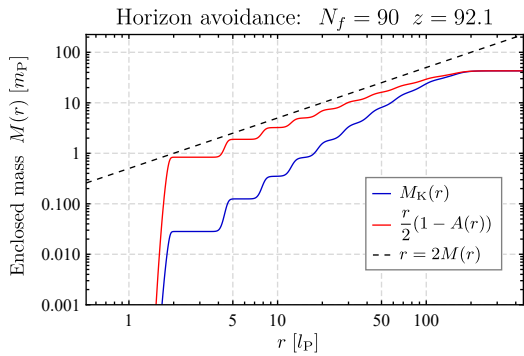


Figure 3.11: A possible indication of horizon avoidance for the $z = 92.1$, $N_f = 90$ state analysed previously. This shows the enclosed gravitational mass $M(r)$ as a function of radius, as obtained from the Komar mass (blue) and the ADM mass (red). We suspect the latter quantity to provide a more accurate measure. The dashed line $r = 2M(r)$ indicates the threshold above which the formation of a horizon proves inevitable.

[88, 89], who construct ultra-compact stars with an arbitrary number of photon spheres, using a physically reasonable equation of state, although no interpretation of the resulting structures is provided.

In the case of Dirac solitons, the appearance of multiple photon spheres may be related to the issue of horizon avoidance. It was shown in refs. [43–47] that systems of this type cannot contain horizons while continuing to remain static, and thus the central compression of states must be increased in such a way that the formation of a horizon is somehow avoided. We tentatively suggest that the emergence of photon spheres, and the associated fermion self-trapping, is the mechanism by which this is achieved. As evidence for this, consider Fig. 3.11, in which we plot the enclosed gravitational mass as a function of radius, for the highest-redshift $N_f = 90$ state shown in Fig. 3.1. Two measures for the enclosed mass are included — the Komar mass, defined in (2.47), and the enclosed ADM mass, which is obtained by calculating the deviation of A from unity. Given that the metric is of an approximate Schwarzschild form between the trapped peaks, it is likely that this latter quantity provides a more accurate measure of the enclosed gravitational mass in this case. Also included is the line $r = 2M(r)$, representing the threshold above which a horizon must inevitably form. As can be seen, at many points the enclosed ADM mass appears on course to overshoot this threshold, only for a sudden plateau to occur before the function rises once again. These plateaus correspond to the regions around the unstable photon spheres, from which the fermion wavefunction is effectively repelled. We therefore surmise that the division of the wavefunction into distinct trapped shells is precisely what prevents the enclosed gravitational mass becoming large enough in any one region for a horizon to form, effectively allowing the central density of the soliton to increase without limit, while ensuring static equilibrium is maintained.

It would be interesting to establish whether the appearance of multiple photon spheres in other systems is similarly accompanied by a feature resembling horizon avoidance. As mentioned, however, these are few and far between, although the models of ultra-compact stars studied by Karlovini *et al.* [88, 89] do exhibit other similarities to Dirac solitons. In particular, a ‘skeleton’ solution is shown to exist (the optical geometry of which is similar to that of our infinite-redshift states), towards which finite-density solutions spiral

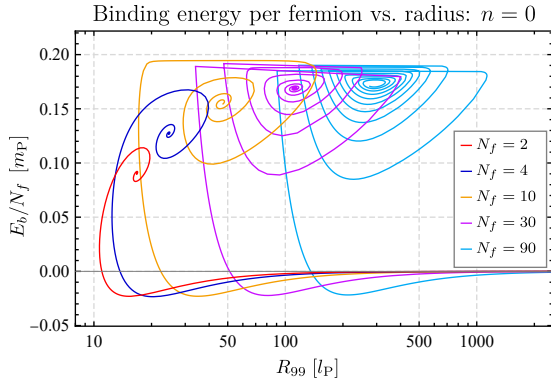


Figure 3.12: The overall binding energy per fermion as a function of radius, for the families of ground-state solutions with the N_f values indicated. In all cases, we would expect stable solutions to be those with central redshift less than that of the state with minimum binding energy. Thus high-redshift states, in which the fermion self-trapping effect is in evidence, are expected to be dynamically unstable, since these are located along the spiralling portions of the curves.

as the number of photon spheres is increased. This lends further evidence to the claim that the spiralling nature exhibited by objects of this type is related to the appearance of photon spheres, and by extension, at least in the Einstein–Dirac system, the fermion self-trapping effect. A possible physical interpretation for this spiralling behaviour is detailed as follows. Recall that one of the surprising aspects associated with Dirac solitons is that, for relativistic states, there can exist multiple ground-state solutions that have the same value of the fermion mass. This is contrary to our naive Newtonian intuition that there should exist a unique equilibrium configuration for a collection of gravitationally interacting fermions of mass m , as proves to be the case in the non-relativistic limit. In general relativity, however, the presence of photon spheres permits the possibility of alternative configurations in which the fermions are trapped within the associated bottleneck structures, but equilibrium is still maintained. Thus one can obtain states for which the fermion mass is identical, but the properties of which differ depending on the number of photon spheres present. It remains to be seen whether this argument can be extended to other systems, or whether the appearance of photon spheres is a common feature of spiralling, self-gravitating systems. The most obvious candidates for testing these hypotheses are boson stars, and indeed recent results obtained for ℓ -boson stars [90] show indications of a similar self-trapping effect to that discovered here.

We finish by briefly addressing the issue of stability. It is important to point out that, despite the perhaps misleading terminology of ‘trapping’, all states exhibiting the self-trapping effect have an overall positive binding energy, and thus we would in fact expect them to be unstable to perturbations. This is illustrated in Fig. 3.12, where we plot the binding energy per fermion as a function of soliton radius. Thus the self-trapping effect is arguably of more academic interest, having little relevance to the application of constructing stable particle-like objects. The possibility remains, however, that the inclusion of additional interactions could act to stabilise states in which the self-trapping is in evidence. It is also likely that the dynamical evolution of unstable solutions is somewhat influenced by its presence.

Chapter 4

Many-fermion excited states

In the previous chapter, we discussed in detail the behaviour of ground-state solutions to the many-fermion Einstein–Dirac system. Here, we shall extend this analysis to the corresponding excited states, i.e. those with multiple nodes in the fermion wavefunction. As we shall show, multivaluedness becomes an issue even for relatively small fermion numbers ($N_f \geq 6$), with the central redshift no longer uniquely identifying excited states in the relativistic regime. Within these multivalued regions, it is observed that the physical properties of states are often very similar to those of lower excited states, a feature that can be explained by analysing their internal structure. In what follows, we shall first briefly discuss the numerical challenges that arise when generating excited-state solutions, before presenting results illustrating the behaviour of the first excited states as the fermion number is varied. An analysis of higher excited states then follows, which are shown to display an even larger multiplicity. Note that we shall again restrict our analysis to states of even parity (even values of n), with the expectation that their odd-parity counterparts will exhibit a similar behaviour.

4.1 Numerical method

4.1.1 Review of the two-fermion case

Before discussing the complications that arise at large fermion number, we first briefly recap the numerical method used to obtain states in the two-fermion Einstein–Dirac system. Since this was previously described in detail in section 2.2, we shall here only highlight the most relevant aspects.

First, recall that directly imposing the boundary conditions of asymptotic flatness and normalisation proves numerically challenging, and thus a rescaling technique is implemented, in which we first seek ‘unscaled’ solutions to the system (denoted by an additional tilde), from which the physical states can subsequently be obtained by rescaling the fields

and parameters as appropriate. For each value of the unscaled parameter $\tilde{\alpha}_1$, there exists a unique ground state with unscaled fermion energy $\tilde{\omega}_0$, preceded by a (presumably infinite) tower of excited states, each with a higher value of $\tilde{\omega}_n$ than the previous. States are distinguished by the total number of fermion nodes n in the fields $\tilde{\alpha}$ and $\tilde{\beta}$, which, for even-parity states, is required to be even. Since each value of n uniquely identifies each state, it is relatively straightforward to obtain the precise values of $\tilde{\omega}_n$, for example by performing a simple binary chop in which solutions with the desired number of fermion nodes are sought.

Upon rescaling, the nodal structure of solutions is preserved, but the values of the physical fermion mass m and energy ω differ from those in the unscaled system. In addition, the physical parameter α_1 no longer uniquely identifies each state, and instead the appropriate parameter to use for this purpose is the central redshift z . One-parameter families of solutions can then be generated, one for each value of n , by varying the value of the central redshift from zero to infinity. Since $\tilde{\alpha}_1$ is found to be in one-to-one correspondence with z , the families could equivalently be parametrised by the former quantity, although it has no physical interpretation. A summary of the two-fermion system can be seen in Fig. 4.1, which shows the behaviour of the ground and first three (even-parity) excited states. Note that, even after the rescaling, the fermion energy-redshift curves for the families of states remain well-separated, i.e. at each redshift, the state with n nodes is always of higher energy than the corresponding state with $n - 2$ nodes. This will cease to be the case when considering systems with larger numbers of fermions. The mass-radius

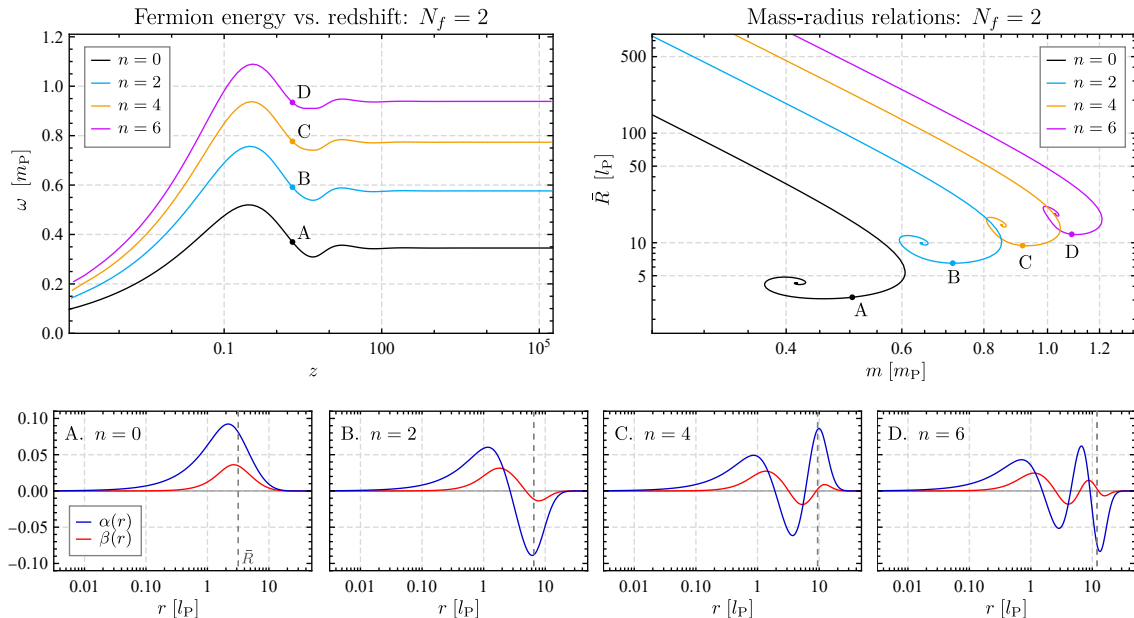


Figure 4.1: Summary of the two-fermion system, showing (**top left**) the fermion energy as a function of redshift, alongside (**top right**) the mass-radius relations, for the families of ground and first three even-parity excited states. Include also (**bottom**) are the radial profiles of the fermion fields for the states located at the points A–D, each having a redshift value of 2.00.

relations are similarly distinct, with the observed general increase in radius with n a consequence of accommodating the additional fermion nodes. Note that here, and throughout this chapter, we use the averaged radius \bar{R} as the primary measure of the radial extent of solutions, since this quantity always remain continuous in z .

Also shown in Fig. 4.1 are examples of four individual solutions, corresponding to the ground and first three even-parity excited states that occur at a redshift value of $z = 2.00$. The internal structure of these is as described in section 2.5, with a clear separation into four distinct zones: the core, the power-law zone, the wave-zone and the evanescent zone. Although too small to be visible on the plots shown, the fields exhibit oscillations within the power-law zone, the appearance of which is attributed to a fermion self-trapping effect, as discussed. These have long since decayed by the time the wave zone is reached, at which point the fermion fields perform large-amplitude oscillations that ultimately define the value of n . As we shall see, the situation is altered somewhat at large fermion number, where we find that fermion field nodes are no longer confined simply to the wave zone, leading to a reassessment of the previously discovered zonal structure.

4.1.2 The many-fermion system

The behaviour described above extends also to the $N_f = 4$ system, as well as being valid in the non-relativistic regime ($z < 1$), regardless of the fermion number. From $N_f = 6$ onwards, however, we find that neither $\tilde{\alpha}_1$ nor z can be used to uniquely identify all states in the relativistic regime, i.e. multiple states (with the same n) may be found with identical values of these parameters. Note that this multiplicity applies only to excited states ($n \geq 2$), and there always remains a unique $n = 0$ ground state for each value of $\tilde{\alpha}_1$ and z . Thus an identical procedure to the two-fermion case can be used to obtain ground state solutions, as detailed in the previous chapter. The same method cannot be used, however, to generate excited-state solutions, since multiple states may exist with the same value of n , making a simple binary chop unworkable.

The severity of the problem is illustrated in Fig. 4.2, where we show the unscaled fermion energy spectra for two values of $\tilde{\alpha}_1$, one non-relativistic (left) and the other relativistic (right), for the $N_f = 40$ system. These plots are to be interpreted as follows. The x -axis shows various values of $\tilde{\omega}$, while the y -axis indicates the number of fermion nodes n obtained when solving the FSY equations (2.28)–(2.31) using each $\tilde{\omega}$ value. Note that normalisable solutions, corresponding to physical states after rescaling, are located only at specific values of $\tilde{\omega}$; for the remainder (and vast majority), the fermion fields ultimately diverge to infinity at large r . Also note that the x -axes here are not linear in $\tilde{\omega}$ (although they are monotonic), and instead simply represent a sequence of unevenly-spaced $\tilde{\omega}$ values. Considering first the non-relativistic case, we observe the familiar behaviour present in the two-fermion system — a gradual monotonic increase in the number of fermion nodes as $\tilde{\omega}$ is increased, with this continuing indefinitely, although here we restrict to $n < 10$. The

locations of the five normalisable states contained within this region (one on each plateau) are indicated on the x -axis, with their values given in the caption.

Turning to the relativistic case, we observe an entirely different, and much more complex, behaviour. Instead of a single tower of states, we now appear to have a series of ascending and descending sequences, where in each there exist a presumably infinite number of excited states (although again we have capped the figure at $n = 9$), contained within a finite $\tilde{\omega}$ range. Here, a normalisable state is located along every even- n plateau, and thus there are a total of 37 individual states depicted in the figure. Note that, although each upward sequence of states extends to $n = \infty$, the downward sequences do not all descend to the same value of n . In fact none of them return all the way to $n = 0$ (and hence the ground state remains unique), with only two returning to $n = 1$. The final sequence, on the far right of the plot, increases indefinitely but never returns, reaching $n = \infty$ at strictly infinite $\tilde{\omega}$. We thus conclude that there exist a total of five $n = 2$ states for this value of $\tilde{\alpha}_1$, but nine $n = 4$ states, and eleven $n = 6$ and $n = 8$ states.

Given the evident high degree of multivaluedness present at large fermion number, how does one go about determining the energies of the excited states, i.e. how does one obtain the energy spectrum shown in Fig. 4.2. As mentioned, a binary chop cannot be used, although this can be employed once the energy spectrum has been obtained, and unfortunately the sequences of states tend never to be particularly well-separated in $\tilde{\omega}$. Indeed, we find that entire sequences often exist within tiny regions of $\tilde{\omega}$, as indicated by the values of $\tilde{\omega}_{a-e}$ given in the caption of the figure. With no other clear options, we are therefore forced to perform a thorough search of $\tilde{\omega}$ -space, starting with an initial sweep of 500 $\tilde{\omega}$ values ranging from the ground state $\tilde{\omega}_0$ to $1.05\tilde{\omega}_0$, with a similar, more focused sweep implemented once the final sequence of states has been identified. Thereafter,

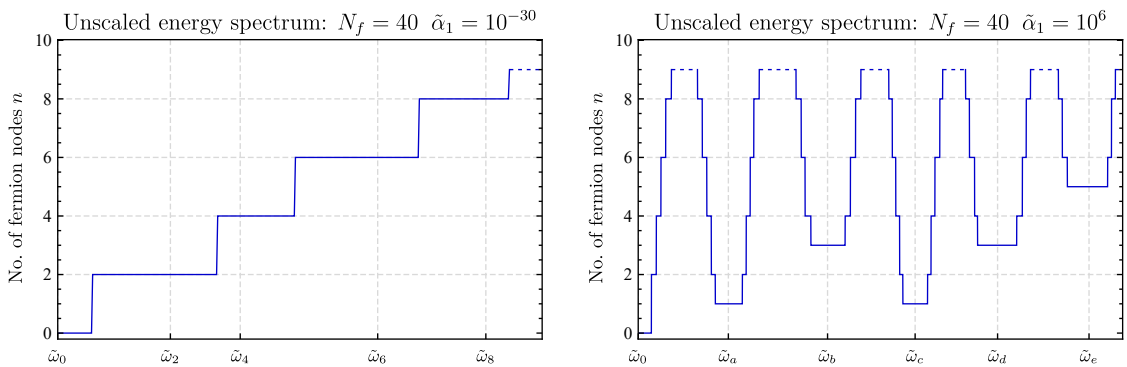


Figure 4.2: The extent of multivaluedness in the 40-fermion system. Shown are the unscaled energy spectra, in which $\tilde{\omega}$ is plotted against the number of fermion nodes n , for **(left)** a non-relativistic value of $\tilde{\alpha}_1 = 10^{-30}$ and **(right)** a relativistic value $\tilde{\alpha}_1 = 10^6$. In the left-hand panel, the values of $\tilde{\omega}$ for the ground and first four even-parity excited states are indicated, these being $\tilde{\omega}_0 = 1.1568$, $\tilde{\omega}_2 = 1.1601$, $\tilde{\omega}_4 = 1.1617$, $\tilde{\omega}_6 = 1.1628$ and $\tilde{\omega}_8 = 1.1637$. In the right-hand panel, $\tilde{\omega}_0 = 40.1196$ and the points $a-e$ are (relative to the ground state): $\tilde{\omega}_a - \tilde{\omega}_0 = 1.434 \times 10^{-7}$, $\tilde{\omega}_b - \tilde{\omega}_0 = 8.575 \times 10^{-6}$, $\tilde{\omega}_c - \tilde{\omega}_0 = 1.037 \times 10^{-5}$, $\tilde{\omega}_d - \tilde{\omega}_0 = 0.3141$ and $\tilde{\omega}_e - \tilde{\omega}_0 = 0.3897$.

various methods are used to ensure that all sequences have been identified. These include, for example, sweeping regions in which a jump is observed in the number of fermion nodes, expanding the regions around observed peaks in the radius at which the final fermion node occurs, or simply searching close to already discovered solutions. This procedure can of course be automated, again using MATHEMATICA, although on occasion a final manual search is required if the number of solutions discovered is fewer than expected. Once the complete energy spectrum is obtained, the precise values of $\tilde{\omega}$ for each state can then be determined by implementing a binary chop within the plateaus in which states are known to exist. These unscaled solutions are then converted into physical states using the rescaling method outlined in section 2.2.2.

Clearly, generating solutions using the method described above comes at a much greater computational cost than a simply binary chop. One way this can be reduced is to search only for solutions in which the number of fermion nodes is less than a specified threshold. Here, as indicated in Fig. 4.2, we set this limit at $n = 9$, allowing us to obtain up to and including fourth even-parity excited states ($n = 8$). Even with this restriction, however, the maximum fermion number that we can obtain is significantly less than in Chapter 3, recalling that as the fermion number increases so must the numerical precision. We are therefore limited, when considering entire families of solutions, to $N_f < 40$, although if the energy spectrum at only a single value of $\tilde{\alpha}_1$ is required, this can be increased to $N_f \approx 70$.

As a final remark, one might be concerned as to whether this multivaluedness is a feature solely of the unscaled equation system, and in fact disappears when considering rescaled, physical states. We emphasise that this is not the case, however. Indeed, $\tilde{\alpha}_1$ and z remain in one-to-one correspondence, and hence the same multivaluedness that appears in $\tilde{\alpha}_1$ is also present in the central redshift.

4.2 1st excited states ($n = 2$)

Having outlined the numerical method used to generate solutions, we now present results illustrating the behaviour of the excited-state many-fermion system as the number of fermions is varied. We shall begin with an analysis of the first even-parity excited states ($n = 2$), before proceeding to discuss higher excited states in section 4.3.

4.2.1 Varying N_f

The evolution of the fermion energy-redshift curves, as we gradually increase the fermion number from $N_f = 4$ to $N_f = 14$ is illustrated in Fig. 4.3. These can be compared with the corresponding plot for the two-fermion system shown previously in Fig. 4.1. Note that here we focus solely on the relativistic regime ($z > 1$), since, as mentioned, the non-

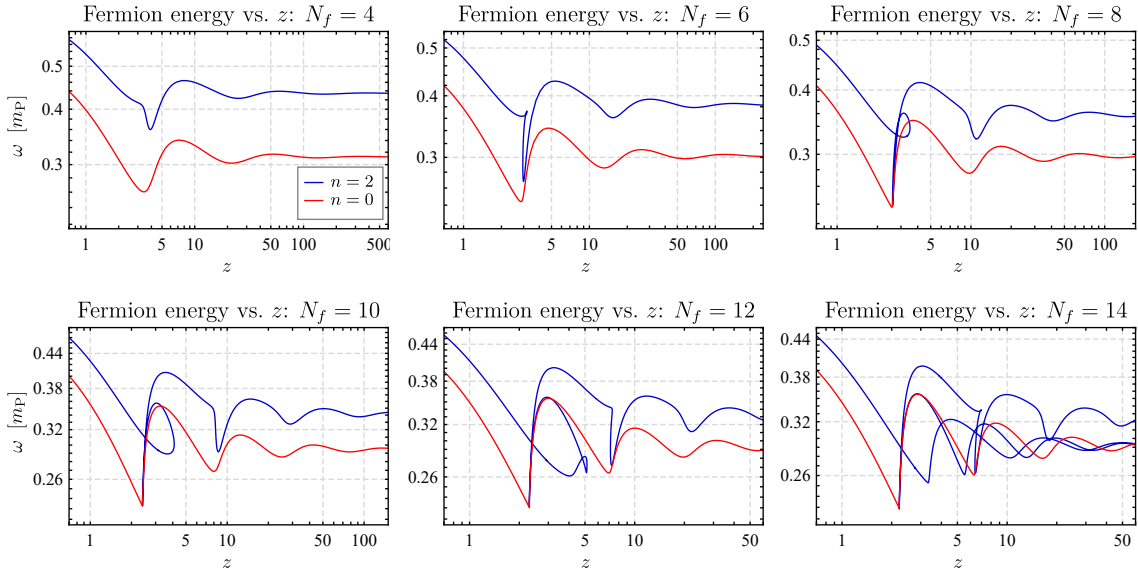


Figure 4.3: Plots summarising the behaviour of the first even-parity excited states as the fermion number is varied, showing the fermion energy as a function of redshift for the families of $n = 0$ and $n = 2$ states, for the values of N_f indicated. Note the gradual emergence of multivaluedness as the fermion number is increased.

relativistic behaviour is identical to that observed for $N_f = 2$. The families of ground-state solutions (shown in red) exhibit the characteristics discussed in Chapter 3, with the first minimum of the curve gradually evolving into a cusp as the fermion number is increased, as a consequence of the increased strength of the fermion self-trapping effect. The behaviour of the first excited states is markedly different, however. Even at $N_f = 4$, we observe a slight distortion around the first oscillation in the $n = 2$ energy-redshift curve, with this developing into a multivalued region (or ‘fold’) as N_f increases. Within these regions, the curves initially reverse direction at some value of redshift, and subsequently approach close to the first minima of the corresponding ground-state curves, before reversing direction once again and continuing to infinite redshift. The extent of this multivalued portion is greater at larger fermion number, and indeed at $N_f = 14$ the reversal occurs at a redshift value beyond the limit of our numerics. In addition, from $N_f = 12$ onwards, a second fold begins to emerge, with a similar structure to the first, but located around the second minimum in the ground-state curve.

It is important to note that, despite this multivalued nature, the energy-redshift curves remain entirely continuous for all values of N_f . Thus they still represent one-parameter families of solutions, but the appropriate parameter to use is no longer the central redshift. We shall discuss an attempt at obtaining a single-valued parametrisation of these curves in section 4.4. Note also that, for $N_f \geq 14$, we cannot be certain that the $n = 2$ energy-redshift curves reverse direction at a strictly finite value of redshift; it may be the case that they extend all the way to $z = \infty$ before eventually reversing. An argument against this is to note that the intermediate portion of each $n = 2$ curve appears to perform damped

oscillations around the corresponding ground-state solution, and thus if this portion were to extend to $z = \infty$, its endpoint would ultimately coincide with the infinite-redshift $n = 0$ solution. This would imply that there exists a degeneracy associated with the infinite-redshift states, a feature that we do not detect during their numerical generation. This argument is not conclusive, however, since it may be the case, for example, that a perturbation analysis is required in order to reveal a previously hidden degeneracy. Whether the curves do indeed reverse at finite redshift therefore remains an open question.

The corresponding behaviour of the mass-radius relations, for the same values of N_f , is illustrated in Fig. 4.4. Here we observe that the family of $n = 2$ states gradually begins to wrap around the ground-state family as the fermion number increases, with the region in which this occurs corresponding to the multivalued portion of the energy-redshift curves. One end of this region becomes fixed at the first sharp turning point of the ground-state curve, whereas the other moves progressively inwards, nearing the centre of the ground-state spiral. Indeed, at $N_f = 14$, it appears (at least visually) that the $n = 2$ curve completes two spirals, the first following the $n = 0$ relation, and the second around the $n = 2$ infinite-redshift state. As mentioned, however, we cannot be certain whether the $n = 2$ curve truly reaches the centre of the ground-state spiral.

4.2.2 The $N_f = 12$ system: a closer look

There are a number of obvious questions that immediately arise concerning the behaviour outlined in the preceding section. In particular, why do the $n = 2$ curves contain multivalued portions, and why do the properties of states within these regions so closely resemble

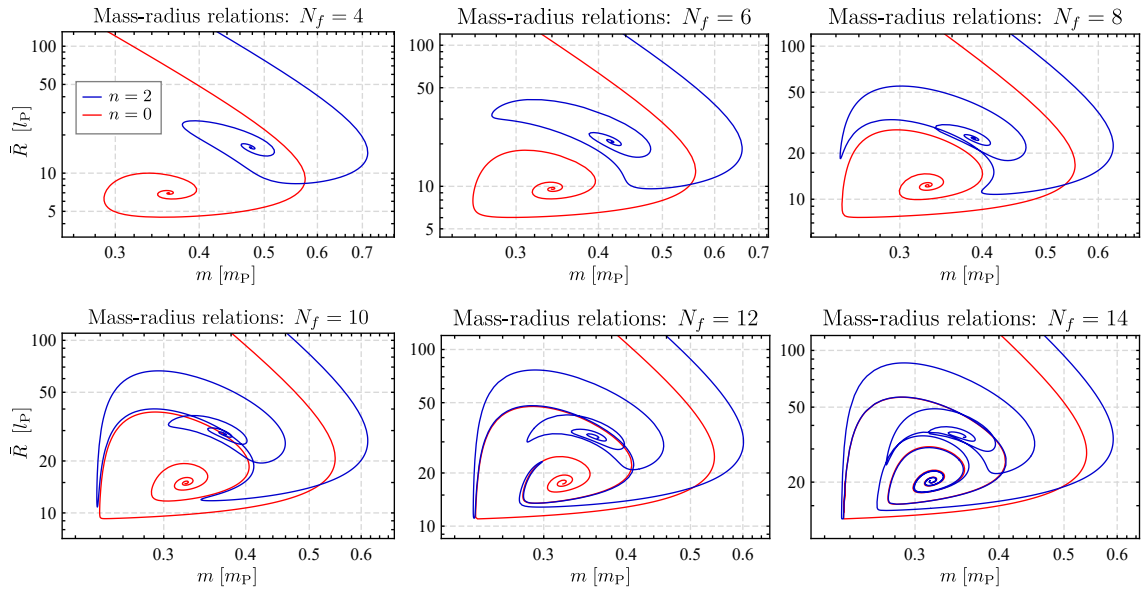


Figure 4.4: Plots illustrating the change in behaviour of the mass-radius relations for the family of first even-parity excited states ($n = 2$), as the fermion number is increased. Note the appearance of a distortion in the $n = 2$ curve that gradually wraps itself around the ground-state spiral.

those of ground-state solutions? The answers to these questions can at least partially be obtained by analysing the internal structure of individual solutions along the curves. In this section, we therefore present a more detailed analysis of the family of $N_f = 12$ states, focusing on the multivalued region.

The results of this analysis are summarised in Fig. 4.5, where we show the radial profiles of the fermion fields α and β for twelve individual states located in and around the multivalued region. The precise positions of these are indicated on the fermion energy-redshift and mass-radius curves. This figure allows us to track the evolution of the fermion wavefunction as the fold in the energy-redshift curve is traversed. The general behaviour can be described as follows. Entering from the low-redshift regime, solution A is encountered first, which exhibits the structure expected from a non-relativistic $n = 2$ state, with a single node in each fermion field separating a pair of extrema. As we move further along

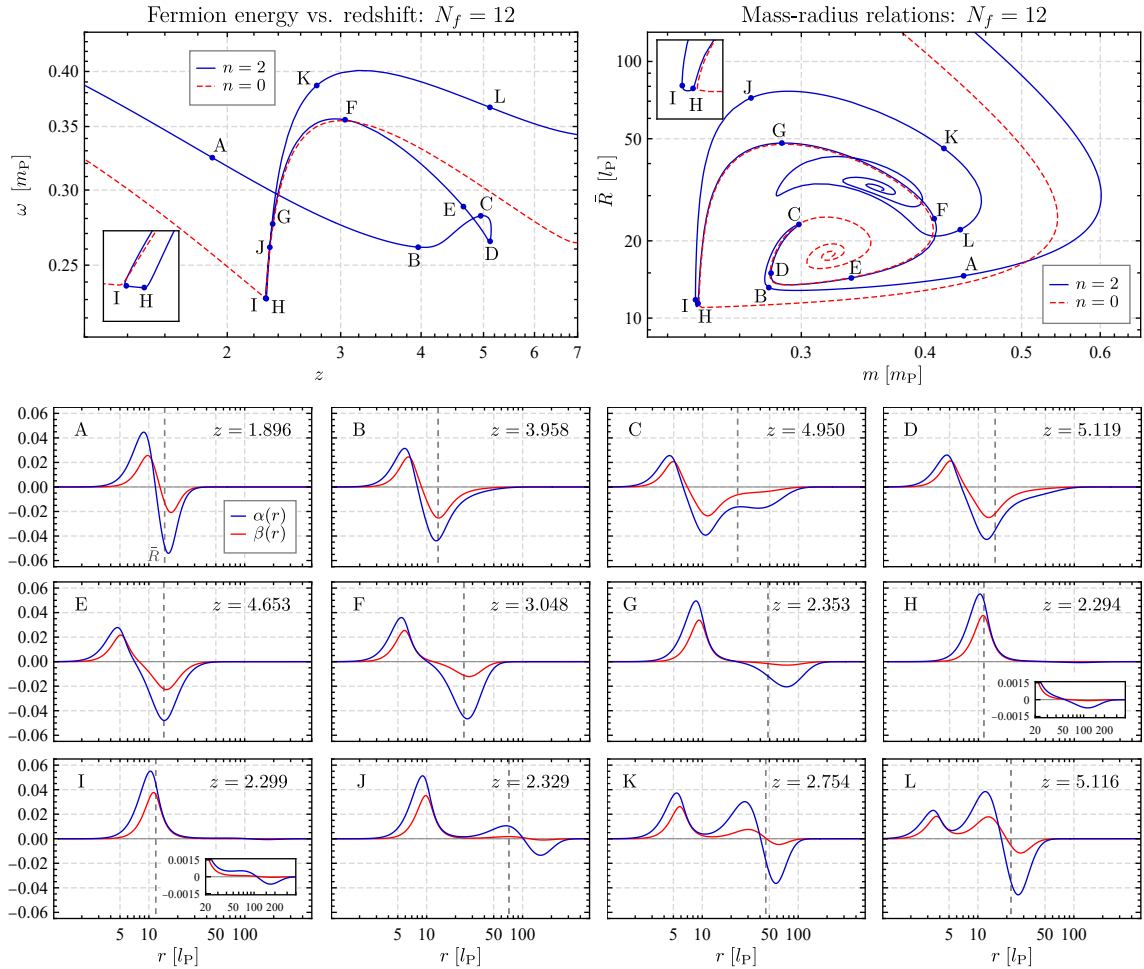


Figure 4.5: A closer look at the behaviour of the $N_f = 12$ system, detailing the evolution of the $n = 2$ fermion wavefunction as the multivalued region is traversed. Shown (**top**) are the fermion energy-redshift and mass-radius relations for the families of $n = 0$ and $n = 2$ states, with the former restricted to the multivalued portion. Shown (**bottom**) are the fermion field profiles for 12 individual $n = 2$ states, with the locations of these indicated on the curves above.

the curve (into the multivalued region), a secondary minimum begins to form, located outside the fermion nodes (i.e. at a larger radius), with the amplitude of this reaching a peak around solution C. This minimum subsequently subsides and in fact disappears entirely by the time solution E is reached, where once again only two extrema in each field are present. Solutions F–H are located along the portion of the curve that travels backwards in redshift, and following these cases, we see that the maxima located inside the fermion nodes (i.e. those at smaller radii) gradually grow in amplitude, while the minima outside shrink. In fact, in solution H, located just prior to the second redshift reversal point, the minima in the fermion fields are almost invisible. These do not, however, disappear entirely, and instead, once the curve has reversed in redshift, a new fermion field peak begins to form at a radius inside the fermion nodes (see solution I). As we now move forwards in redshift, this additional peak increases in amplitude, and upon exiting the multivalued region, the wavefunction structure is as shown in solution L.

In order to interpret this evolution, it is useful first to recap the corresponding behaviour of the two-fermion system. There, non-relativistic states (resembling solution A) evolve gradually into high-redshift states via the addition of successive fermion field oscillations within the power-law zone. The evolution would structurally resemble the sequence $A \rightarrow J \rightarrow K \rightarrow L$, with new peaks forming inside the radius of the fermion field nodes. In the $N_f = 12$ case, however, we note that solutions exist, such as C, in which an additional peak is present *outside* the fermion nodes, a feature never observed in the two-fermion system. In such cases, it appears that the fermion nodes are located within the power-law zone, and not within the expected wave zone; in fact the wave zone here ceases to exist at all. This will become clearer when we consider the $N_f = 20$ system in section 4.2.3.

It is important to note that the evolution of the fermion fields shown in Fig. 4.5 is entirely continuous, i.e. the wavefunction is deformed through solutions A–L in a continuous manner, ensuring a single node in each fermion field is always present. Along this process, there exist solutions (e.g. H and I) in which the oscillations around the fermion nodes are of very small, but non-vanishing, amplitude, and hence the resulting wavefunctions visually resemble ground-state solutions. This explains why the properties of some $n = 2$ states are so similar to their ground-state counterparts. Despite this, we should undoubtedly still refer to these as $n = 2$ states, since they contain a single node in each fermion field, and are continuously connected to the more-conventionally structured $n = 2$ solutions. Whether the terminology of ‘excited state’ remains appropriate, however, is a more debatable issue, given that solutions such as B–E have a lower fermion energy than the ground-state solution with corresponding redshift. Nonetheless, we shall continue to refer to states with $n \geq 2$ fermion nodes as ‘excited’, with this caveat acknowledged.

4.2.3 The $N_f = 20$ system

A clearer appreciation of the situation can be obtained by considering the structure of high-redshift states. To do so, we are required to increase the fermion number such that the multivalued portion of the fermion energy-redshift curve extends to a sufficiently high redshift. From Fig. 4.3, we note that this first occurs at $N_f = 14$, but here we shall consider the $N_f = 20$ system, since the slightly stronger self-trapping effect will aid the analysis.

A summary of the behaviour of the $N_f = 20$ system is presented in Fig. 4.6, where we plot the fermion energy-redshift and mass-radius relations for the families of $n = 0$ and $n = 2$ states, together with four individual solutions located at high redshift. This is again a case in which the intermediate portion of the energy-redshift curve extends beyond the limit of our numerics, and hence there are always at least three $n = 2$ states within the relativistic regime. Since the plots depicting the families of states are somewhat convoluted, we have separated the $n = 2$ curves colour-wise into three distinct sections, as follows. The light-blue portion of the energy-redshift curve enters from the non-relativistic regime and oscillates towards the $n = 0$ infinite-redshift state. The orange section then indicates the part of the curve that subsequently travels backwards in redshift, with the purple portion occurring beyond the redshift reversal at $z = 2.1$. This colour-coding is similarly reproduced for the $n = 2$ mass-radius relation, which as before is seen to exhibit

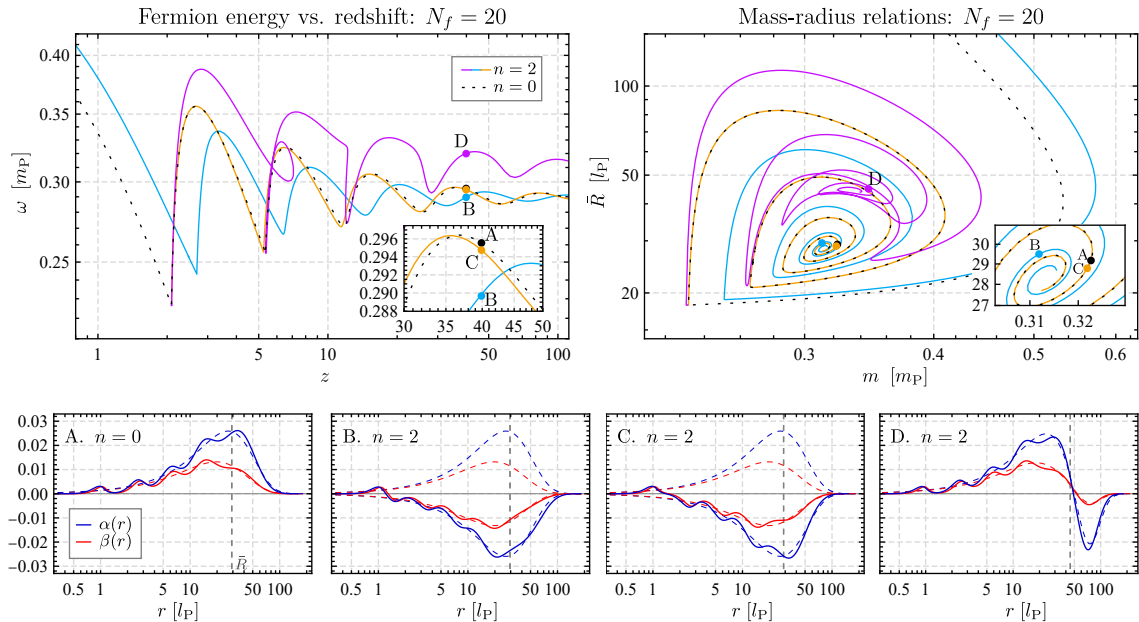


Figure 4.6: Plots summarising the behaviour of the $N_f = 20$ system, showing (**top**) the fermion energy-redshift and mass-radius relations for the families of ground and first even-parity excited states. These are colour-coded to indicate the respective incoming and outgoing branches. Also included (**bottom**) are four individual states (one $n = 0$ and three $n = 2$), each with redshift $z \approx 40$, the positions of which are located on the curves above. The dashed lines represent the radial profiles of the infinite-redshift states around which the solutions oscillate.

two spirals, towards both the $n = 0$ and $n = 2$ infinite-redshift states.

Considering the structure of the individual solutions shown in Fig. 4.6, we see first that the ground state (A) now contains an extended power-law zone, with a series of fermion field oscillations around the $n = 0$ infinite-redshift state (indicated by the dashed curves) in evidence. Note that, due to the strong self-trapping effect, the first minimum in both α and β is very close to zero. Turning to the $n = 2$ states, solution D, which is located along the final outgoing portion of the energy-redshift curve, exhibits the standard behaviour expected for an $n = 2$ state, with a single fermion node in each field located outside the power-law zone. In solutions B and C, however, the fermion nodes clearly occur within the power-law zone, just outside the first trapped peak, with the remainder of the solution now oscillating around the negative version of the infinite-redshift state. These two solutions are located along the intermediate portion of the energy-redshift curve, on the outgoing and incoming branches respectively, and are close to the centre of the $n = 0$ mass-radius spiral.

Why do the properties of solutions B and C so closely resemble those of ground states? To understand this, we first note that the majority of the fermion probability density is contained within the outer regions of the soliton, towards the end of the power-law zone and within the wave and evanescent zones. Thus the properties of a state are determined primarily by the form of the wavefunction at large r . In addition, physical observables are calculated solely using bilinears involving α and β , and are hence entirely insensitive to the sign of the fermion fields. With this information, it becomes clear that the properties of a solution are overwhelmingly dictated by the number of fermion nodes located within the wave zone. Since solutions B and C have zero such nodes (in fact the wave zone does not exist), it is therefore not surprising that their physical properties are similar to those of ground states. Finally, it is worth noting that, although B and C are located along different branches of the energy-redshift curve, there appear to be no discernable structural differences between the two solutions. We find that this is indeed always the case — states on the outgoing and incoming portions of the same fold are identical in nodal structure.

The structural evolution of states can be discerned by considering the radii at which fermion nodes occur, as a function of redshift. This is shown in Fig. 4.7, in which we have colour-coded the curves in the same manner as before. The evolution proceeds as follows. Entering from the non-relativistic regime, the fermion field nodes initially travel radially inwards (along the light blue branches), before reversing at some (possibly infinite) value of redshift, and subsequently moving outwards (along the orange branches). The states along these portions of the curve are those in which the fermion nodes are located within the power-law zone, and thus the nodes are carried radially inwards as the redshift is increased, along with the inner boundary of the zone. Note that, for states on the incoming branch (orange), the fermion nodes consistently occur at a slightly larger radius than for those on

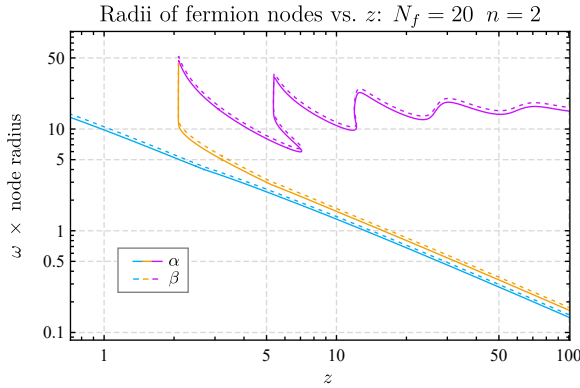


Figure 4.7: The radii of fermion nodes, as a function of redshift, for the family of first even-parity excited states ($n = 2$) that exists in the 20-fermion system. This should be interpreted as one continuous curve, separated into three colour-coded branches as in Fig. 4.6. Note that the radius of the node in α always occurs slightly inside the corresponding node in β . Note also that we have multiplied the node radii by a factor of ω in order to mask the oscillatory features present in the power-law zone.

the outgoing branch (light blue), a feature that is reflected in the structures of solutions B and C in Fig. 4.6. Once the curve returns to $z = 2.1$, the structure of the wavefunction transitions such that the fermion nodes are now located within the wave zone, and the curve then oscillates around a constant radius. Note that along this section, it can be seen that a secondary fold in the curve has begun to develop, indicating that five distinct solutions are present within the region $z \approx 5.3 - 7.1$. This additional fold is also in evidence in the fermion energy-redshift plot shown in Fig. 4.6. To analyse this properly, however, we must increase the fermion number even further.

4.2.4 The $N_f = 38$ system

We now consider the $N_f = 38$ system, which represents the upper limit of our numerics, at least with regard to generating complete families of solutions. For values of the fermion number that are this large, the second fold discussed above now extends into the high-redshift regime, allowing us to analyse the structure of the states located along it. This is shown in Fig. 4.8, where we again plot the fermion energy-redshift and mass-radius relations for the $n = 2$ family of states. These may appear significantly more complicated than their $N_f = 20$ equivalents, but the overall behaviour is in fact similar. The major difference concerns the second fold, which is observed to oscillate around the $n = 0$ infinite-redshift state, in conjunction with the first. Thus the mass-radius curve completes three spirals: two around the ground state, followed by one around the $n = 2$ infinite-redshift solution. In order to make this behaviour more apparent, we have again separated the curves into distinct branches, with the blue section containing the first fold, the orange the second fold, and the purple the final portion that ultimately reaches infinite redshift. The outgoing and incoming portions (with respect to redshift direction) are distinguished by solid and dashed lines, respectively. Note that the second fold does not return to the same value of redshift as the first ($z = 1.95$), but instead reverses direction at $z = 4.68$, which corresponds to the position of the second cusp-like feature in the ground-state curve.

Considering now the individual states shown in the lower panels of Fig. 4.8, these have a common redshift value of $z \approx 41$, and there are now five distinct $n = 2$ states (B–F),

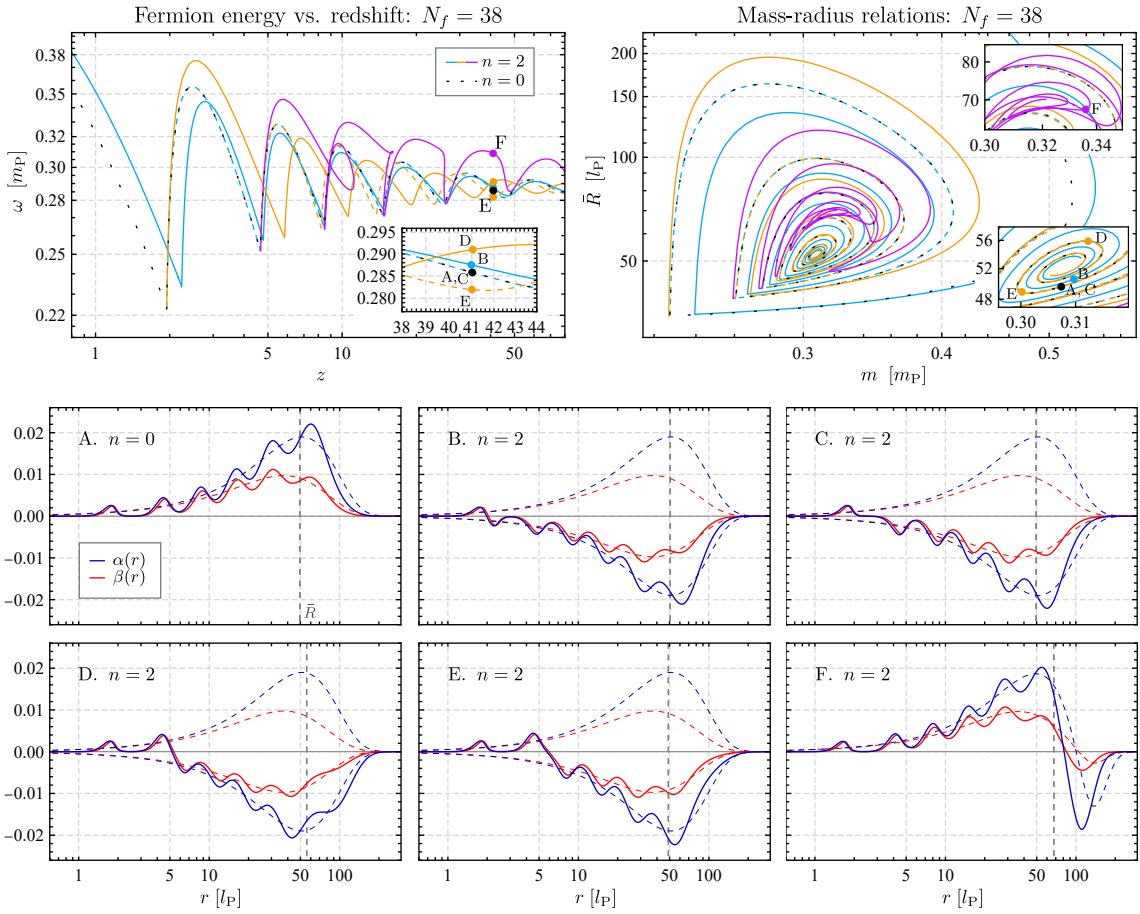


Figure 4.8: Summary of the $N_f = 38$ system, showing (**top**) the fermion energy-redshift and mass-radius relations for the families of $n = 0$ and $n = 2$ states. These are separated into colour-coded branches, which progress in the order light blue \rightarrow orange \rightarrow purple. Shown also (**bottom**) are the fermion field profiles for the six individual states with $z \approx 41$, five of these being $n = 2$ solutions (B–F), and the other the unique ground state (A).

accompanied by a single $n = 0$ ground state (A). As is evident, the fermion self-trapping is even stronger at this value of N_f , with the first peak in the ground-state wavefunction having begun to separate from the remainder, and in addition the second minimum now also approaches close to zero. With regard to the $n = 2$ states, solutions B, C and F are similar in structure to those discovered in the $N_f = 20$ system, with the fermion nodes in B and C occurring just outside the first power-law oscillation, while in F they are contained within the wave zone. Unsurprisingly, we find that states B and C are located along the first fold in the energy-redshift curve, while F is located along the final spiralling portion. The two new solutions, D and E, which lie along the second fold, exhibit a structure unseen at $N_f = 20$ — the fermion nodes are still located within the power-law zone, but they are now preceded by two trapped peaks. This explains why the second fold does not return to the same redshift value as the first. Along the first fold, the transition of the fermion nodes from the power-law to the wave zone occurs at the redshift at which the power-law zone first forms, but for the second fold, this transition can occur only once the second power-

law oscillation is present, i.e. at higher redshift. Using the rationale discussed earlier, we can also explain why the second fold oscillates around the $n = 0$ curve. This is simply due to the lack of fermion nodes within the wave zone of solutions, as is evident from the structure of the $n = 2$ states D and E. Note that since these occur along the same fold, they do not exhibit any clear differences in nodal structure, as mentioned previously.

We are also now in a position to explain, at least partially, why there exists such a multiplicity of states at high fermion number. The key feature once again appears to be the fermion self-trapping effect. As this grows in strength with increasing fermion number, the minima in the wavefunction, associated with the locations of unstable photon spheres, become progressively more pronounced. At some point, the first minimum approaches so close to zero that it becomes possible for the fermion fields to switch sign prematurely, resulting in a pair of fermion nodes occurring within the power-law zone. As N_f is subsequently increased, the value at the second minimum also approaches zero, and there is now a possibility of fermion nodes appearing outside the second trapped peak in the fermion wavefunction, but still within the power-law zone. If the fermion number were to be increased even further, we surmise that, as each subsequent minimum in the wavefunction drops towards zero, new solutions will appear in which the fermion nodes are located outside the third, fourth and fifth power-law peaks, and so on. One might wonder whether, at some point, the power-law zone will ‘run out’ of oscillations that can be converted into fermion nodes. For a fixed value of redshift, this will indeed be the case, but it is solely a reflection of the fact that each additional fold in the energy-redshift curve begins at a slightly higher redshift than the last. Thus given a particular value of redshift, there will be a maximum number of $n = 2$ states that can exist, regardless of the fermion number, but this will become infinite as both z and N_f tend to infinity.

This concludes our analysis of the first even-parity excited states. We turn now to the issue of higher excited states, for which the landscape becomes increasingly complex.

4.3 Higher excited states

As indicated earlier in section 4.1.2, when considering the unscaled $N_f = 40$ system (Fig. 4.2), the number of excited-state solutions present at a particular value of redshift tends to increase with increasing n (provided of course that the system is sufficiently relativistic). Here, we shall analyse the behaviour of higher excited states ($n \geq 4$) in detail, and show that this is indeed the case, with the arrival of new solutions being attributed to the increased number of ways in which multiple fermion nodes can be arranged within the power-law zone. As a result, the system becomes increasingly multivalued as both n and N_f are increased.

4.3.1 Small fermion number

We begin, however, by first reviewing the situation at relatively small fermion number, where the number of solutions remains manageable. A summary of this, for the $N_f = 8$, 8 and 12 systems, is presented in Fig. 4.9. This shows the fermion energy-redshift and mass-radius relations for the families of ground and first three even-parity excited states. These can be compared with the corresponding curves for the two-fermion system, shown previously in Fig. 4.1. At each value of N_f , the overall behaviour of the three excited-state families is qualitatively similar, with a multivalued region emerging in each energy-redshift curve, with these encompassing approximately the same redshift range. By considering the mass-radius relations, we observe that, as the fermion number is increased, each subsequent excited state gradually begins to wrap around the previous, before ultimately spiralling towards the appropriate infinite-redshift solution.

It therefore appears that, for systems containing relatively small numbers of fermions, the higher excited states are no more multivalued than the first. This is a consequence of the weak nature of the self-trapping effect at such fermion numbers, since only the first minimum in the wavefunction is close enough to zero to allow a fermion node to form. Thus only a single fermion node can occur within the power-law zone, irrespective of the value of n , with the remainder confined to the wave zone. Families of states with n total nodes therefore wrap around the families with $n - 1$ nodes, since solutions located in these regions contain exactly one fewer fermion node in their outer wave zones. Hence the total number of excited states at a particular value of redshift is identical for each value of n , a

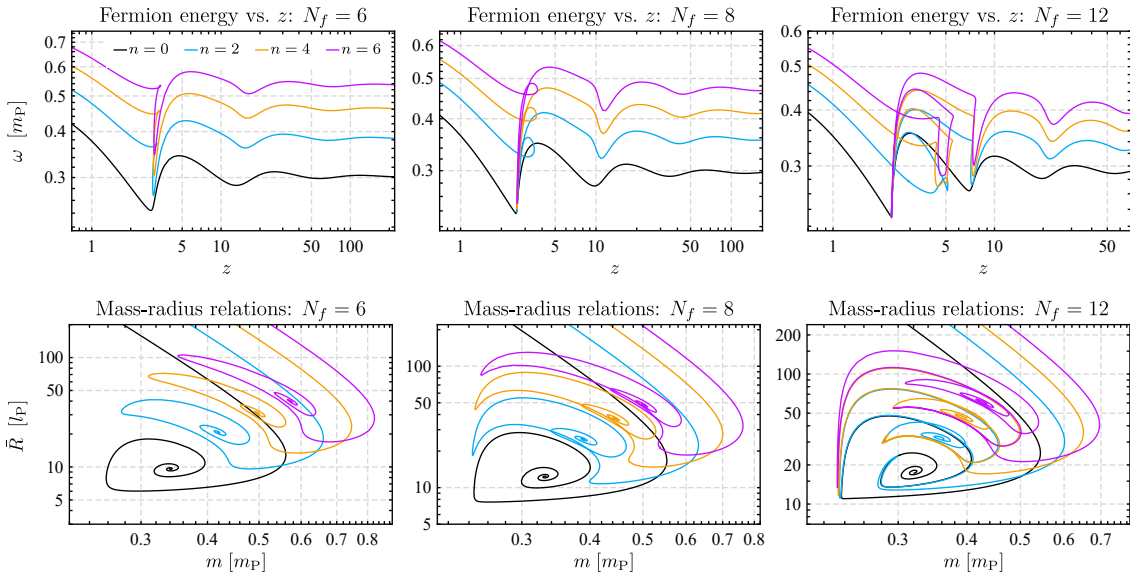


Figure 4.9: Plots illustrating the self-similar behaviour exhibited by the first three even-parity excited states, for systems with relatively small numbers of fermions. Shown (**top**) are the fermion energy-redshift and (**bottom**) mass-radius relations for the families of $n = 0, 2, 4$ and 6 states, at N_f values of $6, 8$ and 12 respectively.

feature that remains the case for systems containing up to and including 24 fermions.

4.3.2 The $N_f = 38$ system: $n = 4$ states

To observe a significant disparity between the relative number of excited states, it proves necessary to once again increase the fermion number. We thus return to the $N_f = 38$ system, now analysing the behaviour of the second even-parity ($n = 4$) states.

A summary of this is presented in Fig. 4.10, in which we again show the fermion energy-redshift and mass-radius relations, in this case solely for the family of $n = 4$ states. These plots are somewhat cluttered, to say the least, but we have once again separated the curves into colour-coded branches, in an attempt to make them intelligible. There are here a total of four folds in the $n = 4$ energy-redshift curve (compared to two in the corresponding $n = 2$ curve), the order of which can be described as follows. Entering from the non-relativistic regime, we first traverse the red portion of the curve, which oscillates towards the $n = 0$ infinite-redshift state, before reversing direction (following the dashed portion) and transitioning to the light blue branch at $z = 2.24$. This oscillates around the $n = 2$ infinite-redshift solution, reverses, and transitions to the orange branch at $z = 4.78$, which once again oscillates around the $n = 0$ infinite-redshift state. This then reverses direction at high redshift, returning to $z = 1.95$, and transitions to the purple branch, which in turn oscillates around the $n = 1$ infinite-redshift state, before the final transition to the black branch occurs at $z = 4.68$. This behaviour is mirrored in the mass-radius relations, in which it can clearly be seen that the curve spirals twice towards the $n = 0$ infinite-redshift state (along the red and orange sections), twice towards the $n = 2$ infinite-redshift state (along the light blue and purple sections), followed by a final spiral towards the $n = 4$ infinite-redshift state. Note that the appearance of these two additional folds (relative to the $n = 2$ case) is not a consequence of a distortion occurring in the region surrounding the third cusp-like feature in the fermion energy-redshift curve, as one might have expected. Indeed, one can see that this has only partially begun to develop, even at $N_f = 38$. These are instead ‘folds within folds’, which emerge along those already present.

How are we to interpret the presence of these new folds? As before, clues can be obtained by considering the structure of individual high-redshift solutions, which, for the case of $z \approx 40$, are shown in the lower panels of Fig. 4.10. In total, there are nine distinct $n = 4$ states, each containing exactly two fermion nodes within each fermion field, as required. Five of these states (C, D, G, H and I) are analogous to the five previously discovered $n = 2$ states (see Fig. 4.8), in that they are identical in nodal structure, except for the presence of an additional fermion node within the outer wave zone. These states are located along the branches that spiral towards the $n = 2$ infinite-redshift state (in the case of C, D, G, and H) and the $n = 4$ infinite-redshift state (in the case of solution I). The four new solutions (A, B, E and F) therefore reside along the sections that spiral towards the $n = 0$ infinite-redshift state, as could be predicted by their lack of wave-zone

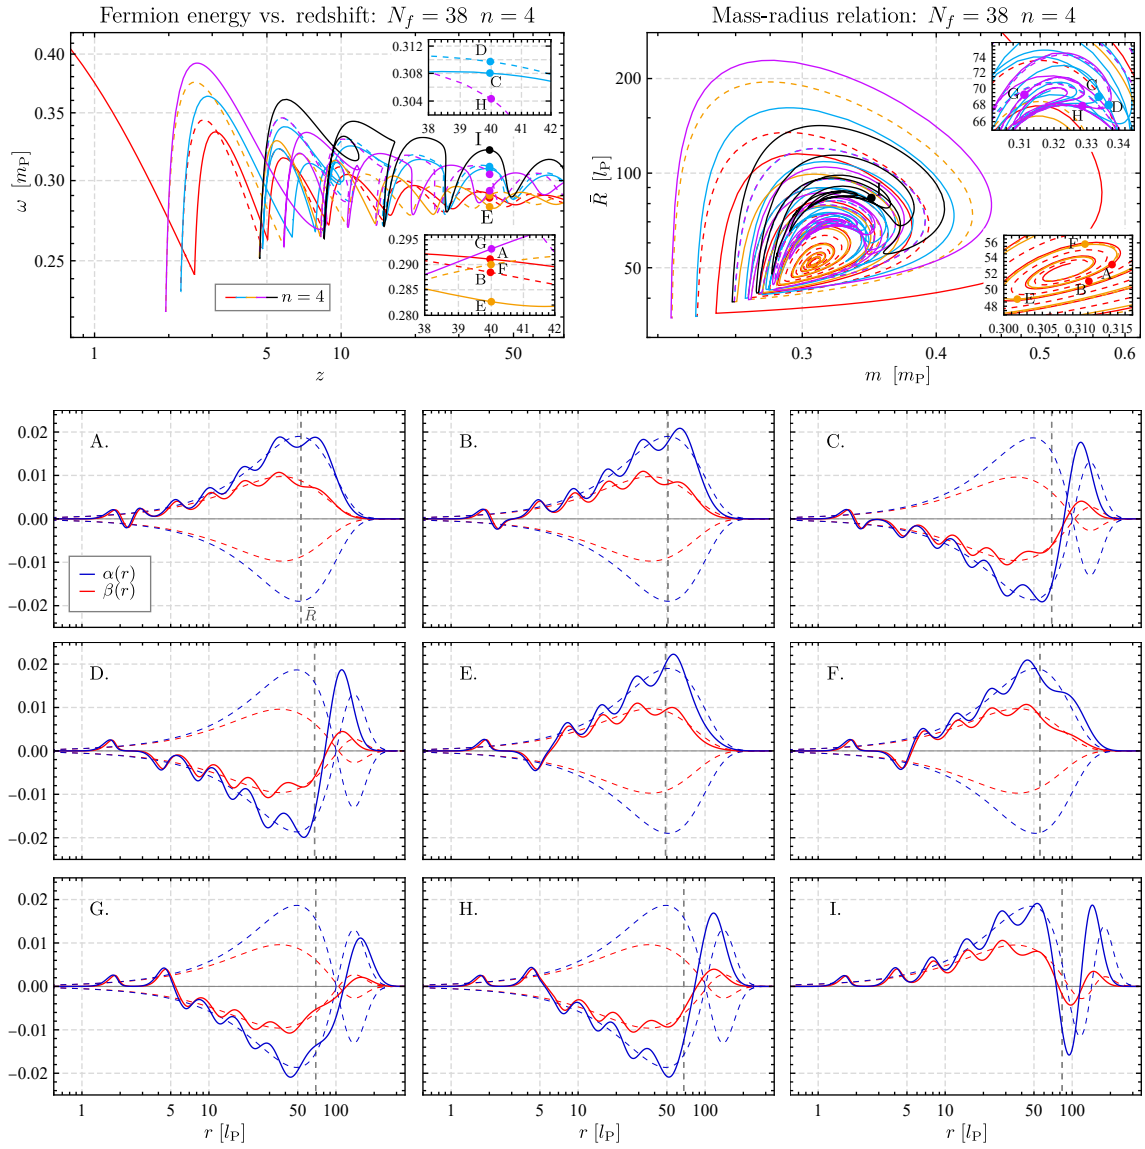


Figure 4.10: Plots summarising the behaviour of the second even-parity excited states, for the $N_f = 38$ system, showing (**top**) the fermion energy-redshift and mass-radius relations for the family of $n = 4$ states. The curves are separated into colour-coded branches, with the outgoing and incoming portions denoted by solid and dashed lines, respectively. The profiles of the fermion fields for nine $n = 4$ states, each with a common redshift of $z \approx 40$, are also shown (**bottom**), with their locations indicated on the plots above.

nodes. In all four cases, the two fermion nodes are located within the power-law zone, but it is not entirely clear how to interpret the almost identical structures exhibited by the fermion wavefunctions. Our best guess is as follows. Solutions A and B, which are located along the outgoing and incoming red branches of the curves, respectively, contain two fermion nodes sandwiched between the first and second trapped peaks, with this being more obvious in solution A. By contrast, solutions E and F contain only one node between the first and second peaks (the latter of which has become a minimum), while the other is located between the second and third peaks.

Assuming this interpretation to be accurate, we might have anticipated the appearance of solutions E and F at $N_f = 38$, since both the first and second minima in the ground state are sufficiently close to zero to allow nodes to form. It seems reasonable, therefore, that two nodes could be contained within the power-law zone, with each minimum converting to a node, as observed. What is somewhat unexpected, however, is the presence of solutions A and B, in which both nodes are seemingly located between the first and second trapped peaks. To explain this, we note that, as mentioned briefly in Chapter 3, the strength of the self-trapping effect at large N_f is such that the first peak in the fermion wavefunction becomes spatially separated from the remainder (which are evenly spaced in $\log(r)$). Due to the large radial extent of the resulting void, it would appear that, for sufficiently large N_f , there is room to fit two fermion nodes between the first two trapped peaks. Given that the spatial separation of peaks continues to increase alongside the fermion number, this suggests that, at even higher N_f , it may be possible to fit three, or potentially even more, nodes between the first two peaks, and indeed that multiple nodes may appear within the corresponding voids between the second, third and fourth peaks, and so on. This would act to magnify the multiplicity of excited states even further as the fermion number continues to increase.

4.3.3 Varying N_f : A summary

Although we cannot be certain as to whether the nodal interpretation discussed above is correct, our analysis of individual states for systems containing up to and including 70 fermions has established that scenarios such as those described can indeed occur. We shall now summarise these results, with particular emphasis on the number and relative structure of states present at each fermion number.

Figure 4.11 provides a detailed overview of the excited-state system as the number of fermions is increased. Here, we plot the soliton radius (strictly speaking $\bar{R}/\sqrt{N_f}$) as a function of fermion number, for all solutions with $z \approx 100$, up to and including fourth even-parity excited states ($n = 8$). This provides a convenient method for separating states based on the number of fermion nodes present within their wave zones. Note that we have here included both non-even and non-integer values of N_f , in order for the behaviour to become apparent, but it should be emphasised that only even values of N_f are physically acceptable. This of course does not prevent one from solving the FSY equations (2.28)–(2.31) with an arbitrary non-integer value of N_f , as has been performed when generating these figures. Consider first the relatively simple behaviour of the $n = 2$ states. For small numbers of fermions, only a single state exists at each value of N_f , with this exhibiting a single fermion node (in each field) located within the wave zone. As the fermion number is increased, recall that a fold appears in the fermion energy-redshift curves, as previously described, with this gradually extending outwards in redshift, reaching $z = 100$ at $N_f = 13.39$. At this point, a new pair of solutions therefore appears, for which the fermion nodes are located within the power-law zone (hence the smaller

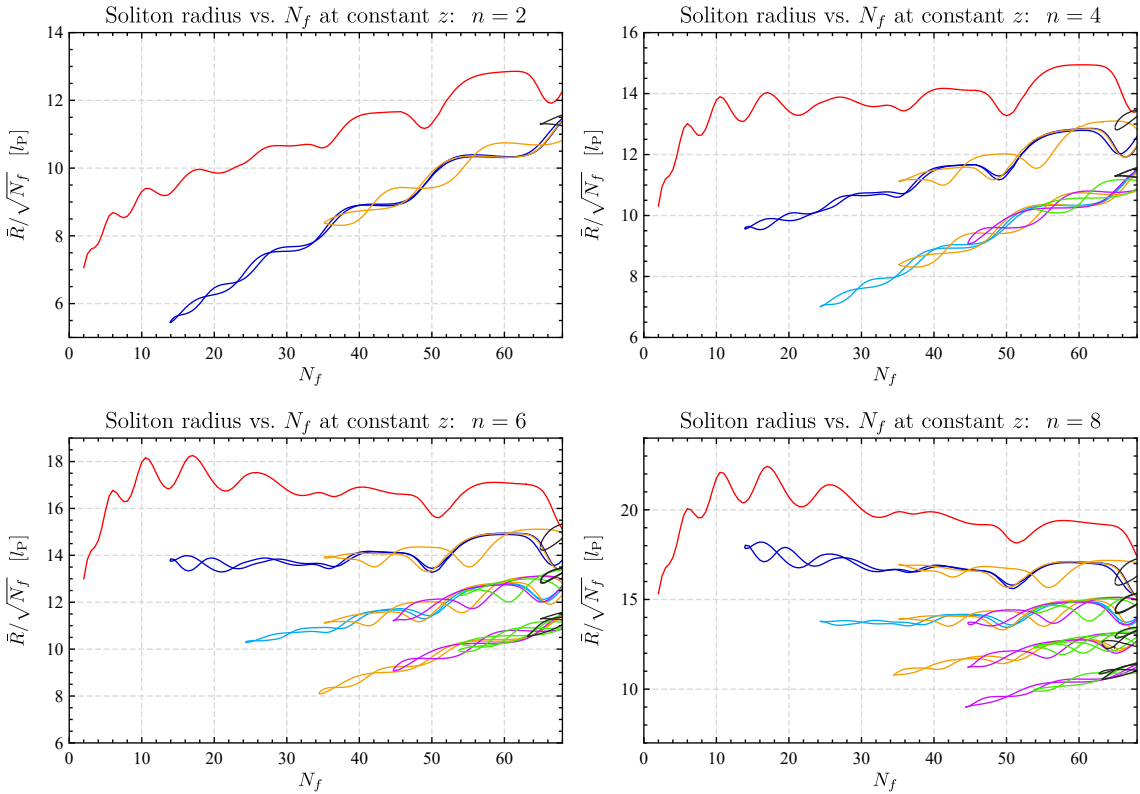


Figure 4.11: Plots summarising the behaviour of the first four even-parity excited states as a function of fermion number, showing the radius of all solutions present at a redshift value of $z \approx 100$. The emergence of new pairs of states as the fermion number is increased can clearly be seen, with the radii of these indicating the number of fermion nodes contained within the wave zone. Note that we have colour-coded the curves such that branches appearing at the same value of N_f have the same colour.

radial extent), just outside the first trapped peak. There remain only three solutions up until $N_f = 35.19$, where a second new pair appears, these corresponding to states along the second fold, which has now extended outwards to $z = 100$. The fermion nodes for these states are similarly located in the power-law zone, but now between the second and third peaks. Finally, at $N_f = 64.98$, a third pair of states emerges, for which the self-trapping effect is now strong enough to allow nodes to form between the third and fourth trapped peaks.

Turning now to the behaviour of the $n = 4$ states, we observe that this is similar to the $n = 2$ case up until $N_f = 24.36$, where an additional pair of solutions is seen to emerge. These are of significantly smaller radius than the three already present, and hence we conclude that the two fermion nodes in these states are located within the power-law zone, both between the first and second trapped peaks. At $N_f = 35.19$, the strength of the self-trapping is such that the second minimum in the fermion wavefunction can now convert to a node. We therefore observe the concurrent appearance of two pairs of new states, one of these having a single fermion node within the power-law zone (located after the second peak), and the other having two (one located after the first peak, and the

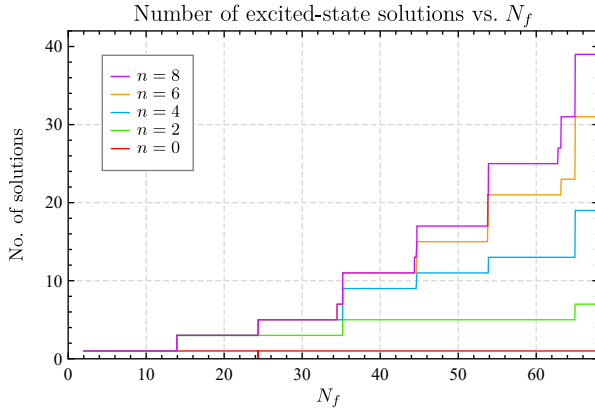


Figure 4.12: The total number of solutions, for the ground and first four even-parity excited states, plotted as a function of fermion number. This indicates the number of states of each type present at $z \approx 100$. There exists always a unique $n = 0$ ground state, but the number of excited states becomes progressively larger as both n and N_f increase.

other the second). Thereafter, subsequent pairs of states emerge at $N_f = 44.7$, $N_f = 53.88$, followed by two pairs at $N_f = 64.98$, making a total of 19 distinct $n = 4$ states.

The corresponding plots for the $n = 6$ and $n = 8$ states are also presented, although we shall not provide a detailed account of their behaviour. Note that, to aid readability, we have colour-coded the curves to indicate branches that emerge at approximately the same value of N_f . Suffice to say the number of states increases substantially with each subsequent excited state, due to the possibility of additional nodal configurations. The overarching nodal structure of states can be read off the plots in a straightforward manner, by noting that the number of wave-zone nodes is related to the soliton radius. For example, at $N_f = 50$, there are a total of 17 $n = 8$ states: 2 with no wave-zone nodes, 4 with one pair, 6 with two pairs, 4 with three pairs, and 1 with four pairs. Beyond this classification, states are distinguished by the precise distribution of nodes within their power-law zones.

The total number of excited states present at each value of N_f is shown in Fig. 4.12 (again at constant $z \approx 100$). From this we see that there is always a unique $n = 0$ ground state, regardless of fermion number, and that the excited states are also unique up until $N_f = 13.939$. Between this and $N_f = 24.36$, three solutions exist for each value of $n > 0$, but beyond this the number of solutions differs with n , with the curves becoming ever more separated as the fermion number increases. Thus at $N_f = 68$, the limit of our numerics, there are a total of 7 $n = 2$ states, 19 $n = 4$ states, 31 $n = 6$ states and 39 $n = 8$ states. Note that jumps in the number of states occur for approximately every 10 fermions added, and that the precise values of these can differ slightly from state to state. We have established that this latter feature is not a numerical artefact, but its interpretation is not entirely clear.

4.4 Discussion

In this chapter, we have analysed thoroughly the excited states of relativistic many-fermion Einstein–Dirac solitons, showing that their behaviour differs significantly from that ob-

served in the two-fermion system. In particular, the central redshift can no longer be used to uniquely identify states, with physical quantities such as the fermion energy becoming multivalued with respect to this parametrisation. This can be attributed to a breakdown in the zonal structure of solutions, in which fermion nodes can form within the power-law zone, a feature permitted at large fermion number due to the strength of the fermion self-trapping effect. The system becomes increasingly multivalued when considering ever higher excited states, a consequence of multiple nodes existing within the power-law zone, and thus the spectrum of states becomes increasingly complex as both n and N_f are increased.

Although we have presented a detailed review of the overall landscape, there remain a number of unanswered questions. First, we can provide no physical interpretation for the nodal structures observed in the excited-state solutions. Certainly the fermion self-trapping effect appears to be relevant, in allowing the fermion wavefunction to drop sufficiently close to zero, but exactly what causes the fermion fields to switch sign prematurely is unknown. Also unclear is why the folds in the fermion energy-redshift curves have only a finite redshift extent (at least in many cases). Perhaps there exists some mechanism that prevents nodes from forming at sufficiently small r . From a purely physical perspective, the most we can say is to note that the Einstein–Dirac system becomes increasingly non-linear as the fermion number is increased, and thus we should not be entirely surprised by the presence of multiple solutions. It would be interesting to determine whether similar systems, such as excited boson or Proca stars, exhibit an equivalent behaviour. There is indeed tentative evidence that this may be the case, in the context of rotating boson stars [91], in which multivaluedness is similarly observed, although the validity of this comparison is debatable.

We also note that, when presenting individual solutions throughout this chapter, we have omitted to display the structure of the corresponding metric fields, partly due to space constraints. One might wonder whether these, or indeed the forms of the optical geometry, are able to shed light on the situation. We therefore present an example of both the metric fields and optical geometry for a many-fermion excited state in Fig. 4.13, corresponding to solution A shown earlier in Fig. 4.10. We also include the corresponding ground-state plots, for comparison. Note that in this case, the two pairs of additional fermion nodes are both located between the first and second trapped peaks in the fermion wavefunction. As one can see, the effect of the fermion nodes is to somewhat dilute the first minimum in the metric field A (relative to its ground-state form), causing it to split into multiple smaller oscillations of the type previously only associated with the wave zone. There is also a corresponding change in the optical geometry, with the first trapping region becoming less pronounced and covering a larger spatial extent. The behaviours of both plots can be understood by considering the excited-state form of the radial fermion energy density $r^2\rho_f(r)$. Due to the presence of the fermion nodes, the first trapped shell

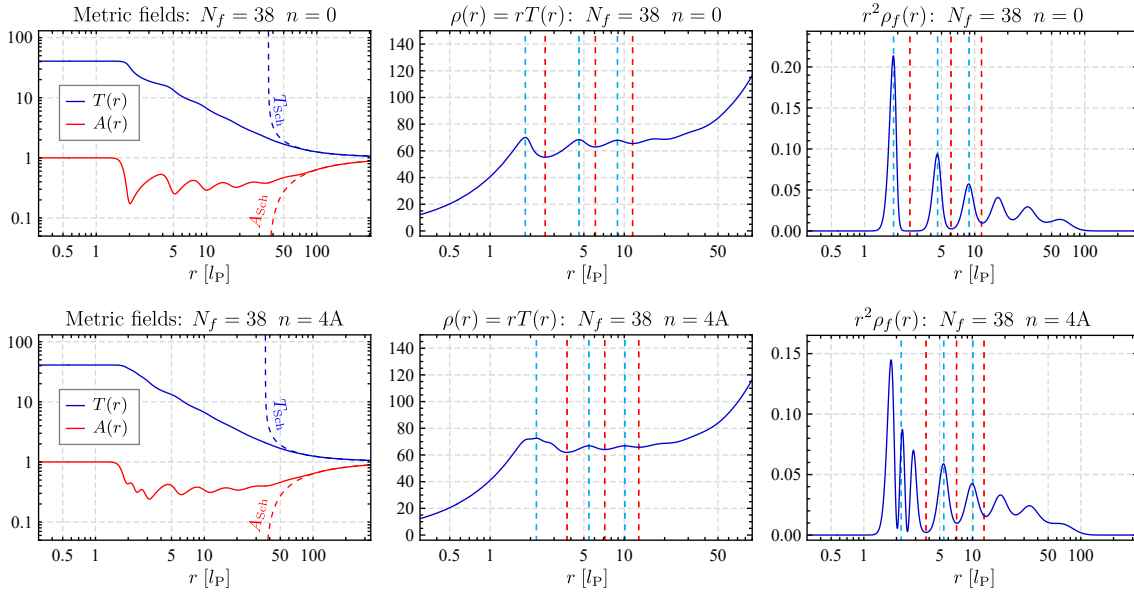


Figure 4.13: Plots showing the radial profiles of the metric fields (**left**), radial optical geometry co-ordinate $\rho = rT$ (**middle**) and radial fermion energy density $r^2\rho_f$ (**right**) for the ground state (**top**) and a fourth even-parity excited state (**bottom**) of the $N_f = 38$ system. Both of these have a central redshift of $z \approx 40$, and the excited state corresponds to solution A shown in Fig. 4.10. The locations of stable and unstable photon spheres are indicated by dashed light blue and red lines, respectively.

splits in a series of three distinct peaks, resulting in a more evenly-spread distribution of energy. Thus the accompanying metric distortion is slightly less prominent, with the triple peaks causing a corresponding oscillatory behaviour in A . Note of course that the separation of these three peaks is not a consequence of the self-trapping effect (as can be seen from the optical geometry), but is purely due to the presence of the fermion field nodes. Overall, although this behaviour is interesting in its own right, it does not appear to clarify any remaining questions regarding interpretation.

Finally, we emphasise that, although the fermion energy-redshift curves at large fermion number are multivalued as a function of redshift, they nonetheless still represent continuous, one-parameter families of states. The complication is that the central redshift z no longer provides the appropriate parametrisation. It should therefore be possible, at least in principle, to employ an alternative parametrisation in which the curves become entirely single-valued. An attempt at this is briefly summarised as follows. Recall that the fermion field nodes transition continuously from the power-law zone to the outer wave zone as the curves are traversed, suggesting that progress could be made by constructing ‘redshift-like’ parameters from the values of T at each node, relative to its central value:

$$z_{p\alpha} = \frac{T(0)}{T(r_{p\alpha})} - 1. \quad (4.1)$$

Here, $r_{p\alpha}$ is the radius of the p^{th} node in the field α , and we have subtracted 1 in order to




Figure 4.14: The results of parametrising families of excited states by the central redshift z_0 (**top**) compared to the quantity $z_{1\alpha}$ (**bottom**). We show three different cases, corresponding to the N_f and n values indicated. Note that the dashed lines contained within the two right-hand plots correspond to redshift regions that are beyond the upper limit of our numerics.

ensure $z_{p\alpha}$ runs from zero to infinity. For nodes which are located within the outer wave zone, this parameter will differ only slightly from the central redshift (which we here refer to as z_0), but for those in the power-law zone, the deviation will be significant. Various results of this new parametrisation are shown in Fig. 4.14, where we plot the fermion energy as a function of both central redshift z_0 and the parameter $z_{1\alpha}$, for three different families of solutions, at three different fermion numbers. As can be seen, this new parametrisation is entirely successful for the $N_f = 12$, $n = 2$ case, with the curve becoming completely single-valued. For the $N_f = 20$, $n = 2$ curve, however, a small multivalued still remains (see inset), although we suspect that this can be removed by constructing a parameter that incorporates nodes in both α and β . For these $n = 2$ states, the application of the parametrisation above is unambiguous, since only one node exists in each fermion field. When considering higher excited states, however, it is not clear how to proceed, since there are now multiple nodes present. The result of parametrising the $N_f = 38$, $n = 4$ state by solely $z_{1\alpha}$ is shown in the figure, but we find that significant multivalued sections remain, corresponding to regions in which the second fermion node transitions from the power-law to the wave zone. It is clear therefore that an alternative is required. Our attempts at constructing composite parameters, involving multiple nodes, however, have proved unsuccessful, although we note that switching parameters between branches does in fact produce curves that are no longer multivalued. This is somewhat outwith the spirit of a single-valued parametrisation, however, as it requires prior knowledge of the entire family of states. The construction of a true single-valued parameter, valid for arbitrary values of n and N_f , therefore remains an open problem.

Chapter 5

Einstein–Dirac–Higgs solitons

The Einstein–Dirac system considered thus far provides a relatively simple toy model in which one can study interactions between neutral, fermionic quantum particles in general relativity. Although this proves useful in analysing purely gravitational effects, its applicability to physically relevant fermions, which interact via the additional mechanisms described within the Standard Model, is limited. Various extensions to the Einstein–Dirac formalism have therefore been considered, namely the addition of electromagnetic interactions [41], the electroweak force [42], and more recently a Higgs mechanism by Leggat [53]. This latter study, however, considers the inclusion of a conformally coupled Higgs field (which is subsequently applied in the context of conformal gravity), and numerical issues prevent a detailed evaluation of the resulting particle-like states.

In this chapter, we therefore present a more thorough analysis of the Einstein–Dirac–Higgs system, in which the addition of a minimally-coupled Higgs field is considered, as is the case in the Standard Model, and demonstrate that indeed particle-like states exist within this context, which we shall refer to as ‘Einstein–Dirac–Higgs solitons’. Although these largely exhibit similar properties to their Einstein–Dirac counterparts, we intriguingly discover the presence of a mass-scale separation at strong fermion-Higgs coupling, in which the ADM mass of a state is no longer proportional to the mass of its constituent fermions. This chapter is structured as follows. We first discuss the application of the Higgs mechanism to this problem, and derive the appropriate equations of motion, before presenting examples of localised particle-like solutions for the two-fermion system. We subsequently analyse the properties of the families of these states, showing that they exhibit a similar spiralling behaviour to that observed in the Einstein–Dirac system. Thereafter follows a detailed discussion concerning the mass-scale separation observed at strong coupling, along with possible reasons for its appearance, including the results of a large- r analysis. We then briefly present results concerning the zonal structure of solutions, along with an example of an infinite-redshift state, before concluding with a short discussion.

5.1 The Einstein–Dirac–Higgs system

We begin by deriving the equations of motion describing static, spherically symmetric configurations of gravitationally-interacting fermions, modelled within the Einstein–Dirac formalism, with a minimally-coupled scalar Higgs field now included. As we shall detail, various subtleties arise concerning the application of the Higgs mechanism within this semi-classical context, which have implications when considering localised particle-like solutions. As noted, this system has previously been studied by Leggat [53] (although with regard to a conformal coupling), and thus much of the proceeding derivation and discussion is contained therein.

5.1.1 Setup

In the Standard Model, the Higgs mechanism is responsible for generating the observed masses of elementary particles (including the Higgs boson itself), via a spontaneous symmetry breaking that occurs at the electroweak energy scale. This is achieved by the Higgs field h relaxing to the minimum of its quartic ‘Mexican-hat’ potential (see Fig. 5.1), and consequently obtaining a non-zero vacuum expectation value (vev) v . For fermionic particles, the Higgs field is incorporated via the addition of a standard Yukawa coupling term, $\mu\bar{\Psi}h\Psi$, and thus the mass of each fermion is given by its respective coupling strength μ multiplied by the value of the Higgs vev.

Although the Standard model Higgs is strictly a complex SU(2) doublet, we shall here consider the simplified case of a real scalar field, denoted h , leading us to write the total action for the Einstein–Dirac–Higgs system as:

$$S_{\text{EDH}} = \int \left(\frac{R}{16\pi G} + \mathcal{L}_m \right) \sqrt{-g} d^4x, \quad (5.1)$$

where the matter Lagrangian density is now:

$$\mathcal{L}_m = \bar{\Psi}(\not{D} - \mu h)\Psi - \frac{1}{2}(\nabla^\nu h)(\nabla_\nu h) - V(h). \quad (5.2)$$

Here, μ is the strength of the fermion-Higgs coupling, which we shall leave unspecified, while the Higgs potential $V(h)$ is taken to be of the usual Mexican-hat form:

$$V(h) = \lambda(h^2 - v^2)^2, \quad (5.3)$$

where λ is an overall scaling factor and v is the Higgs vacuum expectation value, the values of both of which we shall allow to vary. As in the Einstein–Dirac system, the variation of the action (5.1) with respect to the spinor Ψ produces the Dirac equation, while variation with respect to the metric $g_{\mu\nu}$ results in the Einstein equations. There is also an equation of motion for the Higgs field, found by varying the action with respect to h , or equivalently

by evaluating the Euler-Lagrange equations for h using \mathcal{L}_m , as follows:

$$\frac{\delta \mathcal{L}_m}{\delta h} = \nabla_\mu \left(\frac{\delta \mathcal{L}_m}{\delta (\nabla_\mu h)} \right) \implies \nabla_\mu \nabla^\mu h = \mu \bar{\Psi} \Psi + \frac{dV}{dh}. \quad (5.4)$$

Notice that this contains a term that depends on the spinor wavefunction, and therefore does not simply represent a statement asserting the stability of the minima of $V(h)$. Instead, it leads us to define an effective Higgs potential:

$$V_{\text{eff}}(h) = \mu h \bar{\Psi} \Psi + V(h), \quad (5.5)$$

the form of which now governs the dynamics of the Higgs field, as follows. In the absence of the fermion wavefunction (i.e. in vacuum), the effective potential reduces to the usual Mexican-hat form, and hence the Higgs field is pinned at its vacuum expectation value v . At positions where the spinor field is non-zero, however, the effective potential consequently acquires an additional tilt, dependent on the value of the fermion wavefunction at that point, resulting in the minima of V_{eff} no longer coinciding with $h = \pm v$. The Higgs field will therefore tend to deviate from its vacuum expectation value in the presence of a fermion source, such as a localised particle-like state, becoming instead a function of position. This in turn implies that the fermion mass, μh , of such an object must also be position-dependent, with its value varying within the fermion source, but approaching an asymptotic value, μv , as the wavefunction decays and the Higgs field returns to its vacuum expectation value. It is this asymptotic value that we shall define as the fermion mass m_f of a localised object, i.e. that calculated by an observer at spatial infinity:

$$m_f = \mu v. \quad (5.6)$$

In contrast, the mass associated with the Higgs field itself is indeed a fixed quantity, and is calculated by considering the displacement of the field around its vacuum expectation value. In order to obtain an explicit expression, we therefore write $h = v + \delta h$, and substitute this into the matter Lagrangian density:

$$\mathcal{L}_m = \bar{\Psi} (\not{D} - \mu h) \Psi - \frac{1}{2} (\nabla_\mu \delta h) (\nabla^\mu \delta h) - \lambda ((v + \delta h)^2 - v^2)^2 \quad (5.7)$$

$$= \bar{\Psi} (\not{D} - \mu h) \Psi - \frac{1}{2} (\nabla_\mu \delta h) (\nabla^\mu \delta h) - 4\lambda v^2 \delta h^2 - 4\lambda v \delta h^3 - \lambda \delta h^4. \quad (5.8)$$

Comparing this with the standard form for a real scalar field, we conclude that the Higgs mass m_H takes the value:

$$m_H = 2v\sqrt{2\lambda}. \quad (5.9)$$

As a final remark, we note that this type of behaviour does not arise in the Standard model, partly because fermions are treated as point particles in quantum field theory, but also due to the fact that corrections to the Higgs potential due to fermionic interactions

are negligible at leading order. The phenomenon of a locally-varying fermion mass is therefore purely a consequence of considering extended localised objects within a semi-classical framework.

5.1.2 Equations of motion

Having outlined the general problem, we now present a brief summary detailing the derivation of the equations of motion valid for a static, spherically symmetric collection of fermions. This system was first considered in ref. [53], within the context of a conformally coupled Higgs field, and the expressions given therein reduce to those presented here when the conformal coupling ζ is set to zero.

We begin by restating the metric and spinor ansatzes used in the Einstein–Dirac case, which can be applied without modification to the system under consideration here:

$$g_{\mu\nu} = \text{diag} \left(-\frac{1}{T(r)^2}, \frac{1}{A(r)}, r^2, r^2 \sin^2 \theta \right); \quad (5.10)$$

$$\Psi_{jk}^{\pm}(t, r, \theta, \phi) = \frac{\sqrt{T(r)}}{r} \begin{pmatrix} \chi_{j\mp\frac{1}{2}}^k \alpha(r) \\ i\chi_{j\pm\frac{1}{2}}^k \beta(r) \end{pmatrix} e^{-i\omega t}. \quad (5.11)$$

Using these, an explicit expression for the Dirac equation can be derived in precisely the same manner as in the Einstein–Dirac system, with the sole alteration being that the fermion mass is replaced by the now locally-varying quantity μh :

$$\sqrt{A} \alpha' = +\frac{\kappa\alpha}{2r} - (\omega T + \mu h)\beta; \quad (5.12)$$

$$\sqrt{A} \beta' = -\frac{\kappa\beta}{2r} + (\omega T - \mu h)\alpha. \quad (5.13)$$

Turning to the Einstein equations, the contribution to the energy-momentum tensor from the fermionic sector is again identical to the Einstein–Dirac case, taking the form given in (2.15). To determine the contribution from the Higgs terms, we evaluate:

$$\begin{aligned} T_{\mu\nu}[h] &= \frac{-2}{\sqrt{-g}} \frac{\delta}{\delta g^{\mu\nu}} \left[\sqrt{-g} \left(-\frac{1}{2}(\nabla^\sigma h)(\nabla_\sigma h) - V(h) \right) \right] \\ &= (\nabla_\mu h)(\nabla_\nu h) - \left(\frac{1}{2}(\nabla^\sigma h)(\nabla_\sigma h) + V(h) \right) g_{\mu\nu}, \end{aligned} \quad (5.14)$$

where we have used the following identity for the variation of the metric determinant:

$$\delta\sqrt{-g} = -\frac{1}{2}\sqrt{-g} g_{\mu\nu} \delta g^{\mu\nu}. \quad (5.15)$$

Thus the total (mixed) energy-momentum tensor is:

$$T^\mu{}_\nu = - \sum_{k=-j}^j \Re \left\{ \bar{\Psi}_{jk}^{\pm} (i\gamma^\mu \partial_\nu) \Psi_{jk}^{\pm} \right\} + (\nabla^\mu h)(\nabla_\nu h) - \left(\frac{1}{2}(\nabla^\sigma h)(\nabla_\sigma h) + V(h) \right) \delta^\mu{}_\nu, \quad (5.16)$$

the non-zero components of which evaluate to:

$$T^t_t = -\frac{|\kappa|\omega}{r^2}T^2(\alpha^2 + \beta^2) - \frac{1}{2}A(h')^2 - V(h); \quad (5.17)$$

$$T^r_r = \frac{|\kappa|}{r^2}T\sqrt{A}(\alpha\beta' - \beta\alpha') + \frac{1}{2}A(h')^2 - V(h); \quad (5.18)$$

$$T^\theta_\theta = T^\phi_\phi = \frac{\kappa|\kappa|}{2r^3}T\alpha\beta - \frac{1}{2}A(h')^2 - V(h). \quad (5.19)$$

Since the form of the metric remains unchanged from the Einstein–Dirac case, the components of the Einstein tensor are precisely those given in (2.19)–(2.21), and therefore we can write the tt and rr Einstein equations as:

$$\frac{1}{r^2}(-1 + A + rA') = 8\pi G \left[-\frac{|\kappa|\omega}{r^2}T^2(\alpha^2 + \beta^2) - \frac{1}{2}A(h')^2 - V(h) \right]; \quad (5.20)$$

$$\frac{1}{r^2}(-1 + A - 2rA\frac{T'}{T}) = 8\pi G \left[\frac{|\kappa|}{r^2}T\sqrt{A}(\alpha\beta' - \alpha'\beta) + \frac{1}{2}A(h')^2 - V(h) \right]. \quad (5.21)$$

Note that once again only two of the four Einstein equations are independent. It only remains to derive an explicit expression for the Higgs equation (5.4), for which we require the evaluation of $\bar{\Psi}\Psi$:

$$\begin{aligned} \bar{\Psi}\Psi &= \sum_{k=-j}^j \bar{\Psi}_{jk}^\pm \Psi_{jk}^\pm = \frac{T}{r^2} \sum_{k=-j}^j \left(\alpha^2 \bar{\chi}_{j\pm\frac{1}{2}}^k \chi_{j\pm\frac{1}{2}}^k - \beta^2 \bar{\chi}_{j\pm\frac{1}{2}}^k \chi_{j\pm\frac{1}{2}}^k \right) \\ &= \frac{|\kappa|}{r^2}T(\alpha^2 - \beta^2). \end{aligned} \quad (5.22)$$

Then, using the metric ansatz (5.10) to evaluate the covariant derivatives, we arrive at the final expression for the Higgs equation:

$$Ah'' - A \left(\frac{T'}{T} - \frac{A'}{2A} - \frac{2}{r} \right) h' = \frac{|\kappa|\mu}{r^2}T(\alpha^2 - \beta^2) + \frac{dV}{dh}. \quad (5.23)$$

This, along with (5.12), (5.13), (5.20) and (5.21), constitutes a set of five coupled differential equations, for the five unknown fields $\{\alpha, \beta, A, T, h\}$, that fully define the behaviour of a static filled shell of neutral fermions interacting both gravitationally and via a minimally-coupled Higgs field. We shall hereafter refer to these as the ‘EDH equations’.

5.1.3 Particle-like solutions

We now discuss the boundary conditions required to generate particle-like solutions to the EDH equations. As in the Einstein–Dirac system, we require the spacetime of solutions to be asymptotically flat, i.e. $A(r), T(r) \rightarrow 1$ as $r \rightarrow \infty$, as well imposing the normalisation of each spinor wavefunction:

$$4\pi \int_0^\infty \frac{T}{\sqrt{A}}(\alpha^2 + \beta^2) dr = 1. \quad (5.24)$$




Figure 5.1: The intrinsic Higgs potential $V(h)$ (left) alongside the effective Higgs potential $V_{\text{eff}}(h)$ (right). The latter is inverted due to the static nature of the system, and obtains an additional tilt at radii inside the fermion source.

In order to ensure asymptotic flatness, the contribution to the Einstein equations (5.20) and (5.21) from the terms involving the Higgs field must vanish as $r \rightarrow \infty$. This is guaranteed provided the Higgs field asymptotes towards its vacuum expectation value outside the fermion source, matching the expected behaviour outlined earlier. It is not the case, however, that this occurs naturally, in that the Higgs field will not automatically relax to its vacuum expectation value in the absence of the fermion wavefunction. The reason for this can be understood as follows, where we employ the same rationale as Schlögel *et al.* [92]. First, let us consider the situation outside the fermion source, at a point where the spinor fields have decayed sufficiently for their contribution to the Higgs equation (5.23) to be negligible. Let us also temporarily introduce a time-dependence to the system, such that the dynamics of the Higgs field is now governed by the following equation:

$$T^2 \ddot{h} - Ah'' + A \left(\frac{T'}{T} - \frac{A'}{2A} - \frac{2}{r} \right) h' = -\frac{dV}{dh}, \quad (5.25)$$

where a dot represents the derivative with respect to time. The important point to note from this expression is that the time derivative of the Higgs field is of the opposite sign to the spatial derivative, and therefore the two dynamically stable minima in the Higgs potential are in fact unstable maxima from the point of view of spatial variations. Thus, in the purely static case considered here, the Higgs potential is effectively inverted compared to its usual intrinsic form, as is illustrated in Fig. 5.1. In addition, within the fermion source, the Higgs potential also acquires a tilt due to the presence of the fermion wavefunction, as discussed previously, resulting in an effective Higgs potential of the explicit form:

$$V_{\text{eff}}(h) = \lambda(h^2 - v^2)^2 + \frac{\mu|\kappa|}{r^2} T(\alpha^2 - \beta^2)h. \quad (5.26)$$

An example of this tilt is also shown in Fig. 5.1. Note that, for states that are fermion dominated (as opposed to anti-fermion dominated), i.e. those in which the fermion field α is dominant, the cumulative effect of this tilt is always to guide the Higgs field from a lower value to a higher one. Thus, in order to obtain states in which the (asymptotic) fermion mass is positive, the Higgs field is forced to take a value $h_0 < v$ at the centre of the fermion source. Asymptotically flat states are therefore those in which the value of h_0 is such that the fermion tilt guides the Higgs field to precisely its vacuum expectation value $h = +v$ as $r \rightarrow \infty$. Since we consider here only positive parity states, this is precisely the situation that arises, although we note that (negative parity) states with negative fermion mass do exist in which the the value of h_0 is similarly tuned, but now in order to guide

the Higgs field to the opposite minimum $h = -v$.

The final boundary conditions required to define particle-like solutions are those enforced at $r = 0$. As before, we desire states to be singularity free, leading us to consider a similar small- r asymptotic expansion to that utilised in the Einstein–Dirac case, but now including a Higgs field that tends to a constant value as $r \rightarrow 0$. We note that this is not guaranteed to be the unique non-singular expansion, an issue which we have not yet fully addressed, although it is certainly the expansion that will produce states most analogous to Einstein–Dirac solitons. Interestingly, the explicit form for this expansion differs depending on the value of the fermion number $N_f = |\kappa|$. The two-fermion system turns out to be something of a special case, in which the fermion and Higgs energy densities both contribute to the energy-momentum tensor at small r , resulting in the following asymptotic expansion (valid only for positive parity):

$$\alpha(r) = \alpha_1 r + \mathcal{O}(r^3); \quad (5.27)$$

$$\beta(r) = \frac{1}{3}(\omega T_0 - \mu v)\alpha_1 r^2 + \mathcal{O}(r^4); \quad (5.28)$$

$$T(r) = T_0 + \frac{4\pi G}{3}T_0(\lambda[h_0^2 - v^2]^2 + \alpha_1^2 T_0[\mu h_0 - \omega T_0])r^2 + \mathcal{O}(r^4); \quad (5.29)$$

$$A(r) = 1 - \frac{8\pi G}{3}(\lambda[h_0^2 - v^2]^2 + 2\omega\alpha_1^2 T_0^2)r^2 + \mathcal{O}(r^4); \quad (5.30)$$

$$h(r) = h_0 + \frac{1}{3}(2\lambda h_0[h_0^2 - v^2] + \mu\alpha_1^2 T_0)r^2 + \mathcal{O}(r^3), \quad (5.31)$$

where there are three unconstrained parameters: α_1 , T_0 and h_0 . For $N_f > 2$, however, the leading order terms in the fermion fields α and β become negligible compared to that of the Higgs field, and therefore it is only the latter that contributes to the energy-momentum tensor at next-to-leading order. The small- r expansion in this case (valid for $N_f > 2$) is hence:

$$\alpha(r) = \alpha_1 r^{\kappa/2} + \mathcal{O}(r^{\kappa/2+2}); \quad (5.32)$$

$$\beta(r) = \frac{1}{3}(\omega T_0 - \mu v)\alpha_1 r^{\kappa/2+1} + \mathcal{O}(r^{\kappa/2+3}); \quad (5.33)$$

$$T(r) = T_0 + \frac{4\pi G}{3}T_0\lambda(h_0^2 - v^2)^2 r^\kappa + \mathcal{O}(r^{\kappa+2}); \quad (5.34)$$

$$A(r) = 1 - \frac{8\pi G}{3}\lambda(h_0^2 - v^2)^2 r^\kappa + \mathcal{O}(r^{\kappa+2}); \quad (5.35)$$

$$h(r) = h_0 + \frac{1}{3}(2\lambda h_0[h_0^2 - v^2] + \mu\alpha_1^2 T_0)r^2 + \mathcal{O}(r^3), \quad (5.36)$$

which is again valid only for positive parity states. In what follows we shall consider primarily the two-fermion system, although an example of an $N_f = 20$ state is included in section 5.4.2.

5.2 Einstein–Dirac–Higgs solitons

We now proceed to the issue of generating particle-like states of the Einstein–Dirac–Higgs system, which we shall refer to as ‘Einstein–Dirac–Higgs solitons’. We shall first briefly discuss the numerical method via which these can be obtained, before presenting examples of individual solutions (for the two-fermion system), followed by an analysis of the resulting families of states.

5.2.1 Numerical method

The small- r expansions detailed in the previous section exhibit three unconstrained parameters: α_1 , T_0 and h_0 . In addition, the values of ω , μ , v and λ are all taken to be unspecified, although the last three would presumably be fixed by the physical properties of the Higgs potential and the strength of the fermion-Higgs coupling for the particular fermion species under consideration. In total, it appears that there are therefore seven free parameters in the system (for a fixed value of N_f), which is three more than in the Einstein–Dirac case. As in that system, however, three of these supposed degrees of freedom are removed by imposing the conditions of asymptotic flatness and normalisation, and in addition the value of h_0 must here be fixed such that the Higgs field asymptotes towards its vacuum expectation value. Thus only three unspecified parameters remain.

As in the Einstein–Dirac system, the numerical difficulties that arise when imposing normalisation and asymptotic flatness can be avoided by employing a rescaling technique, in which we first construct unscaled solutions (denoted by an additional tilde) where $\tilde{T}_0 = 1$ and $\tilde{\mu} = 1/(2v)$, from which the physical states can be subsequently obtained by rescaling the fields and parameters as follows [53]:

$$\begin{aligned} \alpha(r) &= \sqrt{\frac{\tau}{\chi}} \tilde{\alpha}(\chi r); & \beta(r) &= \sqrt{\frac{\tau}{\chi}} \tilde{\beta}(\chi r); & T(r) &= \frac{1}{\tau} \tilde{T}(\chi r); & A(r) &= \tilde{A}(\chi r); \\ h(r) &= \tilde{h}(\chi r); & \omega &= \chi \tau \tilde{\omega}; & \mu &= \chi \tilde{\mu}; & \lambda &= \chi^2 \tilde{\lambda}, \end{aligned} \quad (5.37)$$

where the quantities τ and χ are defined as:

$$\tau = \lim_{r \rightarrow \infty} \tilde{T}(r); \quad \chi^2 = 4\pi \int_0^\infty (\tilde{\alpha}^2 + \tilde{\beta}^2) \tilde{T} \tilde{A}^{-1/2} dr. \quad (5.38)$$

Note that the value of the Higgs field itself, and hence also its vacuum expectation value, remain unaltered after this rescaling.

In the unscaled system, we once again find that the parameter $\tilde{\alpha}_1$ is the appropriate quantity that parametrises the families of states, with the analogous physical parameter being the central redshift $z = T_0 - 1$. As before, the value of $\tilde{\omega}$ must be tuned in order to obtain normalisable solutions (containing the desired number of fermion nodes n), and in addition we are required to tune h_0 such that $h \rightarrow v$ as $r \rightarrow \infty$. This leaves only μ ,

λ and v to consider, which, as discussed above, together constitute only two degrees of freedom, i.e. the value of one of these is always determined by the values of the other two. Using these, we wish to construct two parameters that can be used to identify the respective families of states. From a purely numerical perspective, it proves significantly more convenient to choose quantities that are invariant under the rescaling procedure, leading us to take the Higgs vev, v , as one as these, and the other as the Higgs-to-fermion mass ratio:

$$\xi \equiv \frac{m_H}{m_f} = \frac{2\sqrt{2\lambda}}{\mu}. \quad (5.39)$$

We note that this choice is of course far from unique, and an example of an alternative parametrisation is discussed in section 5.5. To summarise, each individual state is therefore uniquely identified by its values of the set of quantities $\{N_f, n, \xi, v, z\}$, with the one-parameter families of states defined using the first four.

In order to generate solutions numerically, we again use MATHEMATICA’s built-in differential equation solver with an explicit Runge-Kutta method, initialising the solver at a starting radius of $r_{\text{in}} = 10^{-8}$, using the small- r expansion (5.32)–(5.36). To determine the values of $\tilde{\omega}$ and h_0 , it would appear that a two-parameter shooting is required, but fortunately we are able to circumvent this by instead implementing a sequential one-parameter shooting procedure. This involves a binary chop in both variables, in which we first determine the value of $\tilde{\omega}$ corresponding to normalisability for a specific value of h_0 , and then subsequently tune h_0 such that the Higgs field asymptotes to its vacuum expectation value, determining a new value for $\tilde{\omega}$ at each step. Although this may be slightly less efficient than employing a more formal two-parameter shooting algorithm, it proves robust enough for us to obtain states with widely-varying values of z , ξ and v .

A few complications do arise, however. First, at strong fermion-Higgs coupling, we find that a low degree of multivaluedness develops, where multiple solutions are present at the same value of $\tilde{\alpha}_1$. This is at most a three-fold degeneracy, however, and is also confined to the binary chop in h_0 , making it fairly straightforward to negotiate by simply performing a broad initial sweep of the expected region. Second, we find that, for large values of ξ , it becomes necessary to switch the order of the sequential one-parameter shooting, such that the value of h_0 is now determined for each $\tilde{\omega}$ value along the binary chop. In addition, we also struggle to obtain low-redshift solutions, in which the Higgs field is only slightly displaced from its vacuum expectation value, for any values of ξ and v . The reasons behind both these features will become clearer when we derive the large- r expansion in section 5.3.2. Of more concern is our inability to generate states with $v \lesssim 0.07$, irrespective of the value of ξ . We are unsure whether this represents a limit below which solutions do not exist, or whether it is purely related to numerical issues. We currently suspect the latter, since indications of a further multivaluedness become apparent below this value, which involves not only the binary chop in h_0 , but also that in $\tilde{\omega}$, resulting in a breakdown of our sequential one-parameter shooting method.

5.2.2 Individual states

Having outlined the numerical approach, we now present particle-like solutions to the Einstein–Dirac–Higgs system, for the two-fermion case, and discuss their overall behaviour. A selection of these are shown in Fig. 5.2, where we plot the profiles of the fermion, metric and Higgs fields, for three states with the values of n , z , ξ and v indicated.

Consider first the two $n = 0$ ground states, which are generated with the same values of ξ and v , but have different values of the central redshift. In both cases, the Higgs field rises monotonically from a constant central value h_0 , and asymptotes towards its vacuum expectation value v , in precisely the manner described earlier. Its motion is guided by the effective Higgs potential (5.26), which acquires a tilt proportional to $\alpha^2 - \beta^2$, causing the extrema of V_{eff} to become displaced relative to their vacuum values. Near the centre of the soliton, this tilt is large enough that only a single minimum in the effective potential remains, as indicated by the curves h_{eq} included alongside the Higgs field. Comparing the two states, we observe that, for the non-relativistic case $z = 0.134$, the Higgs field deviates only slightly from its vacuum expectation value throughout the fermion source, whereas

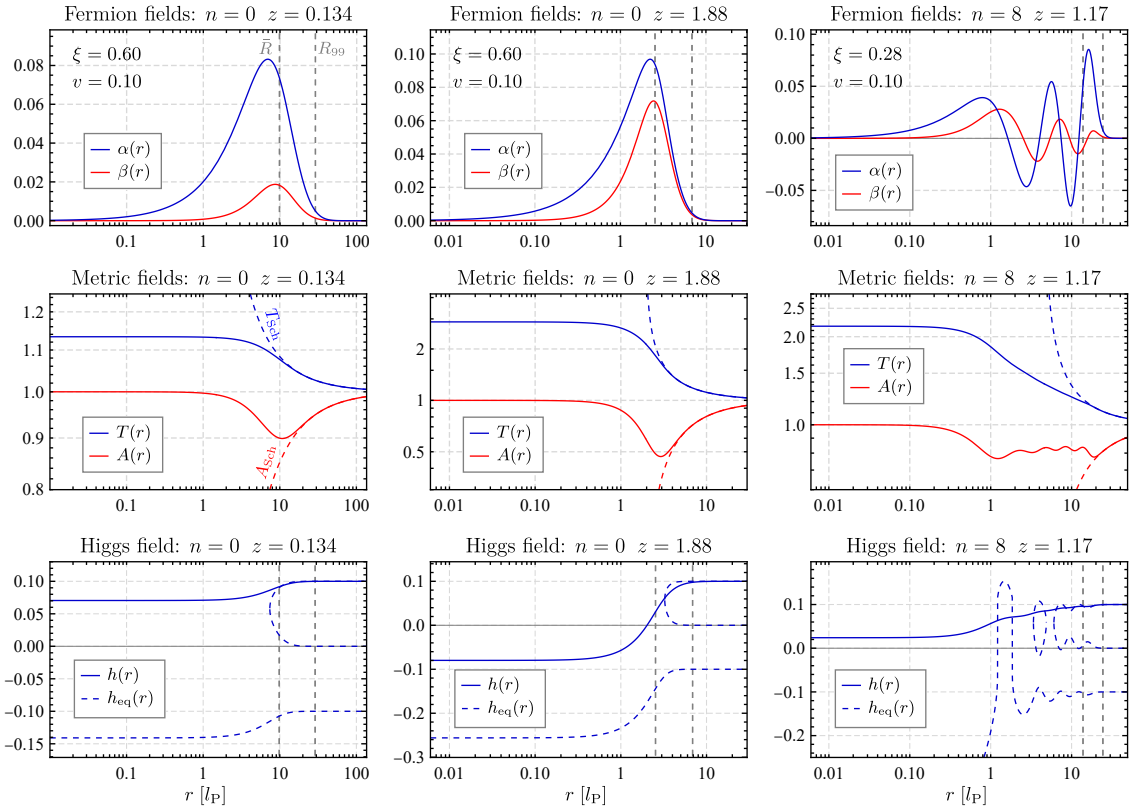


Figure 5.2: Examples of three Einstein–Dirac–Higgs solitons, each constituting a pair of gravitationally-localised neutral fermions interacting with a minimally-coupled Higgs field. Shown are the radial profiles of the fermion, metric and Higgs fields for **(left)** a non-relativistic $n = 0$ ground state, **(middle)** a relativistic $n = 0$ ground state, and **(right)** a fourth even-parity excited state, with the values of v and ξ as indicated. Included alongside the Higgs field profiles are the locations of the extrema of the effective Higgs potential (5.26).

in contrast, for the relativistic state $z = 1.88$, the fermion tilt is such that the Higgs field becomes negative for small values of r . This implies that the local fermion mass μh is also negative within this region, although of course the asymptotic fermion mass $m_f = \mu v$ (as measured by an observer at spatial infinity) remains positive. The reason that the non-relativistic state does not exhibit the same extreme behaviour is that the fermion source is much more diffuse, and hence the tilt in the effective potential acts over a larger radial region, allowing the Higgs field to asymptote gradually towards its maximum. Indeed we find that, as states become increasingly non-relativistic, the required value of h_0 tends towards v .

Also included in Fig. 5.2 is an example of an $n = 8$ excited state, in which eight fermion nodes (four in each field) are located within the wave zone of the soliton. The additional oscillations in the fermion fields have a knock-on effect with regard to the effective Higgs potential, causing it to oscillate around its vacuum configuration, as can be seen from the form of h_{eq} . This does not, however, have a particularly pronounced effect on the profile of the Higgs field itself, with $h(r)$ continuing its monotonic increase, although at a visibly variable rate.

We find that the qualitative behaviour of states is visually similar to that outlined above for all values of ξ and v , at least in the two-fermion system, with the primary difference being the value of h_0 at which the Higgs field must be initiated. The properties of these states, however, and their associated physical observables, are found to vary significantly with ξ and v , as we shall detail in the proceeding sections.

5.2.3 Families of states

Recall that we choose to distinguish Einstein–Dirac–Higgs solitons by their values of the Higgs-to-fermion mass ratio ξ and the Higgs vev v , with each set of states of constant $\{n, \xi, v\}$ forming a one-parameter family of solutions, parametrised by the value of the central redshift z . We note that, at least for the two-fermion system, this parametrisation is always found to be single-valued, even in cases where the unscaled parameter $\tilde{\alpha}_1$ is degenerate. Here we shall consider only the families of $n = 0$ ground states, presenting results outlining their behaviour as both ξ and v are varied.

A summary of this is presented in Fig. 5.3, where we vary the Higgs-to-fermion mass ratio ξ while keeping v at a constant value of 0.1. In the upper panels, we show the fermion mass-energy and ADM mass-radius relations for each family of states, alongside the equivalent curves for the Einstein–Dirac system. As mentioned, generating low-redshift solutions proves numerically challenging, and hence we have been unable to extend the curves significantly in this direction, although it is clear that they will continue to asymptote towards the Einstein–Dirac relations, for the reasons outlined earlier ($h_0 \rightarrow v$ as $z \rightarrow 0$). The most noticeable feature evident from these plots is the rapid increase in

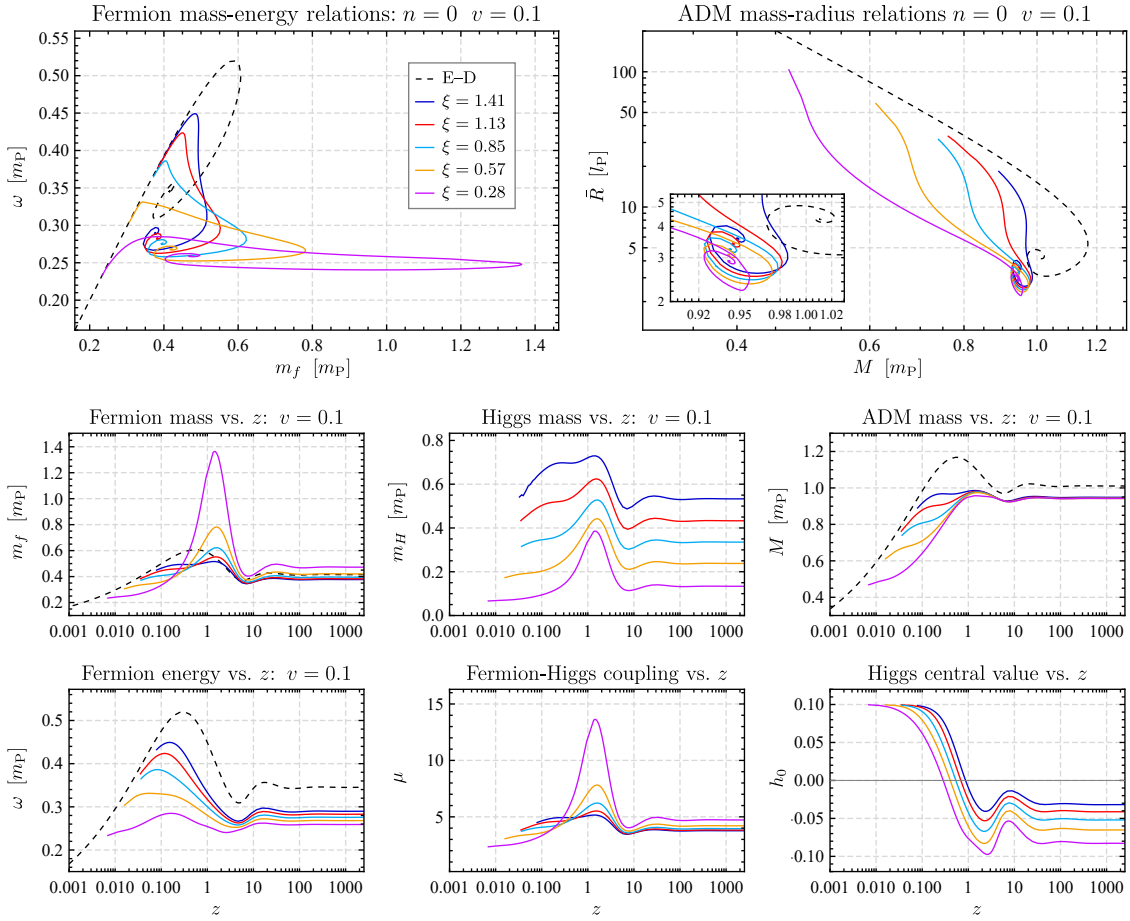


Figure 5.3: Plots summarising the behaviour of the families of $n = 0$ ground-state Einstein–Dirac–Higgs solitons, for a selection of Higgs-to-fermion mass ratios ξ at fixed Higgs vev v . Shown in the upper panels are the fermion mass-energy and ADM mass-radius relations for each respective family, alongside the corresponding Einstein–Dirac curves. Included in the lower panels are the values of m_f , m_H , M , ω , μ and h_0 for each family, all as a function of central redshift z .

fermion mass as ξ is decreased, with the state of maximum mass for the $\xi = 0.28$ curve located at more than twice the fermion mass of its Einstein–Dirac equivalent. Intriguingly, neither the ADM mass nor the fermion energy exhibit a corresponding increase, with in fact a general decrease observed in the latter quantity. We shall discuss this in more detail in section 5.3. Another feature to note from the mass-energy and mass-radius relations is the gradual approach of the curves towards the Einstein–Dirac relations as ξ is increased. This can be understood by recalling that $\xi \propto \sqrt{\lambda}/\mu$, and hence large values of ξ correspond to situations in which the Higgs potential is very steep. Any subsequent tilt to the potential will therefore cause the Higgs field to vary rapidly, and consequently the value of $h(r)$ must never stray far from the Higgs vev. In the limit of $\xi \rightarrow \infty$, we would therefore expect the Higgs field at all radii to be pinned at $h = v$, and thus the Einstein–Dirac case to be recovered.

Also included in Fig. 5.3 are plots showing various quantities, such as the fermion mass

and energy, as a function of redshift. Here, we can see clearly that the fermion mass for small values of ξ exhibits a strong peak, at a redshift just beyond the relativistic transition point, and that this is not accompanied by a similar peak in either the fermion energy or the ADM mass. Since the value of v is held constant for these families, this increase in fermion mass is purely a consequence of entering the strong fermion-Higgs coupling regime (corresponding to large values of μ), which appears only accessible for small values of ξ . Note also from these plots that the central value of the Higgs field, h_0 , shows a general increase with increasing ξ , in agreement with the argument given above, and that the Higgs mass exhibits a similar increase, since $\xi \propto m_H$.

Equivalent plots for families of states with a constant value of ξ but differing values of v are shown in Fig. 5.4. The overall behaviour is similar, with again a strong peak appearing in the fermion mass, now for small values of v , accompanied by a general decrease in the fermion energy. As before, this ultimately arises due to the strong fermion-Higgs coupling in evidence within this redshift region. In addition, the curves are observed to

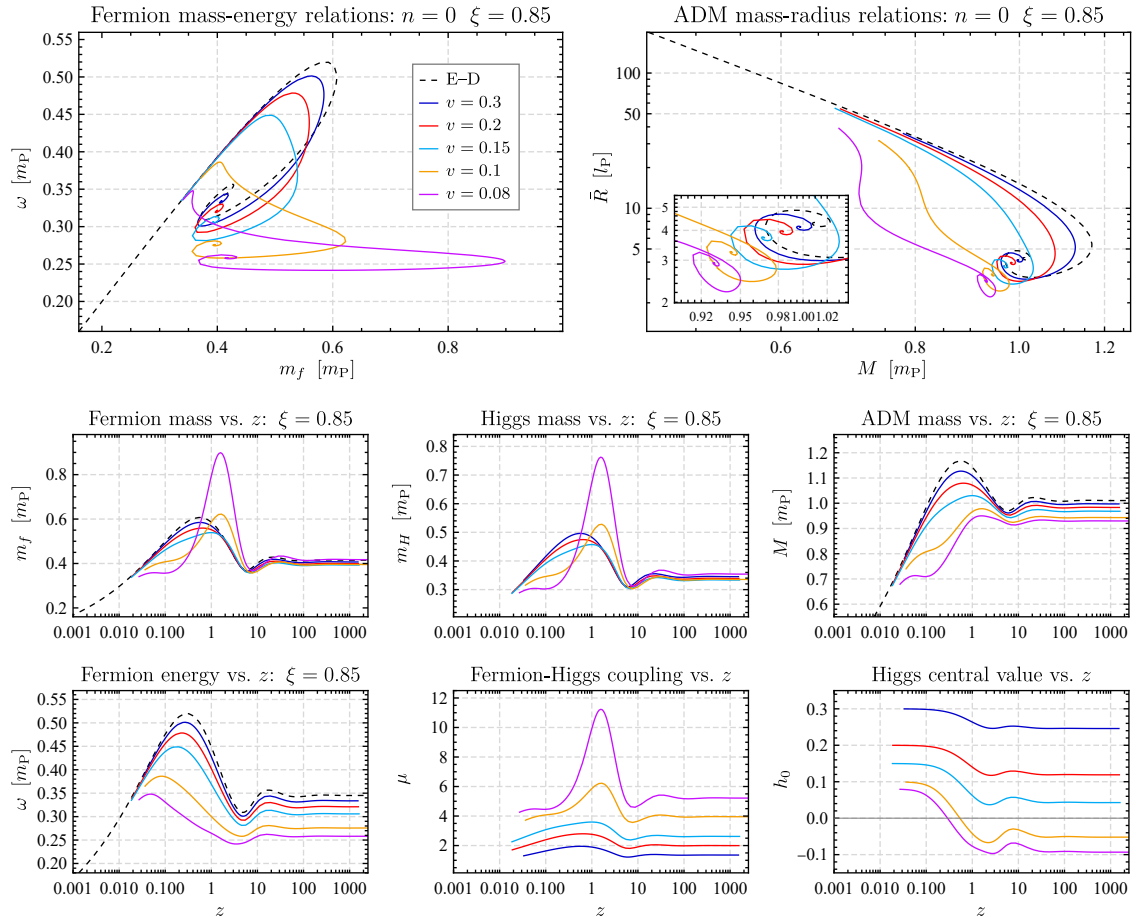


Figure 5.4: Plots summarising the behaviour of the families of $n = 0$ ground states as the value of the Higgs vev v is varied, while keeping the Higgs-to-fermion mass ratio constant. The upper panels show the fermion mass-energy and ADM mass-radius relations for each family, while the lower panels indicate the behaviour of various quantities as a function of redshift.

approach the Einstein–Dirac relation as v increases, a feature that can be understood by noting that an increase in v moves the unstable maxima in the Higgs potential apart, therefore flattening its overall landscape. Finally, we note that the spiralling nature of the families of states implies the presence of an infinite-redshift state, and potentially an accompanying power-law solution, a topic that we shall return to in section 5.4.1.

5.3 Mass-scale separation

The results presented in the preceding section imply the presence of an intriguing feature within the Einstein–Dirac–Higgs system, i.e. that of mass-scale separation, where the ADM mass of a state is no longer approximately proportional to the mass of its constituent fermions.

This is summarised in Fig. 5.5, where we plot the fermion mass of the maximally-bound state against its ADM mass, for a selection of solution families with $v = 0.08$ and the ξ values indicated. Note that the state that is maximally bound in each family is also that of maximum fermion mass. As can be seen, for large values of ξ , the states lie just to the right of the line $M = 2m_f$, indicating that the ADM mass is only slightly lower than the total fermion mass, as anticipated for states which have a small negative binding energy. This is precisely as observed in the Einstein–Dirac system, which as discussed is approached in the limit $\xi \rightarrow \infty$. At $\xi \approx 2$, however, we observe a distinct change in behaviour, with states below this value no longer following the expected relation — instead the fermion mass increases substantially with decreasing ξ , despite the ADM mass continuing to decrease. Thus at the lowest value of ξ shown (representing the limit of our numerics), we find that the total fermion mass of the minimally-bound state outweighs its gravitational mass by a factor of over ten. Although we are unable to verify whether this relationship continues to even lower values of ξ , an extrapolation would suggest that, in the limit of $\xi \rightarrow 0$, we would obtain states in which the total ADM mass is almost negligible, despite the constituent fermions being infinitely massive. Such states would




Figure 5.5: Mass-scale separation in the strongly-coupled Einstein–Dirac–Higgs system, showing the relationship between the fermion mass and ADM mass for the maximally-bound state in each family, with v held constant and ξ varying along the curve. Also included is the corresponding relationship in the Einstein–Dirac system, which is approached as $\xi \rightarrow \infty$, along with the anticipated relation $M = 2m_f$. We note that similar curves exist for larger values of v , although the mass-scale separation in such cases is less pronounced.

correspond to situations in which the Higgs mass also tends to zero, as can be deduced from Fig. 5.3, and where the fermion-Higgs coupling μ is infinitely strong.

There are two main questions that arise from this analysis — why does a strong coupling between the fermion and Higgs sectors result in the somewhat unexpected phenomenon of mass-scale separation, and why is the strong coupling regime only accessible to systems that exhibit small values of both ξ and v . We shall attempt to provide at least partial answers to these questions in the following sections.

5.3.1 Strong coupling

Before exploring possible mechanisms related to the mass-scale separation, we should first verify the claim that this phenomenon is associated with strong coupling μ , and occurs for only small values of both the Higgs-to-fermion mass ratio ξ and the Higgs vev v . Although the figures presented in the preceding sections certainly suggest that this is the case, these only represent specific sections through a 2-dimensional phase space spanned by ξ and v . Exploring this space more thoroughly, however, proves challenging, since it is not computationally feasible to generate entire families of solutions for a large enough selection of parameter values. In order to make progress, we utilise the fact that the minimally-bound state in each family occurs at a roughly constant value of $\tilde{\alpha}_1 = 0.25$, irrespective of the values of ξ and v , and thus we can extract a reasonable comparison between families by considering the properties of states at this value. Note that the minimally-bound state is not the state of maximum mass (this being the maximally-bound state), but is still within the general region where the mass-scale separation is observed.

The results of this analysis are shown in Fig. 5.6, where we plot the values of the fermion mass m_f , ADM mass M and coupling strength μ , for the states with $\tilde{\alpha}_1 = 0.25$, as a function of both ξ and v . These clearly indicate an overall increase in the fermion mass as both ξ and v are decreased, while the ADM mass stays roughly constant (even decreasing somewhat with decreasing v). The fermion-Higgs coupling μ is also observed to be significantly stronger for smaller values of both ξ and v , although the correspondence between this and the fermion mass is perhaps weaker than anticipated. Regardless, it certainly appears that the mass-scale separation is primarily driven by the increase in coupling strength, and that this strong coupling regime is accessible only when ξ and v take small values. It indeed makes sense that small values of ξ should correspond to strong coupling, since recall $\xi \propto 1/\mu$ (assuming the Higgs potential scaling λ remains fairly constant), but the reason why small values of v imply the same is somewhat unclear. Perhaps it is the case that an effective Higgs potential with a strong fermion tilt (recall this is proportional to μ) can only guide the Higgs field to its vacuum expectation value provided v is sufficiently small, although why this should be so is not immediately obvious.

Having established the link to strong coupling, we now consider the question of why




Figure 5.6: Contour plots illustrating the change in the fermion mass m_f (**left**), ADM mass M (**middle**) and fermion-Higgs coupling μ (**right**) as the phase space spanned by the the Higgs-to-fermion mass ratio ξ and the Higgs vev v is explored. Note that the contours become increasingly less smooth as ξ and v are decreased, due to the numerical difficulties in obtaining solutions within this region.

a mass-scale separation occurs at large values of μ . Although we have not been able to obtain a definitive answer to this, we shall discuss a number of features that are potentially related. As a starting point, consider Fig. 5.7, in which we plot the profiles of the fermion and Higgs energy densities, ρ_f and ρ_H , along with the Higgs field itself, for three distinct states with differing values of ξ , but roughly equal values of both v and z . We note that the Higgs energy density is defined using the energy-momentum tensor, and takes the explicit form:

$$\rho_H(r) = \frac{1}{2}A(h')^2 + V(h). \quad (5.40)$$

The three states shown exhibit varying degrees of mass-scale separation (as indicated by the values given in the caption), with the strongest in evidence for the $\xi = 0.03$ solution and none whatsoever for the $\xi = 2.82$ state. There are a number of features to note from these. First, although the Higgs field for the two strongly-coupled states ($\xi = 0.03$ and $\xi = 0.28$) is highly negative within the central regions of the fermion source, this cannot be related to the mass-scale separation since the $\xi = 0.28$ state is consistently more negative than its counterpart. A feature that may be more relevant is the observed outward radial progression of the fermion energy density peak as ξ is decreased. We interpret this as being a direct consequence of the increase in coupling strength, since one would naively expect this to result in an effective attraction between the Higgs and fermion energy densities (the latter of which peaks around the most rapid increase in h), although the precise mechanism via which this occurs is unclear. Another feature to note is related to the relative decay rates of the Higgs and fermion fields. Whereas the Higgs field in all three cases asymptotes towards its vacuum expectation value at approximately the same rate, the fermion density of the $\xi = 0.03$ state decays significantly more quickly when compared to the other two. As a consequence, the fermion density becomes negligible significantly prior to the Higgs field reaching its vacuum expectation value, and thus the fermion source never experiences the asymptotic value of the fermion mass.

We shall return to the issue of whether either of these features can explain the phe-




Figure 5.7: Plots illustrating the change in structure of Einstein–Dirac–Higgs solitons as the fermion–Higgs coupling strength μ is varied. Shown are the radial profiles of the fermion energy density (**left**), the Higgs energy density (**middle**) and the Higgs field itself (**right**), for three states with the values of ξ indicated. All three have the same values of $v = 0.1$ and $z \approx 2.15$. Selected parameter values for these states are: **(a)**: $\{m_f = 3.665, M = 0.8952, \omega = 0.2716\}$, **(b)**: $\{m_f = 1.155, M = 0.9541, \omega = 0.2413\}$ and **(c)**: $\{m_f = 0.4569, M = 0.9894, \omega = 0.3106\}$. Note the appearance of a mass-scale separation for the two states with lowest ξ values.

nomenon of mass-scale separation in section 5.3.3, but before doing so we analyse in more detail the observed disparity between the decay rates of the Higgs and fermion fields, by considering the equations of motion at large r .

5.3.2 Large- r analysis

Our aim in this section is to derive approximate expressions valid for the Higgs and fermion fields at large r , allowing us to analyse more quantitatively their relative decay rates. To do so, we must first construct suitable ansatzes for the general forms of the fields within the evanescent zone. In analogy with the Einstein–Dirac case, we assume the fermion fields to decay exponentially, with an additional subdominant power-law decay, and the metric fields to be approximately Schwarzschild. This leads to the following asymptotic expansions for the fields, valid at large r :

$$\alpha(r) = (D_0 r^\gamma + D_1 r^{\gamma-1} + \dots)e^{-\delta r}; \quad (5.41)$$

$$\beta(r) = (E_0 r^\gamma + E_1 r^{\gamma-1} + \dots)e^{-\delta r}; \quad (5.42)$$

$$T(r) = 1 + M r^{-1} + \dots; \quad (5.43)$$

$$A(r) = 1 - 2M r^{-1} + \dots, \quad (5.44)$$

where γ , δ , D_0 , D_1 , E_0 and E_1 are constants, the values of which are to be determined. This leaves only the Higgs field to consider, for which there appear two options: a simple power-law decay, or an exponential decay similar to that above. We find that a power-law decay is in fact incompatible with the equations of motion, in that there is no way to balance the potential term in the Higgs equation (5.23) at leading order, and thus we quickly conclude that the Higgs field must exhibit an exponential decay towards its vacuum expectation value, taken to be of the following form:

$$h(r) = v - (C_0 r^\sigma + C_1 r^{\sigma-1} + \dots)e^{-\epsilon r}, \quad (5.45)$$

where again the values of the decay rates σ and ϵ , along with the coefficients C_0 and C_1 are to be determined.

Substituting these expansions into the two Dirac equations (5.12) and (5.13), retaining only leading order terms (r^γ), results in the following relations:

$$-D_0 r^\gamma e^{-\delta r} \delta = -(\omega + \mu v) E_0 r^\gamma e^{-\delta r}; \quad (5.46)$$

$$-E_0 r^\gamma e^{-\delta r} \delta = (\omega - \mu v) D_0 r^\gamma e^{-\delta r}, \quad (5.47)$$

which can be combined to give expressions for the relative magnitudes of the leading-order fermion coefficients D_0 and E_0 , along with the fermion exponential decay rate δ :

$$\frac{D_0}{E_0} = \sqrt{\frac{\mu v + \omega}{\mu v - \omega}} \equiv \Lambda; \quad (5.48)$$

$$\delta = \sqrt{\mu^2 v^2 - \omega^2}. \quad (5.49)$$

We can then obtain an expression for the fermion power-law decay rate γ by evaluating the Dirac equations at next-to-leading order ($r^{\gamma+1}$), from which we extract the following two relations:

$$\gamma D_0 + \delta M D_0 - \delta D_1 = \frac{\kappa D_0}{2} - (\omega + \mu v) E_1 - \omega M E_0; \quad (5.50)$$

$$\gamma E_0 + \delta M E_0 - \delta E_1 = -\frac{\kappa E_0}{2} + (\omega - \mu v) D_1 + \omega M D_0. \quad (5.51)$$

Multiplying the second equation by Λ , and combining with the first, we can eliminate both D_1 and E_1 to obtain an explicit expression for γ :

$$\begin{aligned} \gamma D_0 + \delta M D_0 + \gamma D_0 + \delta M D_0 &= \frac{\kappa D_0}{2} - \omega M E_0 - \frac{D_0 \kappa}{2} + \omega M \Lambda D_0 \\ 2\gamma + 2\delta M &= \omega M (\Lambda - \Lambda^{-1}) \\ \gamma &= M \frac{2\omega^2 - \mu^2 v^2}{\sqrt{\mu^2 v^2 - \omega^2}}. \end{aligned} \quad (5.52)$$

Note that the Higgs field (beyond its vacuum expectation value) has not entered this analysis, implying that the precise form of the Higgs decay towards v does not influence the exponential decay of the fermion fields.

We now turn to the Higgs equation (5.23), which for convenience we reproduce below, in order to obtain information regarding the Higgs decay rates:

$$A h'' - A \left(\frac{T'}{T} - \frac{A'}{2A} - \frac{2}{r} \right) h' = \frac{\kappa \mu}{r^2} T (\alpha^2 - \beta^2) + 4\lambda h (h^2 - v^2). \quad (5.53)$$

Considering first the overarching exponential decay, we can immediately eliminate the possibility that the Higgs field decays more than twice as quickly as the fermion fields

($\epsilon > 2\delta$), since this would require α and β to balance each other at leading order, in turn implying $E_0 = D_0$ and consequently $\omega = 0$. This can be interpreted physically by noting that it should not be possible for the Higgs field to reach its vacuum expectation value before the fermion tilt in V_{eff} has been turned off. This leaves two options: (I) $\epsilon < 2\delta$ and (II) $\epsilon = 2\delta$. We shall consider each of these in turn.

(I): $\epsilon < 2\delta$. In this case, the fermion fields do not contribute to the Higgs equation at large r , and we obtain the following expression by considering the leading order terms:

$$-\epsilon^2 C_0 r^\sigma = -8\lambda v^2 C_0 r^\sigma \quad \implies \quad \epsilon = 2v\sqrt{2\lambda}. \quad (5.54)$$

At next-to-leading order ($r^{\sigma-1}$), we then obtain:

$$\begin{aligned} 2\epsilon^2 M C_0 - \epsilon^2 C_1 + 2\epsilon\sigma C_0 + 2\epsilon C_0 &= -8\lambda v^2 C_1 \\ \implies \quad \sigma &= -1 - 2Mv\sqrt{2\lambda} \end{aligned} \quad (5.55)$$

Using these, along with the form for δ given in (5.49), the relation $\epsilon < 2\delta$ can be rewritten to find a condition on ω :

$$\omega^2 < v^2(\mu^2 - 2\lambda). \quad (5.56)$$

This in turn implies that $\mu^2 > 2\lambda$, which can equivalently be written as $\xi < 2$. Thus we conclude that the fermion density can decay more quickly than the Higgs field only when $\xi < 2$ and when the fermion energy obeys the relation specified above.

(II): $\epsilon = 2\delta$. We now consider the case where the Higgs field decays at exactly half the rate of the fermion fields, i.e. at an equal rate to the fermion density. Now both the fermion and Higgs terms contribute at leading order, with the Higgs equation becoming:

$$-\epsilon^2 C_0 r^\sigma = \mu\kappa(D_0^2 - E_0^2)r^{2\gamma-2} - 8\lambda v^2 C_0 r^\sigma. \quad (5.57)$$

As noted above, the fermion terms cannot balance each other, and therefore we arrive at two further options: (a) $\sigma > 2\gamma - 2$ and (b) $\sigma = 2\gamma - 2$. Both of these can achieve balance at leading order, and we shall again consider each in turn:

(a): $\sigma > 2\gamma - 2$. In this case, the fermion terms do not contribute at leading order, and we thus obtain the relation $\epsilon = 2v\sqrt{2\lambda}$ as in case (I). Setting this equal to 2δ , we obtain the following expression for ω :

$$\omega = v^2(\mu^2 - 2\lambda), \quad (5.58)$$

which again implies $\xi < 2$. At next-to-leading order, the Higgs equation then becomes:

$$(2\epsilon^2 M C_0 + 2\epsilon\sigma C_0 + 2\epsilon C_0) r^{\sigma-1} = \mu\kappa(D_0^2 - E_0^2)r^{2\gamma-2}. \quad (5.59)$$

Again, it appears we have the option to choose whether the fermion terms contribute at this order, but in reality one can show that $\omega < 0$ is implied if they are not retained. Thus we must obtain balance by setting $\sigma = 2\gamma - 1$, resulting in an expression for σ in terms of C_0 :

$$\sigma = \frac{\mu\kappa}{2\epsilon C_0}(D_0^2 - E_0^2) - \epsilon M - 1. \quad (5.60)$$

The condition $\sigma = 2\gamma - 1$ then reduces to:

$$C_0 = \frac{\mu\kappa(D_0^2 - E_0^2)}{8M\omega^2}. \quad (5.61)$$

Overall, this case represents the situation where the exponential decay rates of the Higgs field and fermion density are equal, but the fermion fields decay slightly more quickly due to the smaller power-law dependence. This can only occur for states in which $\xi < 2$ and for which the aforementioned condition on ω is satisfied.

(b): $\sigma = 2\gamma - 2$. In this case, both the fermion and Higgs terms contribute to the Higgs equation at leading order, from which we obtain an expression for C_0 (using $\epsilon = 2\delta$):

$$C_0 = \frac{\mu\kappa(D_0^2 - E_0^2)}{4(\omega^2 + v^2[2\lambda - \mu^2])}. \quad (5.62)$$

Since we require this to be strictly positive ($h(r)$ must approach v from below), a condition on ω can again be extracted:

$$\omega^2 > v^2(\mu^2 - 2\lambda). \quad (5.63)$$

Note that this does not here imply a corresponding condition on ξ . Finally, at next-to-leading order, the Higgs equation becomes:

$$2C_0\epsilon(\epsilon M + \sigma + 1) = \mu\kappa [M(D_0^2 - E_0^2) + 2(D_0D_1 - E_0E_1)], \quad (5.64)$$

from which one could in principle obtain an equation relating D_1 and E_1 . Overall, this case represents the situation where the Higgs field and fermion density decay rates are precisely equal at leading order, including their power-law dependences, which can only occur for states in which the above condition on ω holds.

From the above analysis, we conclude that the relative decay rates of the fermion and Higgs fields depend almost entirely on the values of ξ and ω . To summarise, we find that both the fermion and Higgs fields decay exponentially at large r :

$$\alpha(r) = (D_0r^\gamma + D_1r^{\gamma-1} + \dots)e^{-\delta r}; \quad (5.65)$$

$$\beta(r) = (E_0r^\gamma + E_1r^{\gamma-1} + \dots)e^{-\delta r}; \quad (5.66)$$

$$h(r) = v - (C_0 r^\sigma + C_1 r^{\sigma-1} + \dots)e^{-\epsilon r}, \quad (5.67)$$

where the fermion exponential and power-law decay rates are as follows:

$$\delta = \sqrt{\mu^2 v^2 - \omega^2}; \quad \gamma = M \frac{2\omega^2 - \mu^2 v^2}{\sqrt{\mu^2 v^2 - \omega^2}}. \quad (5.68)$$

There are then three options for the decay profile of the Higgs field: (1) the exponential decay of the Higgs field is less than that of the fermion density ($\epsilon < 2\delta$), (2) the exponential decay rates of the two fields are equal but the Higgs power-law dependence results in a slower decay rate ($\epsilon = 2\delta$, $\sigma = 2\gamma - 1$), and (3) both the exponential and power-law decay rates of the Higgs and fermion fields are equal ($\epsilon = 2\delta$, $\sigma = 2\gamma - 2$). The explicit expressions for these, along with the associated conditions, are summarised as follows:

$$(1): \xi < 2; \quad \omega^2 < v^2(\mu^2 - 2\lambda); \quad \epsilon = 2v\sqrt{2\lambda}; \quad \sigma = -1 - 2Mv\sqrt{2\lambda}; \quad (5.69)$$

$$(2): \xi < 2; \quad \omega^2 = v^2(\mu^2 - 2\lambda); \quad \epsilon = 2v\sqrt{2\lambda}; \quad \sigma = Mv\sqrt{\frac{2}{\lambda}}(\mu^2 - 4\lambda) - 1; \quad (5.70)$$

$$(3): \xi > 0; \quad \omega^2 > v^2(\mu^2 - 2\lambda); \quad \epsilon = 2\sqrt{\mu^2 v^2 - \omega^2}; \quad \sigma = 2M \frac{2\omega^2 - \mu^2 v^2}{\sqrt{\mu^2 v^2 - \omega^2}} - 2. \quad (5.71)$$

Note that, for states with $\xi \geq 2$, the only option available is case (3), where the Higgs and fermion fields exhibit precisely the same decay rates, whereas for $\xi < 2$ all options are accessible to the system, with the precise value of ω determining which is followed. Note also that $v^2(\mu^2 - 2\lambda)$ can be equivalently written as $m_f^2 - m_H^2/4$, and hence the conditions on ω correspond to its value relative to the Higgs and fermion masses.

An illustration of this analysis ‘in action’ is shown in Fig. 5.8, where we plot the values of ω^2 alongside the quantity $v^2(\mu^2 - 2\lambda)$, as a function of redshift, for three families of states with differing values of ξ . For the $\xi = 2.26$ case, we observe that ω is always greater than $v^2(\mu^2 - 2\lambda)$, as predicted by the large- r analysis above, indicating that the fermion and Higgs fields for all states decay at precisely the same rate. For the $\xi = 1.41$




Figure 5.8: Plots showing the values of the (squared) fermion energy ω^2 and the quantity $v^2(\mu^2 - 2\lambda)$, as a function of redshift, for three families of states with differing values of ξ . The relative values of these indicate the relative decay rates of the fermion and Higgs fields. Also included on each plot is the corresponding fermion energy-redshift curve for the Einstein–Dirac system.

family, however, the curves cross at two separate redshift values, and hence different decay profiles apply in different regions, with the points of crossing corresponding to option (2) above. As the value of ξ is decreased further, we observe in the $\xi = 0.57$ case that ω is less than $v^2(\mu^2 - 2\lambda)$ for the majority of the redshift range, only dropping below this value at the non-relativistic end of the curve. This indicates that the fermion density decays more quickly than the Higgs field for almost all states in the family, doing so significantly in the mid-redshift regime.

It should be noted that we have verified the analysis in this section by comparing the large- r profiles of our numerical solutions with the predicted forms in (5.69)–(5.71), with an excellent agreement found with respect to both the power-law and exponential decay rates. As a final remark, we also note that the behaviour illustrated in Fig. 5.8 can also explain some of the peculiarities of our numerical method. In particular, the difficulties associated with generating solutions at low redshift, and also at large ξ , can be understood by noting that the decay rates of the Higgs and fermion fields are precisely equal for these states, and as such our sequential one-parameter shooting method proves unreliable, since the behaviour of the Higgs and fermion fields at large r is linked.

5.3.3 Discussion

We now return to the issue of mass-scale separation, and discuss whether the observed disparity in decay rates with varying ξ is at all related. At first glance, it would certainly appear that a connection is present, since for example the change in large- r behaviour occurs at a value of $\xi = 2$, precisely at the point the mass-scale separation is observed to emerge (see Fig. 5.5). Further evidence in favour of this can be seen in Fig. 5.9, in which we analyse the large- r behaviour of states as a function of both ξ and v , where the solutions correspond to the least bound state in each family (as in Fig. 5.6). The leftmost panel shows the value of ω^2 relative to the quantity $v^2(\mu^2 - 2\lambda)$, with a larger disparity between




Figure 5.9: Contour plots indicating the change in relative decay rate of the Higgs and fermion fields as a function of both v and ξ . Shown are (left) the value of ω^2 relative to the quantity $v^2(\mu^2 - 2\lambda)$, (middle) the ratio of the fermion to Higgs exponential decay rate, and (right) the ratio of the fermion to Higgs power-law decay rate. These latter quantities are both always positive, indicating that the fermion density decays more rapidly than the Higgs field within the regions plotted.

these clearly observed for small values of both ξ and v , indicating that the fermion density for such states decays significantly more quickly than the Higgs field. This is corroborated by the two rightmost plots, which show the relative exponential and power-law decay rates of the two fields, with both of these proving important in determining the overall large- r behaviour. Thus we conclude that it is indeed the case that states in which the mass-scale separation is most pronounced exhibit a large disparity between the decay rates of the fermion and Higgs fields.

Of course this does not necessarily imply that the mass-scale separation is a direct consequence of the decay rate disparity; instead it may be that the latter is simply an outcome of strong coupling, and unrelated to the increase in m_f . Given the lack of a definitive mechanism through which these are connected, we should at least entertain the possibility of alternative explanations. It is therefore instructive to consider the mass-scale separation itself in more detail. In particular, we note that the surprising feature of this phenomenon is not necessarily the increase in fermion mass, but the lack of an associated increase in the ADM mass. To investigate this further, we can analyse the various contributions to the ADM mass by considering the Komar integral, which, for the Einstein–Dirac–Higgs system, is found to take the following volume integral form:

$$M = M_K(\infty) = 8\pi \int_0^\infty \left[2\omega T(\alpha^2 + \beta^2) - \mu h(\alpha^2 - \beta^2) - \frac{r^2}{T} V(h) \right] \frac{1}{\sqrt{A}} dr \quad (5.72)$$

$$= 4\omega - 8\pi \int_0^\infty \left[\mu h(\alpha^2 - \beta^2) + \frac{r^2}{T} V(h) \right] \frac{1}{\sqrt{A}} dr, \quad (5.73)$$

where in the second line we have used the normalisation integral (5.24) to evaluate the first term. The relative contributions of the three terms to the integrand of (5.72) are shown in Fig. 5.10, corresponding to solution (a) in Fig. 5.7, where there is a significant mass-scale separation. From this, we observe that by far the dominant contribution to the ADM mass comes in fact from the term involving the fermion energy ω , with both terms involving the Higgs field negligible in comparison. In particular, we note that the overall contribution from the local fermion mass μh is further reduced by the fact that it changes




Figure 5.10: The relative contributions of the terms involving the fermion energy ω , the local fermion mass μh , and the Higgs potential $V(h)$ to the integrand of (5.72), from which the total ADM mass is calculated. Dashed sections indicate regions in which the quantities become negative. Note that the state shown here corresponds to solution (a) in Fig. 5.7, where the mass-scale separation is in evidence.

sign due to $h(r)$ becoming negative in the inner regions of the fermion source. We therefore conclude that the total ADM mass of states exhibiting mass-scale separation is primarily determined by the value of ω , and in fact at strong coupling is approximately equal to 4ω , which can be verified by noting the parameter values in the caption of Fig. 5.7. Hence an increase in the fermion mass does not directly imply a corresponding increase in the ADM mass, and the two therefore become decoupled.

This does not fully explain the origin of the mass-separation, however, since we would clearly expect the fermion energy ω to increase in proportion to the fermion mass. This is not observed to occur, however, with ω instead tending to decrease as the fermion mass becomes larger. With no knowledge of how ω itself is determined, we are therefore no closer to a satisfactory explanation. Perhaps it is indeed the case that the lack of increase in ω is somehow related to the observed disparity in decay rates, or conceivably the redistribution of the fermion density towards the outer regions of the fermion source. From a semi-classical point of view, one might be tempted to conclude that this latter feature may be responsible, since the bulk of the fermion wavefunction is located within a region where the ‘local’ fermion energy ωT is relatively small. This argument does not prove particularly convincing, however, and hence the true origin of the mass-scale separation remains an open question.

5.4 Other results

Before concluding this chapter, we first briefly present a selection of miscellaneous results relating to the Einstein–Dirac–Higgs system. These constitute a discussion concerning the zonal structure of states, including a derivation of the associated power-law solution, as well as the presentation of an $N_f = 20$ state for which the effect of the fermion self-trapping is investigated.

5.4.1 Zonal structure & infinite-redshift states

As mentioned previously, the spiralling behaviour exhibited by the families of Einstein–Dirac–Higgs solitons suggests the presence of a zonal structure similar to that outlined in section 2.5, along with the associated existence of infinite-redshift states. In the Einstein–Dirac case, much of the underlying structure can be understood in relation to an analytic ‘power-law’ solution, around which the fields oscillate within the power-law zone, and which sources the small- r expansion used to obtain infinite-redshift states. We shall here derive the equivalent solution valid in the Einstein–Dirac–Higgs system, the equations of motion for which we restate below, for convenience:

$$\sqrt{A} \alpha' = +\frac{\kappa\alpha}{2r} - (\omega T + \mu h)\beta; \quad (5.74)$$

$$\sqrt{A} \beta' = -\frac{\kappa\beta}{2r} + (\omega T - \mu h)\alpha; \quad (5.75)$$

$$\frac{1}{r^2} (-1 + A + rA') = 8\pi G \left[-\frac{|\kappa|\omega}{r^2} T^2 (\alpha^2 + \beta^2) - \frac{1}{2} A(h')^2 - V(h) \right]; \quad (5.76)$$

$$\frac{1}{r^2} \left(-1 + A - 2rA \frac{T'}{T} \right) = 8\pi G \left[\frac{|\kappa|}{r^2} T \sqrt{A} (\alpha\beta' - \alpha'\beta) + \frac{1}{2} A(h')^2 - V(h) \right]; \quad (5.77)$$

$$Ah'' - A \left(\frac{T'}{T} - \frac{A'}{2A} - \frac{2}{r} \right) h' = \frac{|\kappa|\mu}{r^2} T (\alpha^2 - \beta^2) + \frac{dV}{dh}. \quad (5.78)$$

In the Einstein–Dirac system, the power-law solution is obtained by setting the fermion mass equal to zero, but of course doing so here would imply that the Higgs field vanishes at all radii. We therefore instead seek an approximate solution in which the local fermion mass μh is assumed to be negligible in comparison to the local fermion energy ωT , with the aim of applying this to the inner regions of high-redshift states. In this case, we can neglect the fermion mass terms in the two Dirac equations, and also the terms involving the Higgs field in the Einstein equations, since these are related to the Higgs energy density. The first four equations above hence reduce to precisely the massless Einstein–Dirac system, resulting in the same power-law dependences for the fermion and metric fields as those detailed in section 2.5:

$$\begin{aligned} \alpha(r) &= \sqrt{\frac{\omega}{12\pi G N_f^2 \kappa_-}} r; & A(r) &= \frac{1}{3}; \\ \beta(r) &= \sqrt{\frac{\omega}{12\pi G N_f^2 \kappa_+}} r; & T(r) &= \frac{1}{\omega} \sqrt{\frac{N_f^2}{4} - \frac{1}{3}} r^{-1}, \end{aligned} \quad (5.79)$$

where $\kappa_{\pm} = \kappa/2 \pm \sqrt{1/3}$. This leaves only the Higgs equation to consider. Substituting the power-law forms for α , β , A and T into (5.78), we obtain:

$$\frac{1}{3} h'' + \frac{1}{r} h' = \frac{1}{r} \frac{\mu}{6\sqrt{3}\pi G \kappa \sqrt{\kappa_+ \kappa_-}} \frac{1}{\kappa \sqrt{\kappa_+ \kappa_-}} + 4\lambda(h^3 - v^2 h). \quad (5.80)$$

To make further progress, we assume the solution to be valid only at small r (certainly the power-law zone of high-redshift solutions should occur within that region), and in addition assume the Higgs field to also exhibit a simple power-law dependence at leading order, of the form $h(r) = h_p r^\gamma$. Thus, in order for the Higgs energy density not to contribute to the Einstein equations at small r , we require $\gamma > 0$, and hence we can immediately neglect the terms arising from the potential in (5.80). We can then straightforwardly solve the resulting equation for h to obtain:

$$h(r) = \frac{\mu}{6\sqrt{3}\pi G \kappa \sqrt{\kappa_+ \kappa_-}} \frac{1}{\kappa \sqrt{\kappa_+ \kappa_-}} r - \frac{c_1}{2r^2} + c_2, \quad (5.81)$$

where c_1 and c_2 are constants. As argued, the leading order power of h must be greater than zero, forcing us to take $c_1 = 0$, and we therefore arrive at our final expression for the




Figure 5.11: An illustration of the zonal structure of a high-redshift ground-state Einstein–Dirac–Higgs soliton with the parameter values indicated. The behaviour is similar to the Einstein–Dirac case, with three distinct zones in evidence: the core, the power-law zone and the evanescent zone. Within the power-law zone, the fields approximate a simple power-law dependence, in agreement with the forms derived. Also included is the predicted value of h' within the power-law zone.

Higgs field within the power-law zone:

$$h(r) = \frac{\mu}{6\sqrt{3}\pi G} \frac{1}{\kappa\sqrt{\kappa_+\kappa_-}} r + c_2. \quad (5.82)$$

We can verify the validity of the expressions derived above by analysing the structure of high-redshift states, an example of which is presented in Fig. 5.11. From this we see that indeed a power-law zone develops at high redshift, around which the field perform small-amplitude oscillations, as a consequence of the fermion self-trapping effect discussed in Chapter 3. We also observe that, within the power-law zone, the fields exhibit the general power-law dependences predicted from the above analysis, with in particular h' found to oscillate around a constant value corresponding to that implied from (5.82).

In analogy with the Einstein–Dirac system, we conjecture that the power-law zone of Einstein–Dirac–Higgs states will gradually encroach upon the central core region as the redshift is increased, implying that the power-law zone extends all the way to $r = 0$ at $z = \infty$. This indeed turns out to be the case, and an example of an infinite-redshift solution, located at the centre of one of the spiralling families, is shown in Fig. 5.12. This is generated by replacing the regular asymptotic small- r expansions (5.27)–(5.31) with the power-law expressions derived above. One can clearly see that the Higgs field in this case must be initialised at $r = 0$ with a non-zero slope, and that the metric field T diverges as $r \rightarrow \infty$. In addition, this state contains a central singularity, as can be verified by evaluating the Ricci scalar, which in the Einstein–Dirac–Higgs system takes the following general form:

$$R = 8\pi G \frac{\mu h |\kappa|}{r^2} T(\alpha^2 - \beta^2) - A(h')^2 - 4V(h). \quad (5.83)$$

Substituting in the power-law dependences given in (5.79) and (5.82), we then obtain:

$$R = \pm \frac{4\mu c_2}{\sqrt{9\kappa^2 - 12}} \frac{1}{r} + \mathcal{O}(r^0), \quad (5.84)$$

where \pm indicates positive or negative parity. This indeed diverges as $r \rightarrow 0$, and thus we




Figure 5.12: An example of an infinite-redshift state in the Einstein–Dirac–Higgs system, which lies at the centre of the $\xi = 0.28$, $v = 0.10$ spiral. Shown are the radial profiles of the fermion fields (**left**), metric fields (**middle**) and Higgs field (**right**). Note that $T(r)$ is singular at $r = 0$ and $h(r)$ is initialised with a non-zero slope.

conclude that the infinite-redshift states in the Einstein–Dirac–Higgs system are singular. Note that there is an intriguing caveat to this statement, in that it may be possible for c_2 to equal zero, i.e. for the Higgs field to be initialised at precisely $r = 0$, but with a non-zero initial slope. In such a case, the Ricci scalar would then lead like a constant, and hence no central singularity would be present. This of course relies on the assumption that infinite-redshift states exist where $h(0) = 0$, which is not guaranteed.

5.4.2 Many-fermion states

So far in this chapter, we have presented numerical solutions solely to the two-fermion system, although the equations of motion, along with the majority of the analytic analysis, have been derived for the generalised many-fermion case. Although in principle solutions can therefore be constructed for any (even) value of N_f , numerical complications arise when extending the system to higher numbers of fermions, with in particular our sequential one-parameter shooting method proving less reliable. We are therefore unable to generate many-fermion states consistently, limiting the analysis that can be performed.

Despite this, we here present an example of a high-redshift $N_f = 20$ solution, which is shown in Fig. 5.13. As can be seen, the fermion self-trapping effect is the dominant feature of this state, with the fermion density splitting into a series of peaks located at the radii of stable photon spheres. Considering the Higgs field, we note that, although the effective Higgs potential oscillates significantly throughout the solution, as evidenced by the locations of the extrema in V_{eff} (represented by the curves h_{eq}), the Higgs field itself nonetheless remains relatively close to its vacuum expectation value throughout the fermion source. It no longer, however, increases strictly monotonically, with small decreases observed between the fermion shells. The Higgs energy density is also interesting to note, in that the kinetic part of this becomes concentrated around the locations of the stable photon spheres. Whether this constitutes a trapping of the Higgs field is debatable, however, since it may be the case that its profile is simply dictated by the distribution of the fermion source, as opposed to gravitational considerations.

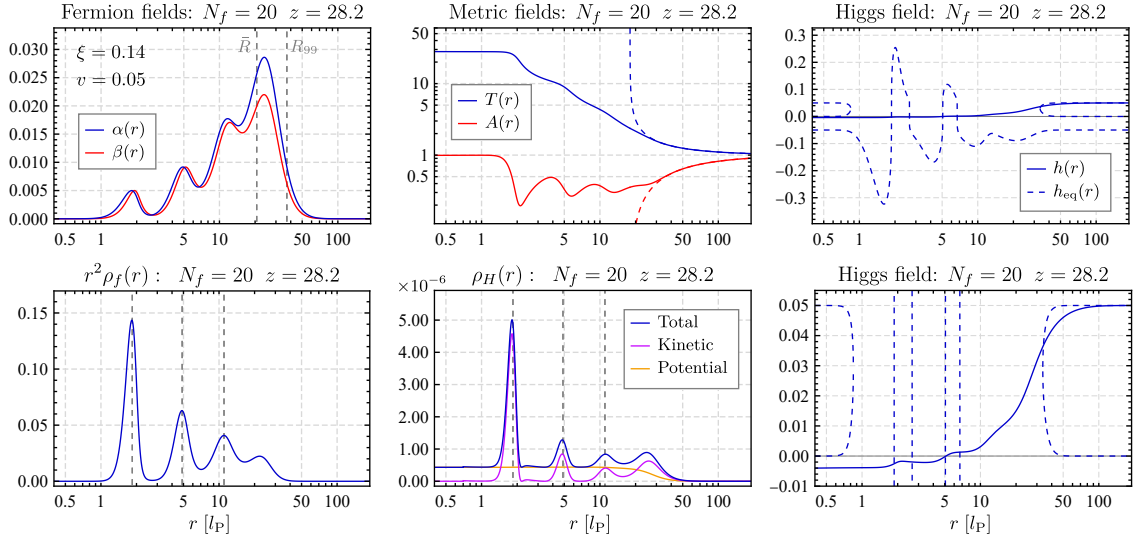


Figure 5.13: An example of a high-redshift $N_f = 20$ state, in which the fermion self-trapping effect is clearly evident. Shown are the radial profiles of **(top)** the fermion, metric and Higgs fields, and **(bottom)** the radial fermion energy density, the Higgs energy density, along with a more detailed view of the Higgs field. Note that the dashed lines in the energy density plots represent the locations of stable photon spheres.

Unfortunately, we have not been able to perform a more thorough analysis of the many-fermion system, and so it remains unclear whether effects such as the mass-scale separation also apply to states with higher fermion numbers. We do note that the $N_f = 20$ state presented here does not appear to exhibit this behaviour (as indicated by the parameter values given in Appendix B), despite having small values of both ξ and v . A detailed analysis, however, is beyond our capabilities.

5.5 Discussion

In this chapter, we have presented particle-like solutions to the minimally-coupled Einstein–Dirac–Higgs system, and analysed the properties of the two-fermion states in detail. Somewhat surprisingly, we discover that, when the fermion-Higgs coupling is sufficiently strong, the various mass scales in the system can become decoupled, with the ADM mass of a state no longer remaining in proportion to the total fermion mass of its constituents. Although we are unable to provide a definitive explanation for this phenomenon, it may be related to a disparity in decay rates that arises at strong coupling, which we have derived using a large- r analysis. There are a number of points concerning this, and other features, that are worthy of further discussion.

First, we should consider in more detail the possible implications of the observed mass-scale separation. In the most extreme cases, we find that, while the overall gravitational mass of the state is roughly of the order of one Planck mass, the total mass of

the constituent fermions exceeds this by a factor of more than ten. Hence the majority of the mass contained within the fermion source does not appear to gravitate, at least from the perspective of an observer at spatial infinity, and thus the results of a gravitational measurement would conclude the fermion mass to be much lower than its true value. Of course this relies on the assumption that the fermion mass should rightly be measured by its asymptotic value. Perhaps instead one should average over the local fermion mass μh , weighted for example by the probability distribution of the fermion source, and hence obtain an alternative measure for m_f . This could potentially lower the overall fermion mass, particularly at strong coupling where the fermion density decays prior to the Higgs field reaching its vacuum expectation value. Which approach is correct depends largely on how the fermion mass is actually measured by an observer - would a non-gravitational measurement probe the inner regions of the fermion source, or simply return its asymptotic value? The answer is somewhat unclear, and may even depend on the type of measurement conducted. Of course it is important to note that this issue does not arise when considering Standard Model fermions, due to their point-like structure, but such complications appear to be an inevitable consequence of considering particles with a finite spatial extent.

On a more technical note, recall that we choose to identify the various families of Einstein–Dirac–Higgs solitons by the values of the Higgs-to-fermion mass ratio ξ and the Higgs vev v . This is purely for numerical convenience, however, and it should in principle be possible to instead parametrise families by any two independent combinations of μ , λ and v . An example of such an alternative parametrisation is presented in Fig. 5.14, which shows the family of states defined by the values $v = 0.30$ and $\lambda = 0.053$. We have only been able to obtain a few states along the curve, due to the difficulty in generating solutions with a specified value of λ , but it is nonetheless clear that a similar spiralling behaviour exists within this parametrisation. Of perhaps more interest would be to construct families where one of the defining parameters is μ , since this would allow a more rigorous analysis of the strong coupling regime, and by extension the mass-scale separation. Unfortunately this does not prove numerically viable, at least via the method utilised here.

Another topic of particular interest is the stability and dynamical evolution of Einstein–




Figure 5.14: An example of an alternative parametrisation for the families of Einstein–Dirac–Higgs solitons, showing the fermion mass-energy relation for a family of states defined by the values of $v = 0.30$ and $\lambda = 0.053$. Also included is the corresponding curve for the Einstein–Dirac system. Note that only a small number of states have been included, due to numerical issues relating to their construction.




Figure 5.15: The binding energy, as a function of soliton radius \bar{R} , for the family of Einstein–Dirac–Higgs states with parameter values $v = 0.08$ and $\xi = 0.14$. Observe that this entire curve is bound, due to the high degree of mass-scale separation present, in contrast to the Einstein–Dirac case (inset). The dashed portions of the curves indicate the branches along which states are assumed to be stable.

Dirac–Higgs solitons. A rough indication of this can be obtained by considering the total binding energy of states, which recall is defined as $E_b = M - 2m_f$, i.e. the difference between the energy of the state and that of two spatially separated fermions. For states in which the mass-scale separation is most pronounced, one would expect the binding energy to be highly negative, since the fermion mass far outweighs the ADM mass. Indeed it is even the case that entire families of solutions can become bound, as is illustrated in Fig. 5.15. Of course a negative binding energy does not necessarily imply stability, and instead we would expect states with redshifts above that of the maximum-mass state to be unstable, as is the case in the Einstein–Dirac system. Nonetheless, we would still expect the stable branch to include states with high degrees of mass-scale separation. The ultimate fate of the states on the unstable branch is unclear, although a negative binding energy should ensure that they at least remain as localised objects, perhaps forming black holes or decaying to the stable branch. A numerical evolution of these would certainly prove enlightening, and in particular it would be interesting to note how the profile of the Higgs field changes as the solutions evolve.

As a final remark, we mention that we have here modelled the Higgs field as a real scalar field. In the Standard Model, however, the Higgs is formally a complex quantity, although it is subsequently made real by employing the unitary gauge. Such a concept does not apply here, however, and hence one might wonder whether a complex scalar field would be more physically realistic. We find, however, that the inclusion of a field $h = H e^{i\varphi}$ is incompatible with the requirement of regularity at $r = 0$, and as a result the phase φ is forced to take a constant value for all r . We are then free to set this constant to zero by a redefinition of φ , and hence h becomes purely real. Thus we are justified in our choice of modelling the Higgs as a real scalar field.

Chapter 6

Singular solutions

Returning to the Einstein–Dirac system, we present in this chapter a number of new singular solutions to the FSY equations (2.28)–(2.31) that nonetheless are localised and entirely normalisable. These are found by varying the small- r expansion used to generate numerical solutions. As we shall show, these singular states often exhibit somewhat unusual properties, such as a negative gravitational mass, but, despite this, many can in fact be related to the original non-singular solutions initially obtained by Finster, Smoller & Yau [3], and discussed here in chapters 2–4. For the purposes of this discussion, we shall here refer to these as the ‘FSY states’. In section 6.2, we shall also present a new analytic solution to the Einstein–Dirac system, valid when $\omega = 0$ and $A(r) = 1$, which we understand constitutes the first analytic particle-like solution to the system.

6.1 Small- r analysis

We begin by noting that the small- r asymptotic expansion (2.35)–(2.38) used to generate the FSY states is not guaranteed to be unique. Indeed, we have already seen that the infinite-redshift states obey an entirely different expansion, as discussed in section 2.5. One might therefore wonder whether other types of localised solutions can be constructed in the Einstein–Dirac system, for which the asymptotic small- r expansions take alternative forms yet to be considered. We have therefore undertaken a rigorous analysis of the small- r behaviour of the FSY equations, and found that there are in fact twelve(!) possible asymptotic expansions, with all but three of these resulting in normalisable, localised states. Only two of the twelve expansions are non-singular, however, and only one of these (the FSY case) can be used to construct normalisable solutions. Thus, despite the multitude of new states presented in this chapter, the FSY states remain the only localised, non-singular solutions in the Einstein–Dirac system.

We shall briefly outline the method through which we obtain these alternative small- r

expansions, noting that the calculation is straightforward but somewhat lengthy. Although this was initially performed by hand, we have in fact been able to fully automate the procedure, allowing us to use MATHEMATICA to verify our results. For convenience, we first restate the FSY equations for general κ (valid for both positive and negative parity):

$$\sqrt{A} \alpha' = +\frac{\kappa\alpha}{2r} - (\omega T + m)\beta; \quad (6.1)$$

$$\sqrt{A} \beta' = -\frac{\kappa\beta}{2r} + (\omega T - m)\alpha; \quad (6.2)$$

$$-1 + A + rA' = -\xi|\kappa|\omega T^2 (\alpha^2 + \beta^2); \quad (6.3)$$

$$-1 + A - 2rA\frac{T'}{T} = \xi|\kappa|T\sqrt{A} (\alpha\beta' - \alpha'\beta), \quad (6.4)$$

where we have defined $\xi = 8\pi G$. In addition, the normalisation condition is:

$$4\pi \int_0^\infty \frac{T}{\sqrt{A}} (\alpha^2 + \beta^2) dr = 1. \quad (6.5)$$

We note that the requirement of normalisability is not necessarily guaranteed even in the case of a localised state, since we must also ensure that the integrand here does not become infinite as $r \rightarrow 0$. With regard to our small- r expansions, one might therefore expect that imposing normalisability would eliminate a number of options, but in fact we find that all expansions compatible with the equations of motion are automatically normalisable at $r = 0$, a perhaps surprising result for which we lack an explanation.

Our strategy for analysing equations (6.1)–(6.4) at small r can be summarised as follows. Assuming that all fields lead with a simple power-law dependence, we can write their small- r asymptotic expansions as:

$$\alpha(r) = \alpha_0 r^{a_0} + \alpha_1 r^{a_1} + \dots; \quad (6.6)$$

$$\beta(r) = \beta_0 r^{b_0} + \beta_1 r^{b_1} + \dots; \quad (6.7)$$

$$A(r) = A_0 r^{c_0} + A_1 r^{c_1} + \dots; \quad (6.8)$$

$$T(r) = T_0 r^{d_0} + T_1 r^{d_1} + \dots, \quad (6.9)$$

where our aim is to determine all possible real values for the leading-order powers a_0 , b_0 , c_0 and d_0 . In order to do so, we substitute these expressions into the FSY equations, keeping only terms that might conceivably contribute at leading order. This requires some degree of care, since it may be the case that one or more fields leads like a constant, and hence next-to-leading order terms must be included when considering derivatives. The four equations at leading order are therefore:

$$a_0 \sqrt{A_0} \alpha_0 r^{a_0 + \frac{c_0}{2} - 1} + \delta_{a_0} \alpha_1 \sqrt{A_0} r^{a_1 + \frac{c_0}{2} - 1} = \frac{\kappa}{2} \alpha_0 r^{a_0 - 1} - \omega T_0 \beta_0 r^{b_0 + d_0} - m \beta_0 r^{b_0}; \quad (6.10)$$

$$b_0\sqrt{A_0}\beta_0r^{b_0+\frac{c_0}{2}-1} + \delta_{b_0}b_1\sqrt{A_0}\beta_1r^{b_1+\frac{c_0}{2}-1} = -\frac{\kappa}{2}\beta_0r^{b_0-1} + \omega T_0\alpha_0r^{a_0+d_0} - m\alpha_0r^{a_0}; \quad (6.11)$$

$$-1 + (1 + c_0)A_0r^{c_0} + \delta_{c_0}c_1A_1r^{c_1-1} = -\xi|\kappa|\omega \left(T_0^2\alpha_0^2r^{2a_0+2d_0} + T_0^2\beta_0^2r^{2b_0+2d_0} \right); \quad (6.12)$$

$$\begin{aligned} & -1 + (1 + c_0)A_0r^{c_0} - 2d_0A_0r^{c_0} - \delta_{d_0}2d_1A_0\frac{T_1}{T_0}r^{c_0+d_1} \\ & = \xi|\kappa|\sqrt{A_0}T_0 \left[a_0\alpha_0\beta_0r^{a_0+b_0+\frac{c_0}{2}+d_0-1} - b_0\alpha_0\beta_0r^{a_0+b_0+\frac{c_0}{2}+d_0-1} \right. \\ & \quad \left. + \delta_{a_0}a_1\alpha_1\beta_0r^{a_1+b_0+\frac{c_0}{2}+d_0-1} - \delta_{b_0}b_1\alpha_0\beta_1r^{a_0+b_1+\frac{c_0}{2}+d_0-1} \right], \quad (6.13) \end{aligned}$$

where the inclusion of the Kronecker delta functions ensures that the next-to-leading order terms are only present when the appropriate leading-order power is equal to zero. We then consider each equation in turn, determining the various possible ways in which balance between terms can be achieved, along with the resulting leading-order equations obeyed by the coefficients. Finally, the constraints arising from each equation are combined, in order to identify the options that are mutually compatible. It should be clear that this calculation quickly gets out of hand, due to the sheer number of possible combinations, and hence it is advantageous to automate the procedure. Even then, care must be taken to ensure that the computational cost remains manageable.

Overall, as mentioned, we obtain a total of twelve asymptotic expansions that are consistent with the FSY equations, which we shall now consider in turn. Note that, when generating numerical solutions for these, we shall aim to construct states where the number of fermion nodes $n = 0$ and the fermion mass is positive, which on occasion forces us to consider negative parity.

6.1.1 Case [A]: $\{\alpha, \beta, A, T\} \sim \{r^{\kappa/2}, r^{\kappa/2+1}, r^0, r^0\}$ (FSY states)

The first case to consider is where both metric fields $A(r)$ and $T(r)$ are non-zero constants at $r = 0$, while the leading-order powers of $\alpha(r)$ and $\beta(r)$ depend on the fermion number in the manner indicated. This is precisely the ansatz used to generate the FSY states, with the small- r expansion taking the form shown previously in section 2.2:

$$\alpha(r) = \alpha_0r^{\kappa/2} + \mathcal{O}(r^{\kappa/2+2}); \quad (6.14)$$

$$\beta(r) = \frac{1}{\kappa+1}\alpha_0(\omega T_0 - m)r^{\kappa/2+1} + \mathcal{O}(r^{\kappa/2+3}); \quad (6.15)$$

$$A(r) = 1 - \frac{\xi\kappa}{\kappa+1}\omega T_0\alpha_0^2r^\kappa + \mathcal{O}(r^{\kappa+2}); \quad (6.16)$$

$$T(r) = T_0 - \frac{\xi}{2\kappa+2}(2\omega T_0 - m)T_0^2\alpha_0^2r^\kappa + \mathcal{O}(r^{\kappa+2}), \quad (6.17)$$

where we note that this is valid only for positive parity, with the corresponding negative parity expansion having the powers of $\alpha(r)$ and $\beta(r)$ interchanged (a common feature of

all the expansions). One can easily verify that this case is non-singular by calculating the corresponding expansions for the Ricci and/or Kretschmann scalar, giving:

$$R = \xi m \kappa \alpha_0^2 T_0 r^{\kappa-2} + \dots; \quad (6.18)$$

$$K = \frac{\xi^2 \kappa^2 \alpha_0^4 T_0^2}{(1 + \kappa)^2} \left[m^2 (3 - 2\kappa + \kappa^2) - 4m\omega T_0 (1 - \kappa + \kappa^2) \right. \\ \left. + 2\omega^2 T_0^2 (2 + 2\kappa + 3\kappa^2) \right] r^{-4+2\kappa} + \dots \quad (6.19)$$

Provided $\kappa \geq 2$, neither of these will diverge at $r = 0$, and hence we conclude that physically relevant FSY states do not contain a central singularity.

Solutions generated using this expansion have of course been presented throughout this thesis, and we shall therefore refrain from doing so here. There are, however, some features of the asymptotic expansion that are worth emphasising, since these can aid our analysis of the remaining cases below. First, note that there are here two unconstrained leading-order coefficients, α_0 and T_0 , and two free physical parameters m and ω . As mentioned previously, imposing normalisation and asymptotic flatness effectively removes three of these supposed degrees of freedom, and we therefore obtain a one-parameter family of states (for constant n and κ). Thus, by a simple parameter-counting, one can straightforwardly infer the dimensionality of the resulting parametrisation of states. Second, the rescaling method outlined in section 2.2.2 allows us to temporarily fix the values of two of the four ‘free’ parameters (chosen here to be m and T_0), with the physical values of these recovered upon rescaling the solution. Since the rescaling itself is a generic feature of the FSY equations, it is guaranteed to apply to any valid small- r expansion, and hence we can utilise the same method for the other forthcoming cases.

6.1.2 Case [B]: $\{\alpha, \beta, A, T\} \sim \{r^{\kappa/2}, r^{3\kappa/2+1}, r^0, r^0\}$

Note that the expansion for the FSY states above fails to hold when $\omega T_0 = m$, since the leading power of $\beta(r)$ will vanish. There is therefore the possibility of an alternative expansion valid in this case, which we find to be (for positive κ):

$$\alpha(r) = \alpha_0 r^{\kappa/2} + \frac{\xi \kappa m^2 \alpha_0^3}{4\omega(\kappa + 1)} r^{3\kappa/2} + \dots; \quad (6.20)$$

$$\beta(r) = -\frac{\xi \kappa m^3 \alpha_0^3}{2\omega \kappa (\kappa + 1)(2\kappa + 1)} r^{3\kappa/2+1} + \dots; \quad (6.21)$$

$$A(r) = 1 - \frac{\xi \kappa m^2 \alpha_0^2}{\omega(\kappa + 1)} r^\kappa + \dots; \quad (6.22)$$

$$T(r) = \frac{m}{\omega} - \frac{\xi \kappa m^3 \alpha_0^2}{2\omega \kappa (\kappa + 1)} r^\kappa + \dots \quad (6.23)$$

The Ricci and Kretschmann scalars here evaluate to:

$$R = \frac{\xi\kappa m^2 \alpha_0^2}{\omega} r^{\kappa-2} + \dots; \quad K = \frac{3\kappa^2 m^4 \alpha_0^4}{\omega^2} r^{2\kappa-4} + \dots + \mathcal{O}(r^\kappa), \quad (6.24)$$

and hence, for $\kappa \geq 2$, there is therefore no central singularity. Note that, for $\kappa > 4$, the Kretschmann scalar will lead like r^κ , but this is of course still non-singular.

The expansion here contains only a single unconstrained coefficient, α_0 , and thus we would expect any resulting states to represent single points in the m - ω plane, since there are zero degrees of freedom in the system. We are unable, however, to generate any localised solutions using this expansion, regardless of the value of κ , although we have no formal proof that they do not exist. Even if such a state could be constructed, it would not necessarily represent a new solution, but rather a special case of the FSY states in which the central redshift $z = T_0 - 1$ happens to coincide precisely with the value of $m/\omega - 1$. We therefore conclude that this expansion does not produce an alternative class of non-singular states in the Einstein–Dirac system.

6.1.3 Case [C]: $\{\alpha, \beta, A, T\} \sim \{r^1, r^1, r^0, r^{-1}\}$ ($z = \infty$ states)

The next case to consider is that obeyed by the infinite-redshift solutions, where the leading-order expansion is independent of the fermion mass, and thus the coefficients are related to those in the analytic massless ‘power-law’ solution. The expansion at small r takes the following form:

$$\begin{aligned} \alpha(r) &= \sqrt{\frac{2\omega}{3\xi\kappa^2\kappa_-}} r + \dots; & A(r) &= \frac{1}{3} + \dots; \\ \beta(r) &= \sqrt{\frac{2\omega}{3\xi\kappa^2\kappa_+}} r + \dots; & T(r) &= \frac{1}{\omega} \sqrt{\frac{\kappa^2}{4} - \frac{1}{3}} r^{-1} + \dots, \end{aligned} \quad (6.25)$$

where $\kappa_\pm = \kappa/2 \pm \sqrt{1/3}$. This is valid for both positive and negative parity, but note that it requires $|\kappa| > 2/\sqrt{3} \approx 1.1547$. Although this of course encompasses all physically acceptable values of κ , one might nonetheless wonder how the system behaves when $|\kappa| < 2\sqrt{3}$, an issue that will be addressed when considering case [D]. The asymptotic forms for the Ricci and Kretschmann scalars (for positive parity) are:

$$R = \frac{8m}{3\kappa\sqrt{3\kappa^2 - 4}} r^{-1} + \dots; \quad K = -\frac{8}{3} r^{-4} + \dots, \quad (6.26)$$

and hence the expansion is singular. As discussed previously, the localised solutions associated with this expansion reside at the centres of the families of FSY states, and an example can be seen in Fig. 2.5. An interesting point to note, however, is that in the above expansion all leading-order coefficients are constrained, and thus it would appear

that only two degrees of freedom (m and ω) are present. Since imposing normalisation and asymptotic flatness is known to remove three degrees of freedom, this feature is somewhat puzzling. Furthermore, when we generate infinite-redshift states numerically, we are forced to include an additional parameter in the small- r expansion, the value of which must then be determined by a shooting procedure (although its value is found to be very close to unity). The reason for this behaviour, however, is not entirely clear.

6.1.4 Case [D]: $\{\alpha, \beta, A, T\} \sim \{r^\Omega, r^1, r^0, r^{-\Omega}\}$

Here, the leading-order powers in the asymptotic expansion are similar to those in case [C], but now $\alpha(r)$ and $T(r)$ are no longer constrained to lead like r and $1/r$ respectively, with instead their dependences allowed to vary. At small r , the relevant leading-order forms for the fields are:

$$\alpha(r) = \alpha_0 r^\Omega + \dots; \quad \beta(r) = \pm \sqrt{\frac{\omega}{\xi\kappa}} \left(\frac{\sqrt{\Lambda}}{2\kappa + \sqrt{16 - 2\Lambda}} \right) r + \dots; \quad (6.27)$$

$$A(r) = 1 - \frac{\Lambda}{8} + \dots; \quad T(r) = \pm \frac{\Lambda}{2\alpha_0 \sqrt{2\kappa\xi\omega}} r^{-\Omega} + \dots, \quad (6.28)$$

where

$$\Omega = \frac{\sqrt{2\kappa}}{\sqrt{8 - \Lambda}}; \quad \Lambda = \kappa(\kappa + \sqrt{16 - 3\kappa^2}). \quad (6.29)$$

Note that these are valid only for positive parity, and that β_0 and T_0 must have the same overall sign, although this can be either positive or negative. In the latter case, α_0 must also be taken negative in order to ensure a positive probability density. We also find that the value of κ is now constrained to $|\kappa| \leq 2/\sqrt{3}$, noting that this condition cannot be ascertained by considering the forms of the above expansion, but instead arises during their derivation. Hence this case represents the analogue of case [C] for fermion numbers less than ≈ 1.1547 . The Ricci and Kretschmann scalars here evaluate to:

$$R = \pm \frac{\sqrt{\xi\kappa m \Lambda} \alpha_0}{2\sqrt{2\omega}} r^{-2+\Omega} + \dots; \quad (6.30)$$

$$K = \frac{\Lambda}{16} \left(3 + \Lambda^2 + \frac{12\Lambda\kappa^2}{(\sqrt{16 - 2\kappa} + 2\kappa)^2} - \frac{2\kappa\sqrt{2\Lambda}(3 + \Lambda)}{\sqrt{16 - 2\kappa} + 2\kappa} \right) r^{-4} + \dots, \quad (6.31)$$

where the sign of the Ricci scalar is governed by that of β_0 and T_0 . The condition on κ restricts $\Omega \leq 1$, and thus the expansion is singular at $r = 0$.

There is only one unconstrained leading-order coefficient in the system, α_0 , and hence any associated localised states are expected to represent single points in phase space. We find that such solutions do indeed exist, apparently only when the signs of β_0 and T_0 are taken positive, and that they are located at the centres of the FSY spirals with $\kappa \leq 2\sqrt{3}$.

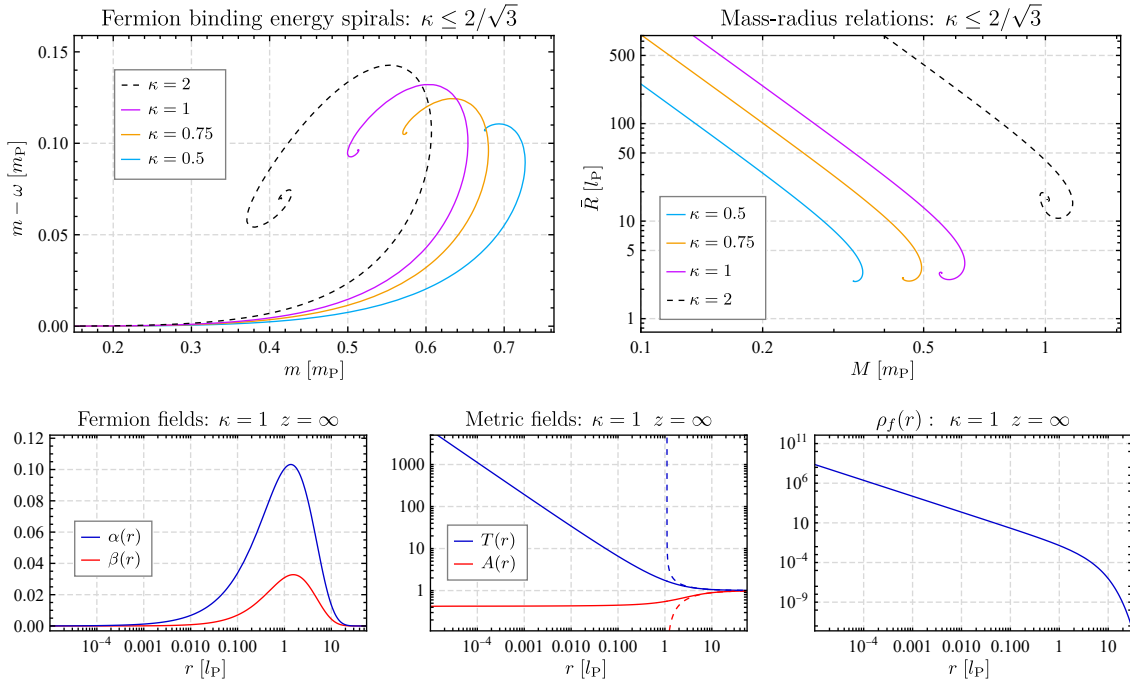


Figure 6.1: Plots summarising the behaviour of the FSY system for $\kappa \leq 2/\sqrt{3}$. The upper panels show the fermion binding energy and mass-radius relations for three families of $n = 0$ FSY states with the values of κ indicated, alongside the equivalent two-fermion curves. At the centre of the spirals lie the states generated using the small- r expansion in case [D], one of which (for $\kappa = 1$) is shown in the lower panels. Note that this is singular, as indicated by the divergence of the fermion energy density ρ_f at small r .

One can therefore interpret these as the infinite-redshift FSY states for fermion numbers below 1.1547. Although such values are entirely unphysical, it is nonetheless interesting to consider the associated behaviour of the system, which is summarised in Fig. 6.1. From this, we observe that the fermion binding energy and mass-radius relations exhibit the usual spiralling structure, although the spiral itself becomes progressively tighter as the fermion number is decreased. Note that, as discussed previously, all states along these curves are in fact singular. We have also included the radial profiles of the infinite-redshift state that lies at the centre of the $\kappa = 1$ spiral, this having been generated using the small- r expansion above.

The existence of separate infinite-redshift states for $\kappa \leq 2/\sqrt{3}$ and $\kappa > 2/\sqrt{3}$ implies that there should be a corresponding difference in the zonal structure of FSY states for the two cases. Indeed, we find that the power-law zone for FSY states with $\kappa \leq 2/\sqrt{3}$ is replaced by an equivalent zone in which the fields approximate the alternative power-law forms given in (6.27)–(6.28). This appears to be related to the fact that neither the mass term nor the fermion energy term contribute to the Dirac equations (6.1) and (6.2) within this new zone, although the precise reason for this is unclear. Given the unphysical nature of the system for $\kappa \leq 2/\sqrt{3}$, however, it is perhaps not worth considering at length.

6.1.5 Case [E]: $\{\alpha, \beta, A, T\} \sim \{r^0, r^0, r^{-2}, r^{-1}\}$

So far, the expansions we have presented have either been previously considered, or equate to slight alterations of limited interest. The remainder of the cases, however, prove more interesting, with their application resulting in new families of states that are entirely separate from those previously studied. The first of these is case [E], where the relevant asymptotic expansion takes the following form (valid for both positive and negative parity):

$$\alpha(r) = \alpha_0 + \frac{\alpha_0^2 + 2\beta_0^2}{\beta_0} \sqrt{\frac{\omega}{\xi|\kappa|(\alpha_0^2 + \beta_0^2)}} r + \mathcal{O}(r^2); \quad (6.32)$$

$$\beta(r) = \beta_0 - \frac{2\alpha_0^2 + \beta_0^2}{\alpha_0} \sqrt{\frac{\omega}{\xi|\kappa|(\alpha_0^2 + \beta_0^2)}} r + \mathcal{O}(r^2); \quad (6.33)$$

$$A(r) = \frac{|\kappa|^3 \xi \alpha_0^2 \beta_0^2}{4\omega(\alpha_0^2 + \beta_0^2)} r^{-2} + \mathcal{O}(r^{-1}); \quad (6.34)$$

$$T(r) = -\frac{\kappa \alpha_0 \beta_0}{2\omega(\alpha_0^2 + \beta_0^2)} r^{-1} + T_1 + \mathcal{O}(r). \quad (6.35)$$

Here, both $\alpha(r)$ and $\beta(r)$ approach non-zero constants at small r , the values of which are unconstrained, whereas $A(r)$ and $T(r)$ now both diverge as $r \rightarrow 0$. Note that, for positive parity, one of either α_0 or β_0 must be negative, in order to ensure $T(0)$ remains positive. The Ricci and Kretschmann scalars for this case are:

$$R = -\frac{\xi \kappa^2 m \alpha_0 \beta_0 (\alpha_0^2 - \beta_0^2)}{2\omega(\alpha_0^2 + \beta_0^2)} r^{-3} + \dots; \quad K = \frac{3\xi^2 \kappa^6 \alpha_0^4 \beta_0^4}{2\omega^2(\alpha_0^2 + \beta_0^2)} r^{-8} + \dots, \quad (6.36)$$

and hence the expansion is singular for all values of κ .

Considering now the potential localised states associated with this expansion, we note that there are two unconstrained coefficients, α_0 and β_0 , which is precisely the same number as in the FSY case, and hence we should expect a single one-parameter family of solutions. For numerical convenience, we choose the quantity that parametrises this family to be $p_0 \equiv \alpha_0/\beta_0 - 1$, since this remains invariant under the rescaling procedure. We find that indeed localised solutions can be constructed for any value of this parameter, from zero to infinity, and in addition that positive fermion mass states can only be obtained for negative parity. We shall here therefore restrict our numerical analysis to the states with $\kappa = -2$, although we emphasise that solutions exist for all values (and both signs) of κ . Note that the numerical procedure for obtaining states is very similar to that employed in the FSY case, with a one-parameter shooting in $\tilde{\omega}$ required to initially obtain an unscaled solution (where $\tilde{\beta}_0 = 1$), after which the physical state is recovered by rescaling the fields and parameters.

The overall behaviour of the family of states is illustrated in Fig. 6.2, where we plot the fermion mass-energy and mass-radius relations, along with a variety of quantities plotted as a function of p_0 . It is clear even from a cursory glance that these states are

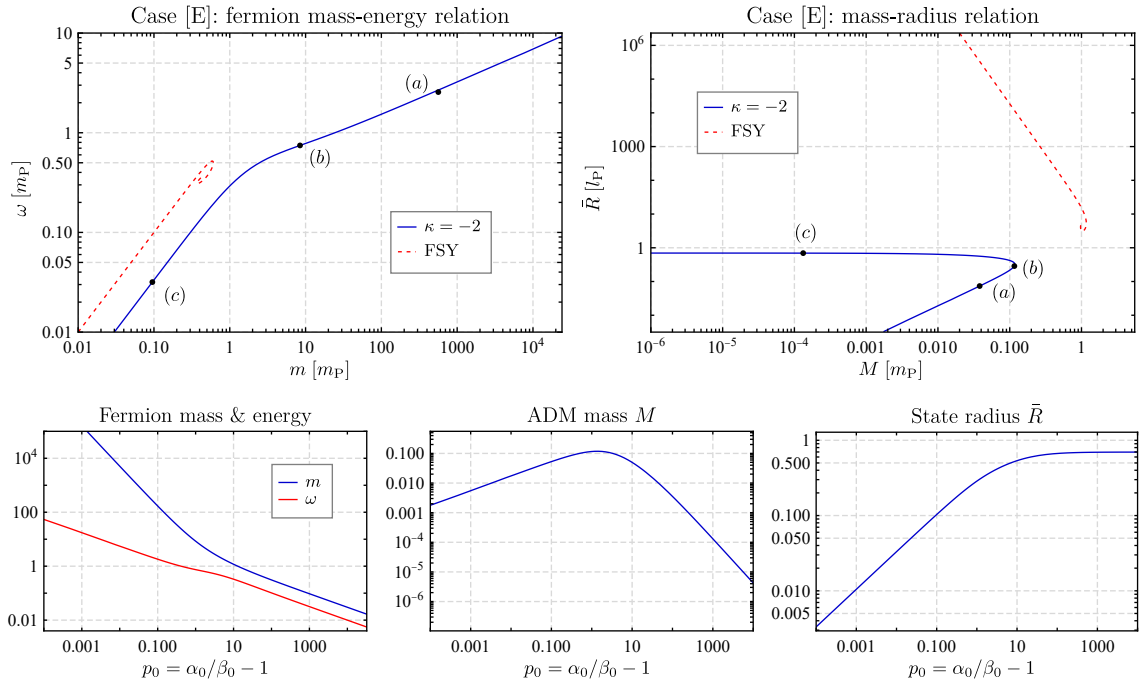


Figure 6.2: Plots summarising the behaviour of the singular states generated using the asymptotic expansion [E]. In the upper panels, we show the fermion mass-energy and mass-radius relations for the one-parameter family of states with $\kappa = -2$, alongside the corresponding curves for the FSY case ($\kappa = 2$, $n = 0$). Included in the lower panels are plots of the fermion mass and energy, the ADM mass, and the average radius as a function of the parameter $p_0 = \alpha_0/\beta_0 - 1$.

far removed from their FSY equivalents, with the only similarity being that both m and ω tend towards zero at large values of p_0 . We note, however, that the relationship in this region approximates $m = 3\omega$, as opposed to $m = \omega$ for the FSY states. Furthermore, despite the fermion mass and energy both vanishing at large p_0 , along with the ADM mass, the radius of the states somewhat confusingly tends towards a constant value of $\bar{R} \approx 0.697$. This suggests that perhaps there exists an analytic solution at infinite p_0 , towards which the solutions approach, although we have been unable to obtain it. At the other extreme, we observe that both m and ω become infinite as $p_0 \rightarrow 0$, whereas both the ADM mass and the radius tend towards zero. This is perhaps an even stranger scenario, since the gravitational mass of these states is almost non-existent, despite the presence of two ultra-high mass fermions. This is somewhat reminiscent of the mass-scale separation observed in the Einstein–Dirac–Higgs system, although here the degree of separation is significantly greater.

Examples of three individual states are shown in Fig. 6.3, where for each we plot the radial profiles of the fermion fields, metric fields and enclosed Komar mass. The locations of these states correspond to the points labelled in Fig. 6.3. Considering first the $p_0 = 0.05$ solution, we see that the fermion fields $\alpha(r)$ and $\beta(r)$ are almost equal throughout the fermion source, initially following an almost elliptical trajectory before abruptly transitioning to an exponential decay. The metric fields exhibit a similarly odd behaviour,

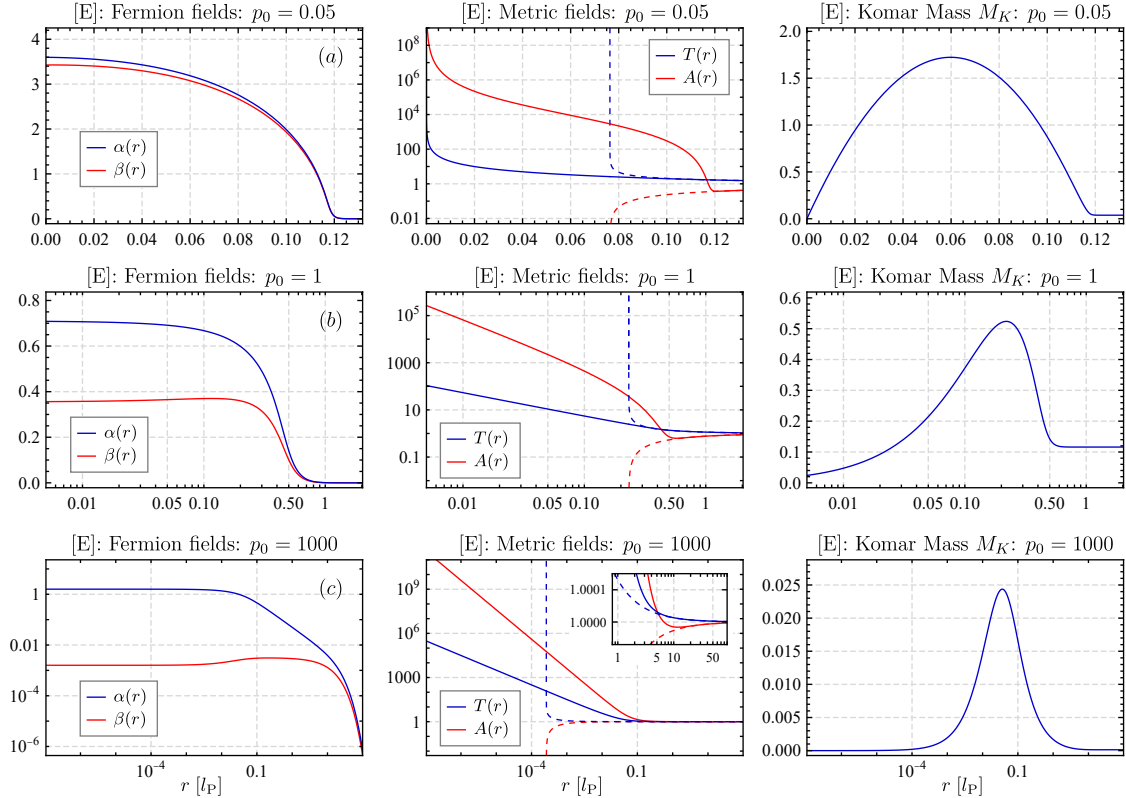


Figure 6.3: Examples of three individual states generated using the asymptotic expansion [E], showing for each the radial profiles of the fermion and metric fields, along with the enclosed Komar mass. The locations of these states are indicated on the curves shown in Fig. 6.2. Note that the parameter values for these solutions can be found in Appendix B.

with $A(r)$ and $T(r)$ gradually decreasing from their central infinite values, before crossing at large r , after which $A(r)$ starts to increase, following its asymptotic Schwarzschild form. This crossing ensures that the ADM mass of the state remains positive, although its actual value is infinitesimal, due to the approximately flat slope of the metric fields at large r . The enclosed Komar mass indicates even a stranger behaviour, since it no longer remains monotonic. This implies that an observer within the fermion source would experience the gravitational effects of a much more massive object than would an observer at infinity, i.e. much of the gravitational mass of the fermion source is hidden from external measurement. The overall qualitative behaviour does not change significantly as p_0 is increased, although we note that the total ADM mass is much larger for the intermediate state, but once again almost infinitesimal for the $p_0 = 1000$ state.

Unfortunately, we have no satisfactory explanation for the somewhat counter-intuitive properties exhibited by this family of states. Our only comment is to emphasise that these solutions are singular, and that the presence of a central singularity can seemingly result in a wide variety of unusual phenomena. This will become a common theme as we continue our review of the various asymptotic expansions.

6.1.6 Case [F]: $\{\alpha, \beta, A, T\} \sim \{r^{\kappa/2}, r^{-\kappa/2}, r^0, r^1\}$

In this case, the fermion field $\beta(r)$ diverges at small r , but all other fields are regular. The accompanying asymptotic expansion (valid for both positive and negative parity) is:

$$\alpha(r) = \alpha_0 r^{\frac{\kappa}{2}} + \mathcal{O}\left(r^{1-\frac{\kappa}{2}}\right); \quad \beta(r) = \beta_0 r^{-\frac{\kappa}{2}} + \mathcal{O}\left(r^{1+\frac{\kappa}{2}}\right); \quad (6.37)$$

$$A(r) = 1 + \mathcal{O}\left(r^{2-|\kappa|}\right); \quad T(r) = \frac{2}{\xi\kappa|\kappa|\alpha_0\beta_0} r + \mathcal{O}\left(r^{2-|\kappa|}\right), \quad (6.38)$$

where the value of κ is constrained to lie within $-1 < \kappa < 1$. The Ricci and Kretschmann scalars take the following forms:

$$R = -\frac{2m\beta_0}{|\kappa|\alpha_0\beta_0} r^{-\kappa-1} + \dots; \quad K = 24r^{-4} + \dots, \quad (6.39)$$

and hence indeed a central singularity is present. This is one of only three expansions for which we have been unable to construct localised solutions, although we have no formal proof that they do not exist. Given the unphysical constraint on κ , however, we shall not dwell further on this case.

6.1.7 Case [G]: $\{\alpha, \beta, A, T\} \sim \{r^0, r^{-1}, r^0, r^2\}$

Here, we once again have a divergent $\beta(r)$, while all other fields are regular, with the small- r expansion valid for positive parity taking the following form:

$$\alpha(r) = -\frac{2\sqrt{2\omega}}{m\sqrt{6\xi}} + \alpha_1 r + \dots; \quad \beta(r) = -\frac{2\sqrt{2\omega}}{m^2\sqrt{6\xi}} r^{-1} + \frac{2\sqrt{2\omega}}{\sqrt{6\xi}} r + \dots; \quad (6.40)$$

$$A(r) = 1 - 2m^2 r^2 + \dots; \quad T(r) = \frac{3m^3}{2\omega} r^2 + \frac{3\sqrt{6\xi}m^4\alpha_1}{\sqrt{2\omega^3}} r^3 + \dots, \quad (6.41)$$

where the derivation of this enforces $\kappa = 2$. Note that the corresponding negative parity expansion can be obtained by interchanging the power-law dependences of $\alpha(r)$ and $\beta(r)$. Considering the Ricci and Kretschmann scalars, we obtain the following simple expressions:

$$R = -4r^{-2} + \dots; \quad K = -16r^{-4} + \dots, \quad (6.42)$$

indicating that indeed a central singularity is present. The situation for this expansion is somewhat unusual, in that the leading-order coefficients are constrained to take definite values, but the next-to-leading order coefficient α_1 is not. Even stranger, it turns out that there is a further redundancy at fourth order, resulting in a second unconstrained coefficient, which we take to be T_4 . We shall not here detail the full expansion up to this order, but simply note that, contrary to appearances, there are indeed four free parameters in the system, and we therefore expect any localised states associated with this expansion to constitute a one-parameter family of solutions. Since there does not appear to be a

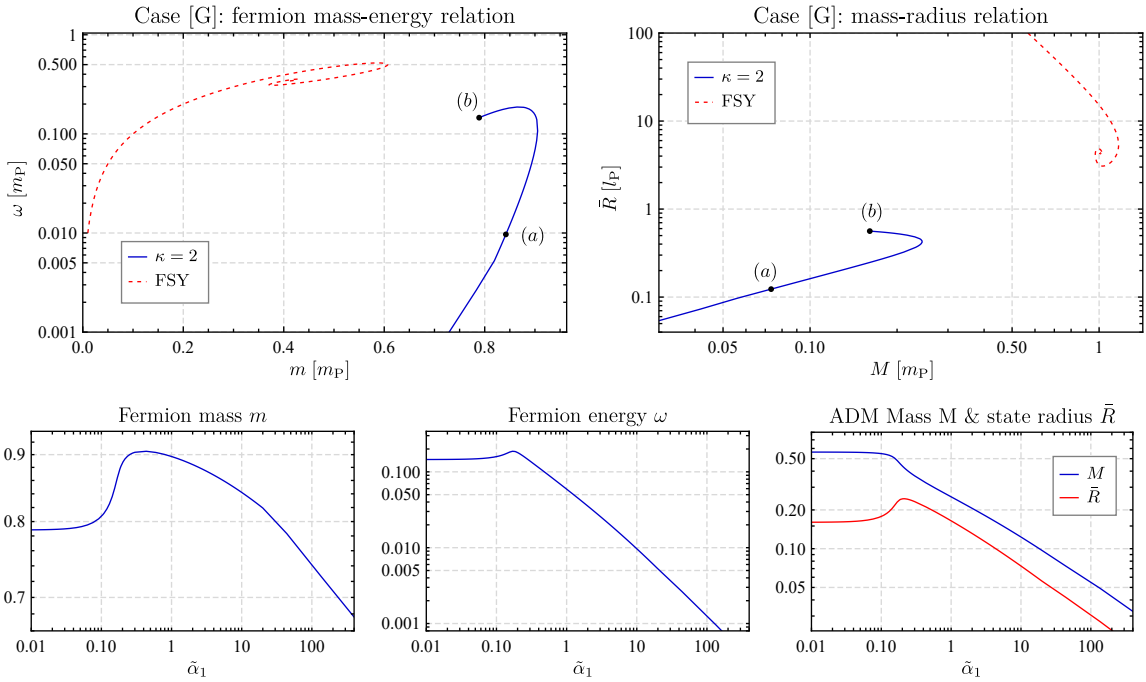


Figure 6.4: The one-parameter family of solutions with small- r behaviour governed by the asymptotic expansion [G]. Shown in the upper panels are the fermion mass-energy and mass-radius relations, alongside the corresponding curves in the FSY case. The lower panels indicate the behaviour of the fermion mass, fermion energy, ADM mass and average radius as a function of the unscaled parameter $\tilde{\alpha}_1$, which we use to parametrise the family of solutions.

straightforward quantity that remains invariant under the rescaling, we simply parametrise this family using the unscaled coefficient $\tilde{\alpha}_1$, while fixing the value of $\tilde{T}_4 = 1$.

There indeed exist localised solutions in this case, which we have successfully obtained numerically, with the behaviour of the one-parameter family of states being summarised in Fig. 6.4. Again we plot the fermion mass-energy and mass-radius relations, along with the values of various physical quantities as a function of $\tilde{\alpha}_1$. As in case [E], we find that the family of solutions is entirely disconnected from the FSY states, and does not exhibit a spiralling behaviour. Instead, at small values of $\tilde{\alpha}_1$, all physical properties of the solutions become constant, and thus the mass-energy and mass-radius curves terminate at a specific point in phase space. In contrast, at large $\tilde{\alpha}_1$, all quantities tend towards zero, and hence the solutions within this region represent highly-localised states, consisting of low-mass fermions confined around a central singularity.

The locations of two individual states are labelled on the mass-energy and mass-radius relations, with the radial profiles of these shown in Fig. 6.5. Consider first the solution with $\tilde{\alpha}_1 = 10$, which is located along the portion of the curves that tends towards zero. For this we plot the profiles of not only the fermion fields, metric fields and enclosed Komar mass, but also the optical geometry radial co-ordinate $\rho = rT$, the radial fermion energy density, and the Ricci scalar. For all quantities, there is a distinct change in behaviour at a radius

of $r \approx 0.15$, with the forms of the metric fields indicating that a horizon is very nearly formed at this point. Indeed, $A(r)$ drops to a very small but ultimately non-zero value, before latching on to its asymptotic Schwarzschild dependence. This radius is associated with a highly-pronounced bottleneck in the optical geometry, due to the maximum in $T(r)$, with the fermion energy density heavily peaked around the location of the accompanying stable photon sphere. Hence the fermion wavefunction is trapped within an exceedingly narrow radial region. Considering the profile of the Komar mass, we note a somewhat unusual behaviour, in that the enclosed gravitational mass is negative within the central regions of the state (in fact tending to negative infinity as $r \rightarrow 0$), only becoming positive at around the radius of the pseudo-horizon. Recall that the Komar mass at a radius r is calculated by considering the gravitational force on a test particle held stationary at that radius, and hence a negative value indicates that this force will be repulsive. The fermion source appears therefore to have two distinct regions — first, an extended inner core, in which the probability of finding a fermion is negligible and there exists effectively a bare spacetime singularity (with an overall repulsive effect), and second, a narrow region within which the fermion wavefunction is highly trapped, with the outer boundary of this

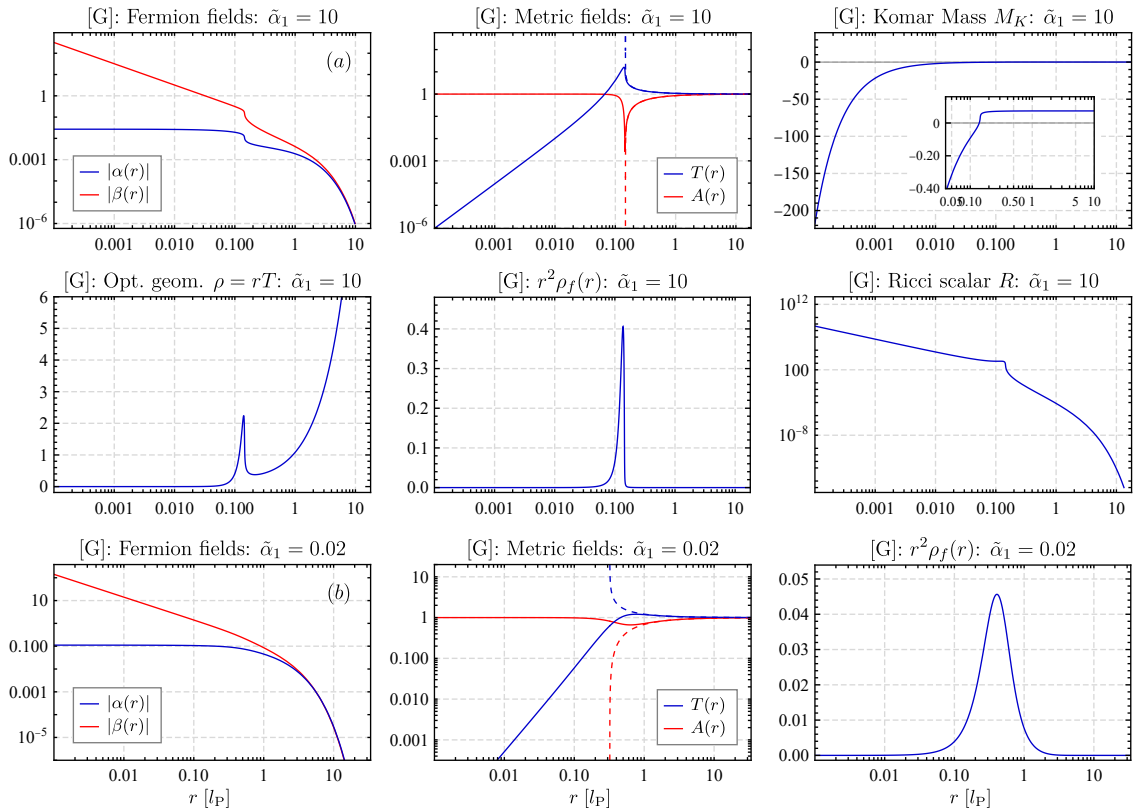


Figure 6.5: Examples of two singular solutions generated using the asymptotic expansions [G]. For both states, we plot the radial profiles of the fermion fields, metric fields and enclosed Komar mass, and in addition, for the $\tilde{\alpha}_1 = 10$ solution, the profiles of the radial co-ordinate in the optical geometry, $\rho = rT$, the radial fermion energy density $r^2\rho$, and the Ricci scalar R . Note that this last quantity diverges at $r = 0$, indicating the presence of a central singularity. Also note that the fermion fields are both negative for all values of r , and hence here we plot their absolute values.

corresponding to a pseudo-horizon. Another interesting feature to note from the Komar mass is that the transition from negative to positive implies that a radius exists at which the enclosed gravitational mass is measured to be zero, and thus an observer located at that point will experience no gravitational force whatsoever. We have verified that this is indeed the case using a geodesic analysis.

To summarise, although the overall behaviour of the $\tilde{\alpha}_1 = 10$ state is somewhat unusual, it is entirely self-consistent. The apparent strong repulsive effect of the central singularity pushes the central portion of the fermion wavefunction outwards in radius, resulting in the outer region of the fermion source becoming highly dense. This in turn creates a significant distortion in the metric, accompanied by the appearance of a stable photon sphere, around which the fermion density becomes trapped. Questions remain, however, such as why the central singularity is effectively repulsive, and indeed how we should interpret this type of solution. Finally, we note that the qualitative behaviour of the states does not change significantly as $\tilde{\alpha}_1$ is decreased, as indicated by the profiles of the $\tilde{\alpha}_1 = 0.02$ solution also included in Fig. 6.5. Rather, we find that the pseudo-horizon gradually disappears, and the fermion energy density becomes trapped within an increasingly broad radial region. As a final remark, we note that, as mentioned when considering the family of states, the solutions at small $\tilde{\alpha}_1$ approach a constant overall form, similar in structure to the $\tilde{\alpha}_1 = 0.02$ state shown, although we can identify no features suggesting why this should be the case.

6.1.8 Case [H]: $\{\alpha, \beta, A, T\} \sim \{r^{-\kappa/2+1}, r^{-\kappa/2}, r^0, r^\kappa\}$

At first glance, this case appears to be the many-fermion generalisation of case [G], since the power-law dependences coincide when $\kappa = 2$. Although this is indeed correct, and in fact this case is valid only when $\kappa \neq 2$, the behaviour of the expansion here is significantly different, and we are consequently unable to construct localised solutions. At small r , we have the following leading-order expansion, valid for positive parity:

$$\alpha(r) = \alpha_0 r^{-\frac{\kappa}{2}+1} + \mathcal{O}\left(r^{-\kappa/2+3}\right); \quad \beta(r) = \frac{\kappa-1}{m} \alpha_0 r^{-\frac{\kappa}{2}} + \mathcal{O}\left(r^{-\kappa/2+2}\right); \quad (6.43)$$

$$A(r) = 1 + \mathcal{O}\left(r^\kappa\right); \quad T(r) = \frac{2m\kappa}{\xi\kappa(\kappa-1)\alpha_0^2} r^\kappa + \mathcal{O}\left(r^{\kappa+2}\right). \quad (6.44)$$

The Ricci and Kretschmann scalars for this case are:

$$R = 2\kappa(\kappa-1)r^{-2} + \dots; \quad K = 4\kappa^2(3+2\kappa+\kappa^2)r^{-4} + \dots, \quad (6.45)$$

and thus there is a central singularity. Although it might appear that there is a single unconstrained parameter, α_0 , in the expansion above, this in fact becomes constrained due to a redundancy at higher order, and hence the system contains only two free parameters. This is incompatible with the construction of localised solutions, since the boundary

conditions in principle remove three degrees of freedom, although we note that this did not prevent the generation of the infinite-redshift states in case [C], for reasons that are unclear. Such an exception does not appear to apply here, however, and our attempts at obtaining localised numerical solutions have been unsuccessful.

6.1.9 Case [I]: $\{\alpha, \beta, A, T\} \sim \{r^0, r^0, r^{-1}, r^0\}$

Here we now return to a situation in which the leading-order powers are independent of the fermion number, with now all but the metric field $A(r)$ leading like a constant. The small- r expansion for this case (valid for both positive and negative parity) is:

$$\alpha(r) = \alpha_0 + \frac{\kappa\alpha_0}{\sqrt{-\xi\kappa^3\alpha_0\beta_0T_0}} r^{1/2} + \dots; \quad (6.46)$$

$$\beta(r) = \beta_0 - \frac{\kappa\beta_0}{\sqrt{-\xi\kappa^3\alpha_0\beta_0T_0}} r^{1/2} + \dots; \quad (6.47)$$

$$A(r) = -\xi\kappa|\kappa|\alpha_0\beta_0T_0r^{-1} + 1 - \xi|\kappa|\omega T_0^2(\alpha_0^2 + \beta_0^2) + \dots; \quad (6.48)$$

$$T(r) = T_0 + \frac{T_0}{3\kappa\alpha_0\beta_0} (2\omega T_0[\alpha_0^2 + \beta_0^2] - m[\alpha_0^2 - \beta_0^2]) r + \dots, \quad (6.49)$$

where we note that, for positive κ , one of either α_0 or β_0 is forced to be negative, and that there are no constraints on κ that arise when deriving these expressions. The Ricci and Kretschmann scalars here take the following forms:

$$R = \xi\kappa m T_0(\alpha_0^2 - \beta_0^2)r^{-2} + \dots; \quad K = 6\xi^2\kappa^3(3\kappa - 2|\kappa|)\alpha_0^2\beta_0^2T_0^2r^{-6} + \dots, \quad (6.50)$$

implying that there is indeed a central singularity. The expansion above contains a total of three unconstrained coefficients, α_0 , β_0 and T_0 , none of which are subject to any further constraints at higher order (as far as we can tell). Any localised states associated with this expansion should therefore constitute a 1-dimensional set of 1-parameter solutions, i.e. each family of states is itself defined by the value of a second parameter. For numerical convenience, we specify the two free parameters in the system as $p_0 = \alpha_0/\beta_0$ and $p_1 = \alpha_0\sqrt{mT_0}$, which are both invariant under the rescaling. We shall take p_1 to be the parameter that identify the respective families, and p_0 as the quantity that distinguishes the states within them. In addition, in order to obtain positive mass states, we are forced to consider the negative parity asymptotic expansion.

We have been successful in generating localised solutions using this expansion for a variety of κ values, with the overall behaviour for $\kappa = -2$ illustrated in Fig. 6.6. Here, we plot the mass-energy and mass-radius relations, along with quantities as a function of p_0 , for three families of states with p_1 values of 0.4, 0.5 and 0.6. We note that there appears only to be a narrow range of parameter values for which solutions can be obtained. Considering the mass-energy curves, we see that once again these singular solutions are

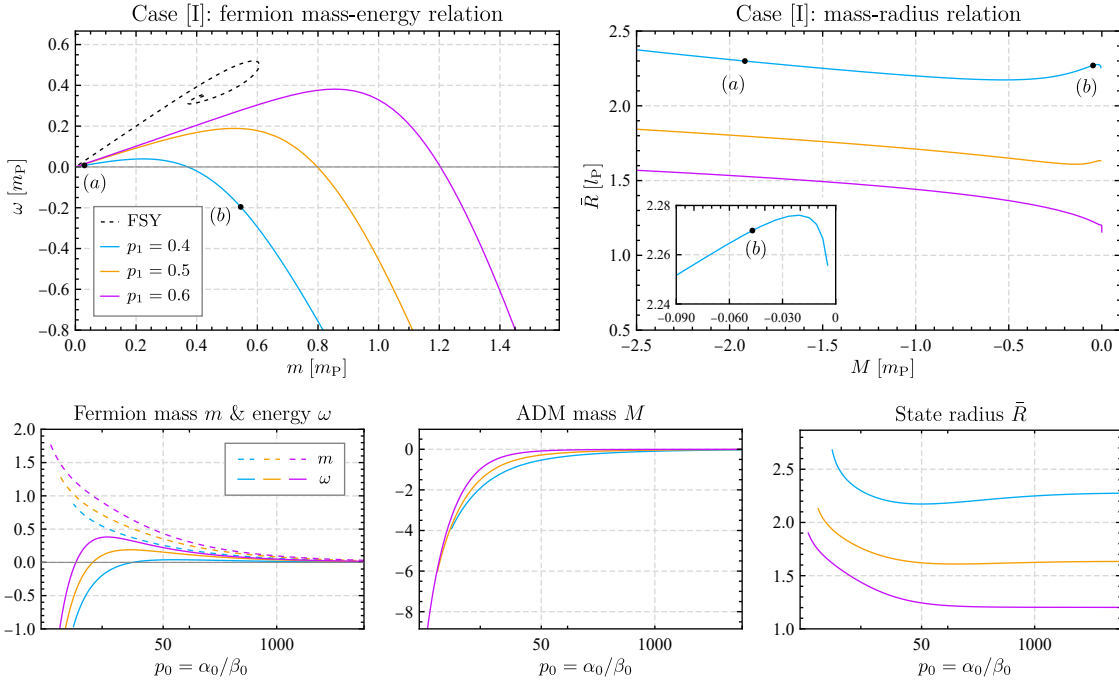


Figure 6.6: Families of states generated using the small- r asymptotic expansion [I]. There are now multiple families of solutions, distinguished by their respective values of the parameter $p_1 = \alpha_0 \sqrt{mT_0}$, with the three cases shown corresponding to $p_1 = 0.4, 0.5$ and 0.6 , all with $\kappa = -2$. In the upper panels, we plot the fermion mass-energy and mass-radius relations, and in the lower panels, m , ω , M and \bar{R} as a function of the second parameter $p_0 = \alpha_0/\beta_0$. Note that all states have a negative ADM mass, and in addition there exist those in which the fermion energy is also negative.

unrelated to the FSY states, although indeed m and ω both tend to zero as $p_0 \rightarrow \infty$. Interestingly, at small values of p_0 , the value of the fermion energy becomes negative (one could argue that this is somewhat unphysical, however), and the curves appear to asymptote towards finite values of p_0 . The most intriguing feature of these states, however, is that their total ADM masses are negative, irrespective of the value of p_0 . This implies that they are gravitationally repulsive, despite containing positive mass fermions with positive energy (at least in the majority of cases). Again, the only explanation for this appears to be the presence of the central singularity, although the physical mechanism through which this permits negative gravitational mass is unclear.

Examples of two individual states are shown in Fig. 6.7, both corresponding to p_1 values of 0.4 , but with different values of p_0 . Their locations are indicated on the $p_1 = 0.4$ curves in the previous figure. There are a few interesting features to note from these plots. First, unlike in some of the previous cases, the metric fields here do not cross, and hence they approach their asymptotic forms from the ‘wrong’ direction. This explains why the states exhibit a negative gravitational mass. What is perhaps surprising, however, is that the Komar mass in fact decreases towards this value, and is precisely zero at the centre of the fermion source. This implies that, although the overall gravitational effect of these

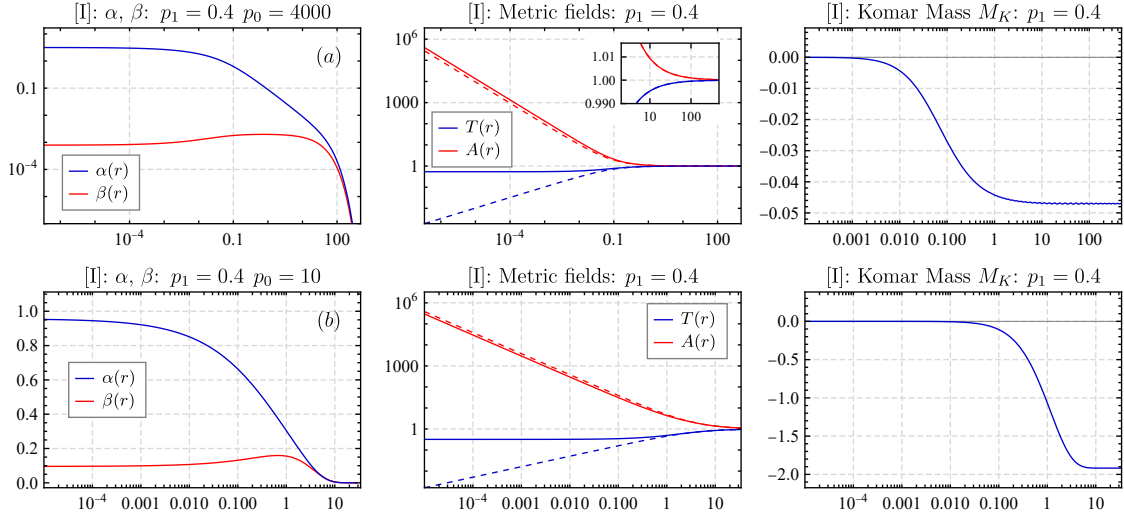


Figure 6.7: Examples of two individual solutions constructed using the small- r expansion [I]. These have a common value of $p_1 = 0.4$, but differing values of p_0 , with their locations along the family of $p_1 = 0.4$ states indicated in Fig. 6.6. Note that the ADM mass of the states here is negative, since the metric fields approach their asymptotic forms from the opposite direction to that usually encountered.

states is repulsive, the strength of this decreases as one moves further within the fermion source. This is precisely the situation one would expect for an object consisting of negatively gravitating particles, but as we emphasise, there is no such exotic matter present here.

6.1.10 Case [J]: $\{\alpha, \beta, A, T\} \sim \{r^{3/2}, r^0, r^{-1}, r^{1/2}\}$

We now encounter another case in which the leading-order powers do not depend on the fermion number, but here only $\beta(r)$ leads like a constant. The small- r expansion takes the following form (for positive parity):

$$\alpha(r) = -\frac{2m\beta_0}{3\sqrt{A_0}}r^{\frac{3}{2}} + \frac{(2m\kappa - 3\omega\sqrt{A_0}T_0)\beta_0}{6\sqrt{A_0}}r^2 + \dots; \quad (6.51)$$

$$\beta(r) = \beta_0 - \frac{\kappa\beta_0}{\sqrt{A_0}}r^{\frac{1}{2}} + \dots; \quad (6.52)$$

$$A(r) = A_0r^{-1} + 1 + \dots; \quad (6.53)$$

$$T(r) = T_0r^{1/2} - \frac{T_0}{2A_0}r^{3/2} + \dots, \quad (6.54)$$

where we note that, for positive mass solutions, α_0 and β_0 must have the opposite sign, and that there are no constraints on the value of κ . The expansions for the Ricci and Kretschmann scalars for this case are:

$$R = -\xi m\beta_0^2 T_0 r^{-3/2} + \dots; \quad K = 12A_0^2 r^{-6} + \dots, \quad (6.55)$$

indicating the presence of a central singularity. There are here three unconstrained parameters in the expansion (β_0 , A_0 and T_0), and hence any localised states should constitute a one-dimensional set of one-parameter families of solutions, similar to the situation observed in case [I]. Again, we choose to use parameters that are invariant under the rescaling, these being $p_0 = mA_0$ and $p_1 = m_0 T_0 \beta_0^2$, with the former parametrising the families and the latter distinguishing between them.

We find that indeed localised solutions can be generated for this case, although positive mass states are required to exhibit at least one fermion node. The behaviour of the families of states (for $\kappa = 2$) is illustrated in Fig. 6.8, where we plot the fermion mass-energy and mass-radius relations for the families with p_1 values of 0.01 and 0.1, along with various quantities as a function of p_0 . We observe that the singular solutions here are not disconnected from the FSY states, instead approaching them in the limit $p_0 \rightarrow 0$, and indeed both the mass-energy and mass-radius curves closely follow their FSY equivalents. This behaviour only continues up until a particular value of p_0 , however, after which the curves separate from the FSY case and diverge to infinity. Note that the $p_1 = 0.1$ family follows the FSY curve for a slightly longer period than the $p_1 = 0.01$ family, and we therefore surmise that the point at which the curves detach from the FSY states becomes progressively closer to the centre of the FSY spiral as p_1 increases. It is difficult to verify this, however, since states with large values of p_1 are not straightforward to gen-

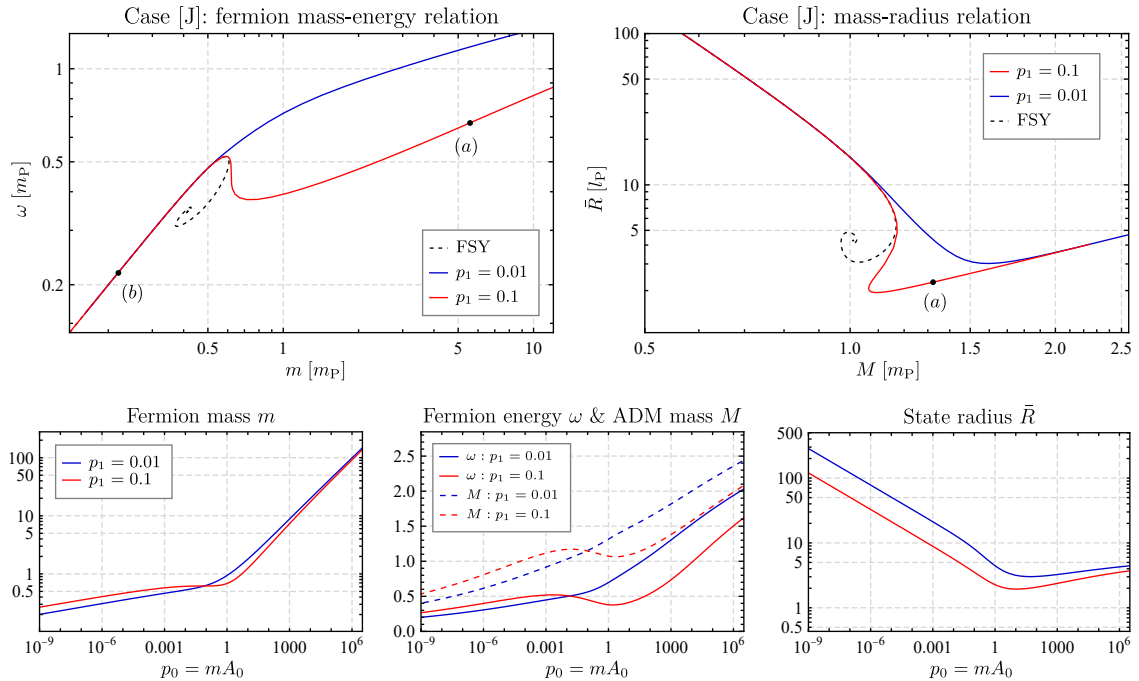


Figure 6.8: Plots summarising the behaviour of the families of singular states generated using the asymptotic expansion [J], with $\kappa = 2$. Shown are the fermion mass-energy and mass-radius relations for two families of states, with parameter values $p_1 = 0.01$ and $p_1 = 0.1$, respectively, along with plots of the quantities m , ω , M and \bar{R} as a function of p_0 . Note that the curves follow closely the low-redshift portion of the FSY spiral, before diverging at large values of p_0 .

erate numerically, due to the appearance of multivaluedness. This behaviour is certainly surprising, even more so when one considers the fact that the families of singular states contain a single fermion node, while the FSY states shown do not.

To understand this behaviour, we must consider the profiles of individual solutions, two of which are presented in Fig. 6.9, corresponding to states (a) and (b) marked on the previous figure. Considering first the $p_0 = 500$ state, which is located beyond the point at which the curve detaches from the FSY spiral, we note that the metric fields again cross at a reasonably large radius, and thus the total ADM mass of the state is positive. The profile of the Komar mass indicates, however, that the enclosed gravitational mass become negative within the inner regions of the fermion source, and consequently there exists a radius at which an observer will experience no gravitational force whatsoever. Turning to the $p_0 = 0.0001$ state, we observe from the fermion field profiles that the solution has effectively split into two distinct regions: a core, in which $\beta(r)$ is positive and $\alpha(r)$ is negligible, and an outer region, located beyond the single fermion node. Somewhat surprisingly, the profile of the fermion fields (and indeed the metric fields) within the outer zone is almost precisely that of an $n = 0$ FSY state, explaining why the family of solutions approaches so closely the FSY curve for small values of p_0 . Note that this argument relies on the fact that the inner regions of the fermion source contain a negligible amount of fermion density, and thus the properties of the solution are primarily determined at large r . Unfortunately, we have no explanation for this overall behaviour, other than the plots presented.

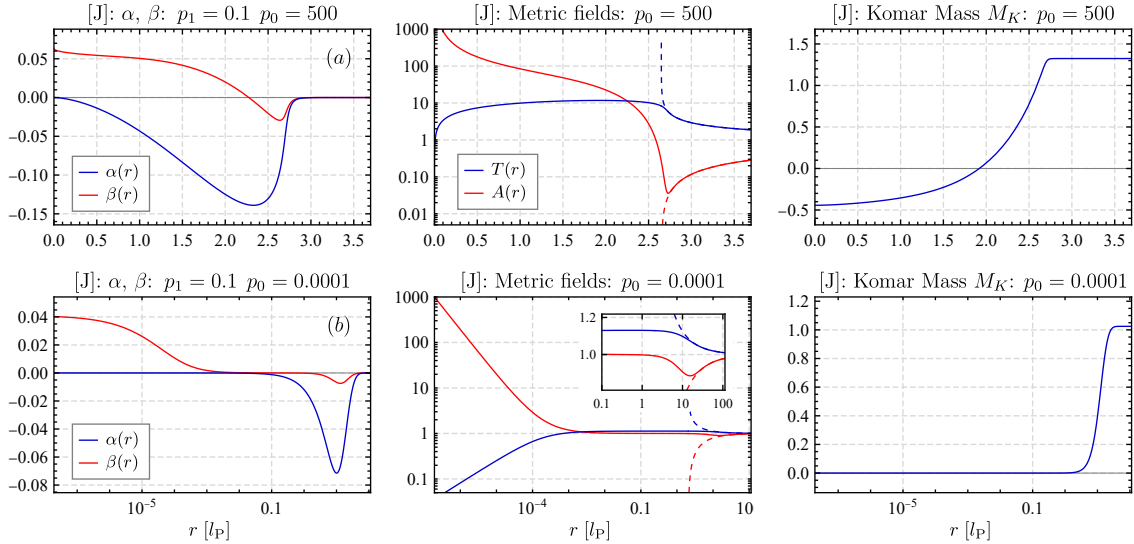


Figure 6.9: Examples of two singular states constructed using the asymptotic expansion [J], these both having the same value of the parameter p_1 , but differing in their values of p_0 . Shown for each state are the radial profiles of the fermion fields, metric fields and Komar mass. The locations of these solutions along the family of $p_1 = 0.1$ states is indicated in Fig. 6.8.

6.1.11 Case [K]: $\{\alpha, \beta, A, T\} \sim \{r^0, r^{3/2}, r^{-1}, r^{1/2}\}$

This case appears to be a variation on case [J], with simply the leading-order dependences of the two fermion fields switched. The expansion is therefore similar to that above:

$$\alpha(r) = \alpha_0 - \frac{\kappa\alpha_0}{\sqrt{A_0}}r^{\frac{1}{2}} + \dots; \quad (6.56)$$

$$\beta(r) = -\frac{2m\alpha_0}{3\sqrt{A_0}}r^{\frac{3}{2}} - \frac{(2m\kappa - 3\omega\sqrt{A_0}T_0)\alpha_0}{6\sqrt{A_0}}r^2 + \dots; \quad (6.57)$$

$$A(r) = A_0r^{-1} + 1 + \dots; \quad (6.58)$$

$$T(r) = T_0r^{1/2} - \frac{T_0}{2A_0}r^{3/2} + \dots \quad (6.59)$$

The corresponding expansions for the Ricci and Kretschmann scalars are:

$$R = \xi m\alpha_0^2 T_0 r^{-3/2} + \dots; \quad K = 12A_0^2 r^{-6} + \dots, \quad (6.60)$$

and indeed there is a central singularity. We find that the states associated with this expansion behave in an identical manner to those of case [J], as might be expected, and thus we shall not dwell further on this case.

6.1.12 Case [L]: $\{\alpha, \beta, A, T\} \sim \{r^0, r^0, r^{-1}, r^{1/2}\}$

We finally arrive at the twelfth and final small- r expansion, for which the fields again have no leading-order dependence on κ , with now both $\alpha(r)$ and $\beta(r)$ leading like constants. The small- r expansion in this case is (valid for positive parity):

$$\alpha(r) = \alpha_0 + \frac{\kappa\alpha_0}{\sqrt{A_0}}r^{1/2} + \dots; \quad \beta(r) = \beta_0 - \frac{\kappa\beta_0}{\sqrt{A_0}}r^{1/2} + \dots; \quad (6.61)$$

$$A(r) = A_0r^{-1} + 1 + \dots; \quad T(r) = T_0r^{1/2} + \xi\kappa^2 \frac{T_0^2\alpha_0\beta_0}{A_0}r + \dots, \quad (6.62)$$

which is found to exist for any value of κ . The corresponding expansions for the Ricci and Kretschmann scalars are:

$$R = \xi\kappa m T_0 (\alpha_0^2 - \beta_0^2) r^{-3/2} + \dots; \quad K = 12A_0^2 r^{-6} + \dots, \quad (6.63)$$

and consequently a central singularity is present. Note that here we have a total of four unconstrained parameters in the expansion (α_0 , β_0 , A_0 and T_0), and as such any resulting localised solutions should constitute a 2-dimensional set of 1-parameter families of solutions. It is again convenient to construct parameters that are invariant under the rescaling procedure, and to this end we therefore choose to identify families by their values of $p_1 = mA_0$ and $p_2 = \beta_0^4 T_0^2 A_0^{-1}$, and distinguish the solutions contained therein by the parameter $p_0 = \alpha_0/\beta_0$.

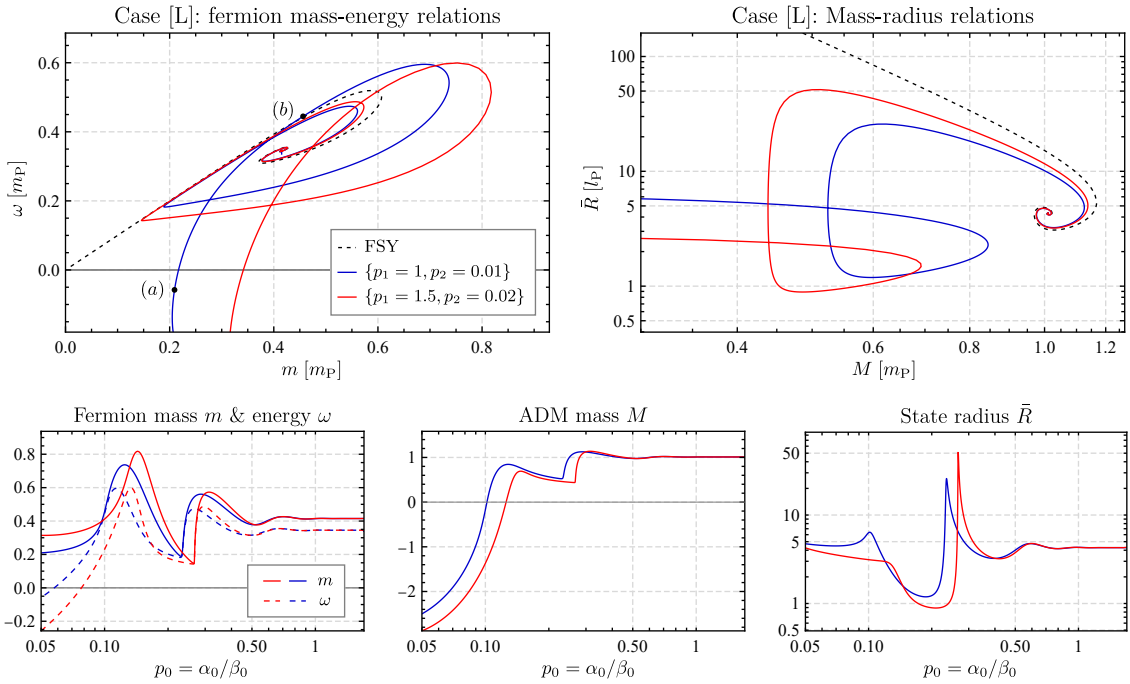


Figure 6.10: Plots indicating the behaviour of the families of singular states generated using the small- r expansion [L]. Shown in the upper panels are the fermion mass-energy and mass-radius relations for two families of states with the values of p_1 and p_2 indicated. These are observed to spiral towards the same infinite-redshift state as the FSY family. The lower panels plot the values of m , ω , M and \bar{R} , as a function of the parameter $p_0 = \alpha_0/\beta_0$, for the same two families.

Given the high dimensionality of the parameter space, we have been unable to perform a thorough analysis of this class of solutions, and therefore simply present here a sample of the observed behaviour. This is illustrated in Fig. 6.10, where we plot the mass-energy and mass-radius relations for examples of two families, along with the behaviour of physical quantities as a function of p_0 . Again, we see that the singular states closely follow the FSY curve, this time spiralling towards the infinite-redshift state at large values of p_0 , but detaching from the spiral at low redshift. Thus the infinite-redshift solution appears to be a limiting case not only for the FSY states but also the states contained in class [L]. On the other hand, for small values of p_0 , we find that the families terminate at seemingly arbitrary states, which exhibit both negative fermion energy and negative ADM mass.

We can again partially explain this behaviour by analysing the structure of individual states, two of which are shown in Fig. 6.11, these corresponding to solutions located within the $\{p_1 = 1, p_2 = 0.01\}$ family shown in the previous figure. Solution (a) lies along the portion of the curve that has detached from the FSY spiral, and we observe a fairly straightforward structure, noting that the metric fields do not cross at any radius, and hence the total ADM mass of the state is negative. Interestingly, the enclosed gravitational mass is observed to decrease as one moves towards $r = 0$. In contrast, solution (b), which is located towards the centre of the FSY spiral, exhibits a more complex behaviour, and we see once again a separation into two distinct regions. These appear analogous to the

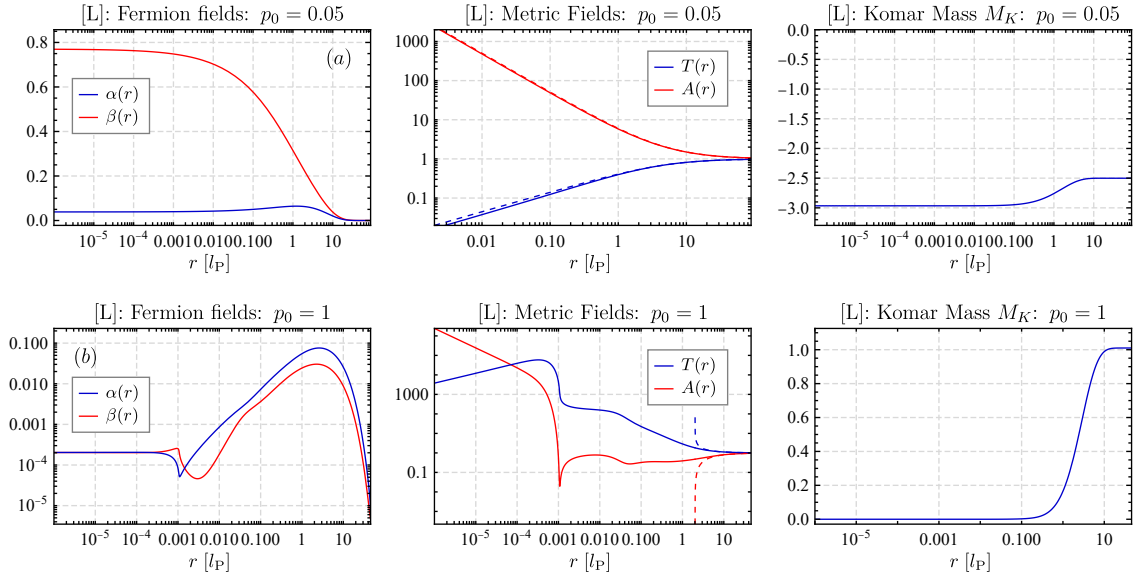


Figure 6.11: Examples of two singular solutions generated using the asymptotic expansion [L], both with parameter values $p_1 = 1$ and $p_2 = 0.01$, but differing values of p_0 . Shown for each state are the radial profiles of the fermion fields, metric fields and Komar mass. The locations of these two states are indicated on the mass-energy curve in Fig. 6.10.

core and power-law zones of the FSY states, and indeed we observe a similar oscillatory behaviour around the pure power-law solution to that which occurs in the FSY case. We therefore conclude that these singular solutions differ only from the FSY states with respect to the small- r expansion that the fields approximate within the core zone. This explains why the two cases exhibit similar properties, with the appearance of a common spiralling behaviour related to the fermion self-trapping effect that occurs within the power-law zone. Again, however, a physical interpretation for this similarity, or indeed these families of singular states, is lacking.

This concludes our somewhat lengthy summary of the various singular solutions obtained by varying the small- r expansions for the equations of motion. In total, we have identified 12 possible expansions, with 9 of these resulting in localised states, 8 of which are singular. Overall, therefore, we conclude that the FSY states are the only non-singular localised states in the Einstein–Dirac system.

6.2 $\omega = 0$, $A = 1$ analytic solution

Before concluding this chapter, we first discuss a new analytic solution to the Einstein–Dirac equations that does not appear to have been presented elsewhere. Although this solution contains a central singularity, it is entirely normalisable, and hence to our knowledge the resulting states represent the first analytic, localised solutions to the Einstein–Dirac system. The solution is valid for the somewhat unusual specific case of zero fermion

energy ($\omega = 0$) and entirely flat space ($A(r) = 1$), noting that this latter condition does not of course imply flat *space-time*. We shall first outline the derivation of the solution, for both the two-fermion system and its many-fermion generalisation, before discussing its properties and interpretation.

We begin by setting $\omega = 0$ and $A(r) = 1$ in the FSY equations (6.1)–(6.4), reducing them to the following set of three differential equations:

$$\alpha' = \frac{N_f}{2r}\alpha - m\beta; \quad (6.64)$$

$$\beta' = -m\alpha - \frac{N_f}{2r}\beta; \quad (6.65)$$

$$-2r\frac{T'}{T} = N_f\xi T (\alpha\beta' - \alpha'\beta), \quad (6.66)$$

where $\xi = 8\pi G$ and we have assumed positive parity. In addition, the normalisation condition (6.5) simplifies to:

$$4\pi \int_0^\infty T (\alpha^2 + \beta^2) dr = 1. \quad (6.67)$$

Differentiating (6.64) and eliminating $\beta(r)$ using (6.65), we can obtain the following differential equation for $\alpha(r)$:

$$\alpha'' = \left(m^2 - \frac{N_f}{2r^2} + \frac{N_f^2}{4r^2} \right) \alpha. \quad (6.68)$$

Note that this simplifies considerably for the case of $N_f = 2$, in which the final two terms cancel, and we shall therefore consider this case separately, before tackling the many-fermion generalisation.

6.2.1 The two-fermion solution

For $N_f = 2$, the equation for α simplifies to $\alpha'' = m^2\alpha$, which can be easily solved to give:

$$\alpha(r) = c_1 e^{-mr} + c_2 e^{mr}. \quad (6.69)$$

where c_1 and c_2 are arbitrary constants. Since our aim is to construct normalisable states, we eliminate the growing mode by setting $c_2 = 0$. Then, substituting this back into (6.64), we can solve for $\beta(r)$:

$$\beta(r) = \frac{1}{m} \left(\frac{1}{r}\alpha - \alpha' \right) = c_1 \left(\frac{1}{mr} + 1 \right) e^{-mr}. \quad (6.70)$$

Using these expressions, the remaining Einstein equation (6.66) reduces to the following differential equation for $T(r)$:

$$\begin{aligned} \frac{T'}{T^2} &= -\frac{\xi}{r} \left(c_1^2 e^{-2mr} \left[-m - \frac{1}{r} - \frac{1}{mr^2} \right] - c_1^2 e^{-2mr} \left(1 + \frac{1}{mr} \right) (-m) \right) \\ &= \frac{c_1^2 \xi}{mr^3} e^{-2mr}, \end{aligned} \quad (6.71)$$

which can be solved by separation of variables:

$$\int_r^\infty \frac{dT}{T^2} = c_1^2 \xi \int_r^\infty \frac{1}{ms^3} e^{-2ms} ds. \quad (6.72)$$

Note that we are integrating radially inwards from infinity, in order to avoid the singularity at $r = 0$. The left-hand side of this can be straightforwardly evaluated, noting that normalisable states must be asymptotically flat, i.e. $T(\infty) = 1$. We can solve the integral on the right-hand side using integration by parts, with a change of variable to $u = 2ms$ proving useful. Our equation therefore becomes:

$$\begin{aligned} -1 - \frac{1}{T(r)} &= 4mc_1^2 \xi \int_{2mr}^\infty \frac{1}{u^3} e^{-u} du \\ &= 4mc_1^2 \xi \left(\frac{1}{8m^2 r^2} e^{-2mr} - \frac{1}{2} \int_{2mr}^\infty \frac{1}{u^2} e^{-u} du \right) \\ &= 4mc_1^2 \xi \left(\frac{1}{8m^2 r^2} e^{-2mr} - \frac{1}{4mr} e^{-2mr} + \frac{1}{2} \int_{2mr}^\infty \frac{1}{u} e^{-u} du \right) \\ &= c_1^2 \xi \left(\left[\frac{1 - 2mr}{2mr^2} \right] e^{-2mr} - 2m \text{Ei}(-2mr) \right), \end{aligned} \quad (6.73)$$

where $\text{Ei}(x)$ is the exponential integral, defined as:

$$\text{Ei}(x) = - \int_{-x}^\infty \frac{1}{t} e^{-t} dt. \quad (6.74)$$

Note that this tends to negative infinity as $x \rightarrow 0$ and zero as $x \rightarrow -\infty$. Thus we arrive at the following final expressions for $\alpha(r)$, $\beta(r)$ and $T(r)$, valid in the two-fermion system:

$$\alpha(r) = c_1 e^{-mr}; \quad (6.75)$$

$$\beta(r) = c_1 \left(\frac{1}{mr} + 1 \right) e^{-mr}; \quad (6.76)$$

$$T(r) = \frac{2mr^2}{2mr^2 + c_1^2 \xi [(1 - 2mr)e^{-2mr} - 4m^2 r^2 \text{Ei}(-2mr)]}. \quad (6.77)$$

Whether this solution is normalisable depends on the behaviour of the normalisation integrand in (6.67). Certainly no issues arise at large r , since $\alpha(r)$ and $\beta(r)$ both decay exponentially, while $T(r)$ asymptotes towards 1, but we must also consider the possibility of a singularity at $r = 0$. To analyse this, we must determine the behaviour of $T(r)$ at

small r , for which we can utilise the following series expansion of the exponential integral:

$$\text{Ei}(-x) = \log(x) + \gamma_E - x + \frac{x^2}{4} + \dots, \quad (6.78)$$

where γ_E is Euler's constant. This allows us to perform a small- r expansion of (6.77), for which we obtain:

$$\begin{aligned} T(r) &= \frac{2mr^2}{2mr^2 + c_1^2 \xi [(1 - 2mr)(1 - 2mr + \dots) - 4m^2 r^2 (\log(2mr) + \gamma_E + \dots)]} \\ &= \frac{2mr^2}{c_1^2 \xi} (1 + 4mr + 4m^2 r^2 \log(2mr) + \dots). \end{aligned} \quad (6.79)$$

We can now consider the corresponding small- r behaviour of the normalisation integrand, this evaluating to:

$$\begin{aligned} T(\alpha^2 + \beta^2) &= \frac{2mr^2}{c_1^2 \xi} (1 + 4mr + 4m^2 r^2 \log(2mr)) c_1^2 (1 - 2mr) \left(\frac{1}{m^2 r^2} \right) + \dots \\ &= \frac{2}{m\xi} (1 + 4mr + 4m^2 r^2 \log(2mr) + \dots). \end{aligned} \quad (6.80)$$

This is entirely well-behaved at $r = 0$ and hence we conclude that the solution presented here is indeed normalisable. We can therefore rewrite the normalisation condition (6.67) by substituting in the forms for $\alpha(r)$, $\beta(r)$ and $T(r)$, which gives:

$$1 = \frac{8\pi c_1^2}{m} \int_0^\infty \frac{(2m^2 r^2 + 2mr + 1)e^{-2mr}}{2mr^2 + c_1^2 \xi [(1 - 2mr)e^{-2mr} - 4m^2 r^2 \text{Ei}(-2mr)]} dr. \quad (6.81)$$

This represents an implicit equation that can be solved to determine the value of c_1 for which the solution is correctly normalised. This therefore leaves only a single free parameter in the system, the fermion mass m , which our numerical analysis indicates can take any positive value.

As mentioned, we find that this solution is singular at $r = 0$, which can be confirmed by considering the form of the Ricci scalar at small r . Using the expression for this given in (2.51), we obtain the following expansion:

$$\begin{aligned} R &= \frac{2\xi m}{r^2} T(\alpha^2 - \beta^2) \\ &= \frac{2\xi m}{r^2} \frac{2mr^2}{c_1^2 \xi} (1 + 4mr + \dots) c_1^2 (-1 - 2mr) \frac{1}{m^2 r^2} + \dots \\ &= -\frac{4}{r^2} - \frac{16m}{r} - 16m^2 \log(2mr) + \dots \end{aligned} \quad (6.82)$$

This diverges at $r = 0$, and hence the spacetime contains a central singularity. Another quantity to consider is the total ADM mass of the solution. Since the metric is not asymptotically Schwarzschild, however, we cannot extract this information from the large r form, and instead must consider the enclosed Komar mass (2.47), which can be simply

read off the Einstein equation:

$$M_K(r) = -r^2 \frac{T'}{T^2} = -\frac{c_1^2 \xi}{mr} e^{-2mr}. \quad (6.83)$$

Note that this diverges to negative infinity at $r = 0$ and decays exponentially at large r . We therefore conclude that the total ADM mass of the state is zero, but at a finite radius the enclosed gravitational mass is negative. Hence, from a purely gravitational point of view, the fermion source is entirely invisible to an observer at $r = \infty$, but produces a repulsive effect on test particles at finite r , with this increasing in strength as $r \rightarrow 0$.

The overall behaviour of the solution is illustrated in Fig. 6.12, for states with fermion mass values of $m = 0.1$, 1 and 10. The most notable feature here is that the metric field




Figure 6.12: Plots summarising the behaviour of the localised analytic solution to the two-fermion Einstein–Dirac system, which is valid when $\omega = 0$ and $A(r) = 1$. Shown are the radial profiles of the fermion fields, metric fields, radial fermion number density and enclosed Komar mass, for states with fermion masses of **(left)** $m = 0.1$, **(middle)** $m = 1$ and **(right)** $m = 10$. Note that in all three cases the Komar mass tends to zero as $r \rightarrow \infty$.

$T(r)$ becomes progressively steeper as m increases, with this being accompanied by a peak developing in the radial fermion number density at a non-zero radius. Hence the fermion source becomes confined to an increasingly narrow region as m becomes larger. We note that the radius of the state (measured by the quantity \bar{R}), however, does not exhibit a monotonic behaviour in m , instead initially rising at small m , with $R \sim m^{1/2}$, before reaching a maximum radius at $m \approx 1.5$, after which the radius subsequently decreases as $R \sim m^{-1/4}$. Hence states comprising fermions with masses of $m = 0$ and $m = \infty$ both have zero radial extent, somewhat surprisingly. We shall discuss further the various properties of this solution in section 6.2.3.

6.2.2 The many-fermion solution

We now consider the many-fermion system, in which an equivalent analytic solution can be found, exhibiting many of the same properties as that in the two-fermion case. The derivation is slightly more involved, however, since there is no longer the same cancellation between terms in (6.68), with the differential equation for $\alpha(r)$ now taking the form:

$$\alpha'' = \left(m^2 - \frac{N_f}{2r^2} + \frac{N_f^2}{4r^2} \right) \alpha. \quad (6.84)$$

In order to solve this, we first change variables to $\gamma = \alpha/\sqrt{r}$, allowing us to rewrite the equation as:

$$\begin{aligned} \frac{1}{\sqrt{r}}\gamma' + \sqrt{r}\gamma'' - \frac{1}{4r\sqrt{r}}\gamma &= \left(m^2 - \frac{N_f}{2r^2} + \frac{N_f^2}{4r^2} \right) \sqrt{r}\gamma \\ r^2\gamma'' + r\gamma' &= \left(m^2r^2 + \left[\frac{N_f - 1}{2} \right]^2 \right) \gamma. \end{aligned} \quad (6.85)$$

The solutions to this are the modified Bessel functions I_n and K_n , and hence we can write $\alpha(r)$ as a linear combination of these:

$$\alpha(r) = \sqrt{r} \left(c_1 K_{\frac{N-1}{2}}(mr) + c_2 I_{\frac{N-1}{2}}(mr) \right), \quad (6.86)$$

where we have dropped the subscript f on the fermion number for purposes of presentation. Note that I_n grows exponentially at large r , while K_n decays, and thus to obtain normalisable solutions we should set $c_2 = 0$. We can then perform a similar calculation to obtain $\beta(r)$, the result of which is:

$$\beta(r) = c_1 \sqrt{r} K_{\frac{N+1}{2}}(mr). \quad (6.87)$$

Substituting these expressions into the Einstein equation (6.66), we now obtain the following equation for $T(r)$:

$$-2r \frac{T'}{T^2} = \frac{1}{2} \xi N m c_1^2 r \left(K_{\frac{N+1}{2}}(mr) \left[K_{\frac{N-3}{2}}(mr) + K_{\frac{N+1}{2}}(mr) \right] - K_{\frac{N-1}{2}}(mr) \left[K_{\frac{N-1}{2}}(mr) + K_{\frac{N+3}{2}}(mr) \right] \right). \quad (6.88)$$

For general N_f , this does not appear to reduce to anything tractable, although MATHEMATICA is able to find an analytic (although lengthy) solution involving hypergeometric functions. For physical acceptable states, however, the value of N_f must be even, which fortunately simplifies the problem considerably due to the properties of the half-integer Bessel functions. Again using MATHEMATICA, we can then ascertain that $T(r)$ must take the following form, valid in the case of even N_f :

$$T(r) = \left[1 + \pi^2 c_1^2 (mr)^{-N} e^{-2mr} \sum_{n=0}^{N-1} \lambda_n (mr)^n + 4(-1)^{\frac{N}{2}+1} \pi^2 c_1^2 N \text{Ei}(-2mr) \right]^{-1}, \quad (6.89)$$

where λ_n are numerical coefficients that vary depending on the value of N_f . We have been unable to obtain an explicit expression for these, although we suspect that in principle a closed form exists, which could be ascertained by undertaking a full analytic derivation of the above equation. The general form of $T(r)$ is similar to the two-fermion case, with the addition of a power series presumably arising from a repeated integration by parts. Note that we have verified that this solution does indeed reduce to that presented in the previous section when $N_f = 2$. As before, the value of c_1 here becomes constrained when imposing normalisation, resulting in a similar integral condition to (6.81).

We find that the properties of this many-fermion solution are qualitatively similar to the two-fermion case, with the exception that $\alpha(r)$ now also diverges as $r \rightarrow 0$ (for $N_f \geq 4$). Again, we find that the solution is singular, and that states can exist for any value of the fermion mass, with the behaviour as a function of m mirroring that presented in Fig. 6.12.

6.2.3 Interpretation

The precise interpretation of the analytic solution presented here is not entirely clear. Both the conditions $\omega = 0$ and $A(r) = 1$ are admittedly somewhat unusual, although neither necessarily proves problematic from a physical perspective. Certainly zero energy does not cause any issues in the type of semi-classical situation considered here, although it does imply that the fermion wavefunction is entirely static, and thus has no internal conception of time. The condition $A(r) = 1$ is also perfectly acceptable, with this simply corresponding to the absence of spatial variations in the metric. It is indeed unusual for this to be accompanied by anything other than a flat space-time, with the scenario here implying that the fermion source creates a time dilation, but without an accompanying

length contraction. We note of course that this impression of flat space is entirely restricted to inertial observers at $r = \infty$, and that a Lorentz boost could in principle be applied to remove this imbalance.

The physical properties of the solution are also difficult to interpret. As mentioned, the total ADM mass is always zero, and is therefore entirely decoupled from the fermion mass, this being a free parameter within the system. An observer far outside the fermion source will therefore measure the gravitational mass of the state to be almost negligible, despite it containing fermions of potentially infinite mass. The behaviour as one approaches the fermion source is even more counter-intuitive, with the state producing an overall repulsive force which increases with decreasing radius, ultimately preventing any test particle from ever reaching $r = 0$. We unfortunately have no physical explanation for these features, and simply remark that, for states with a central singularity, any type of behaviour appears possible.

6.3 Discussion

In this chapter, we have presented a variety of new singular solutions to the Einstein–Dirac system, both numerical and analytic, and shown that these often exhibit unusual properties, the only explanation for which appears to be related to the presence of the central singularity. Although the majority of these are far removed from the original FSY states, there are a number of cases where certain families of singular solutions can be continuously connected with the FSY curves in some limit, due to the peculiarities of their zonal structures.

There are a number of points to discuss. First, we note that the small- r analysis performed here assumes that all fields lead with a simple power of r , and hence we cannot exclude the possibility of alternative expansions existing in which one or more fields exhibit, for example, a logarithmic dependence at small r . It may even transpire that a non-singular solution can be constructed using such an expansion, although we consider this unlikely. Assuming this proves not to be the case, our analysis here has established that the FSY states are the only non-singular localised solutions to the Einstein–Dirac system, thereby cementing their position as the only potential particle-like objects within the theory.

How, then, should we view the various singular solutions presented in this chapter – do they represent possible physical states, or are they simply mathematical curiosities? The answer to this question most probably depends on their stability. If any prove to be perturbatively stable, then we see no reason why they should not be taken seriously as alternatives to the FSY states, despite the somewhat undesirable presence of singularities. To establish this would require either an analytic perturbation analysis or a full dynamical evolution, both of which are entirely feasible, although outwith the scope of this thesis.

Even if none of the cases presented here prove to be stable, they nonetheless represent interesting examples of the way in which quantum wavefunctions (and indeed classical fields) can interact with singularities in general relativity.

As a final remark, we note that the solutions presented here contain ‘naked singularities’, i.e. singularities for which there is no accompanying horizon, and hence the interior region of the spacetime is visible to an external observer. These are of considerable interest in the study of gravitational collapse, and have long been considered alternatives to black holes, although the general consensus is that such situations are either unphysical or unlikely to occur (see e.g. [93] for a review). Although the appearance of naked singularities in the context of the Einstein–Dirac system appears unrelated to this question, it is nonetheless interesting to observe that such structures do indeed exist, and that they do not require the presence of unphysical or exotic types of matter, at least in this scenario.

Chapter 7

Summary & Discussion

Within this thesis, we have presented a number of new discoveries related to the study of Dirac solitons. In particular, we have analysed in detail the many-fermion extension to the Einstein–Dirac system, and generated the first soliton-like solutions for configurations containing up to 90 fermions. We have demonstrated that these exhibit a fermion self-trapping effect, in which the fermion wavefunction becomes trapped around regions of spacetime containing stable photon spheres, resulting in the appearance of a multiple-shell-like structure. The corresponding excited states for the many-fermion system have also been thoroughly analysed, and found to become increasingly multivalued as the fermion number is increased, a feature that can be explained by considering the internal structure of individual solutions. Furthermore, we have presented particle-like solutions to the minimally-coupled Einstein–Dirac–Higgs system, and shown that these exhibit the same qualitative behaviour as those in the Einstein–Dirac case. The exception to this is the discovery of a mass-scale separation at strong coupling, in which the gravitational mass of the state and the mass of the constituent fermions become effectively decoupled. Finally, we have presented a selection of new singular solutions to the Einstein–Dirac system, generated by varying the leading-order behaviour of the fields at small radii, and shown that these exhibit somewhat counter-intuitive features.

Since we have presented a detailed discussion at the end of each chapter, we here confine our remarks to more general aspects related to Dirac solitons. First, we should briefly discuss the applicability of the Einstein–Dirac approach itself. As mentioned, this provides a semi-classical framework in which one can study the interaction between gravity and a quantum wavefunction, but the extent to which it can be considered a low-level approximation to a full theory of quantum gravity is perhaps debatable. In particular, one might be concerned as to the applicability of describing the quantum sector by a simple wavefunction, as opposed to a quantum field, although it is difficult to quantify precisely what information is lost by doing so. There are certainly some features that the Einstein–

Dirac system does indeed model, such as particle number conservation, but concepts such as the quantum vacuum are entirely absent. The implications for Einstein–Dirac solitons is unclear, and it remains to be seen whether these objects have counterparts within the context of quantum field theory. Taking an optimistic point of view, it might even be the case that the consideration of objects such as Dirac solitons can help to regularise the divergences that occur in quantum field theory, as discussed in ref. [94]. Alternatively, it may be that the inclusion of quantum vacuum fluctuations destroys the delicate balance required for localised states of this type to form, in which case it would be difficult to envisage any physical applications for Dirac solitons.

Regardless, the semi-classical Einstein–Dirac system still provides a rare context in which to study the interaction of quantum objects while including the full dynamics of general relativity. In particular, it automatically incorporates the effects of back-reaction, one of the most important features of general relativity, and one which is often (necessarily) neglected when considering quantum effects. Even from a purely classical point of view, the spacetimes associated with Dirac solitons are interesting in their own right, as we have demonstrated here with the self-trapping effect for example. Furthermore, much of the analysis in general relativity is confined to objects assumed to obey specific equations of state, or to vacuum solutions upon which matter is subsequently added. More exotic structures are often discarded due to concerns over physical applicability, but as we have seen here, it is entirely possible to obtain highly non-trivial spacetimes in which the matter content nonetheless remains physically acceptable. The study of systems such as the Einstein–Dirac could therefore potentially prove useful in the field of classical general relativity.

We conclude with some brief remarks concerning future directions. There are a number of potentially fruitful avenues related to Dirac solitons that are yet to be explored. These include issues relating to their interactions, particularly with regard to head-on collisions, which could establish whether these localised states are indeed solitons, in the true mathematical sense. It would also be interesting to establish the ultimate fate of unstable Dirac solitons, and whether they do indeed decay to corresponding stable states, or simply collapse to black-hole type structures. On a more ambitious note, we mention that constructing a soliton-like object corresponding to a Standard Model particle is not necessarily out of reach, since the electromagnetic field, electroweak interaction and Higgs mechanism have now all been individually addressed, and localised states demonstrated to exist in the context of each. The notable absence is of course the strong force, but this would not be required to model, say, a Standard Model electron. If a particle-like state could be found in a system combining these interactions, it would represent a significant step towards Dirac solitons being considered seriously as models for fundamental particles.

Appendix A

Low-redshift relations

In this appendix, we shall derive various scaling relations that exist between properties of Einstein–Dirac solitons in the non-relativistic limit (low redshift).

We begin with the relationship between the ADM mass of a state and its radius, which is observed to approximate $M \sim R^{-1/3}$ in the low-redshift regime. To derive this, we note that in the non-relativistic limit, localised states are held in equilibrium by the balance between their mutual Newtonian gravitational attraction and their kinetic energy. Equating these, for a system of N_f fermions each of mass m , gives:

$$N_f \frac{p^2}{2m} = N_f \frac{GMm}{R}, \quad (\text{A.1})$$

where p is the momentum of each fermion. Since we are dealing with quantum particles, we can relate momentum and displacement via the uncertainty principle: $\Delta p \Delta x \sim 1$. For a filled shell of fermions, the exclusion principle does not apply, and hence each fermion wavefunction occupies an effective volume of R^3 , implying $\Delta x \sim R$ and $p \sim 1/R$. Substituting this into the expression above gives:

$$\frac{1}{2mR^2} \sim \frac{GMm}{R}. \quad (\text{A.2})$$

In the Einstein–Dirac system, the fermion mass m is not a constant within the families of states, and should be eliminated in favour of the ADM mass M . To do so, we note that, in the non-relativistic limit, $M \approx N_f m$, and hence we obtain the following relation:

$$\frac{N_f}{2MR^2} \sim \frac{GM^2}{N_f R} \quad \implies \quad M \sim \left(\frac{2GR}{N_f^2} \right)^{-1/3}. \quad (\text{A.3})$$

Thus the families of states at constant N_f will approximate the relation $M \sim R^{-1/3}$ at low-redshift, as observed.

We now proceed to derive explicit non-relativistic relationships for the scalings of ω , m and R , as a function of central redshift z . We do so by considering the form of the metric field $T(r)$ for low-redshift states, noting that it will remain at a relatively constant value of $T(0) = 1 + z$ throughout the fermion source, before latching on to the Schwarzschild solution at approximately the radius of the soliton R . Matching these two regimes at the boundary, we obtain:

$$1 + z = \left(1 - \frac{2GM}{R}\right)^{-1/2} \approx 1 + \frac{GM}{R} \quad \implies \quad z \approx \frac{GM}{R}. \quad (\text{A.4})$$

Using the mass-radius relationship above, and the fact that $M \approx N_f m$, we can directly infer the following scaling relations:

$$R \sim z^{-3/4}; \quad M \sim z^{1/4}; \quad m \sim z^{1/4}. \quad (\text{A.5})$$

In order to include the fermion energy ω in this analysis, we require information concerning the ground state of the Newton-Schrödinger system. We can obtain this by invoking an analogy with the Bohr model of the Hydrogen atom, where we replace the electrostatic force with gravity, i.e.

$$\frac{e^2}{4\pi\epsilon_0} \rightarrow GMm. \quad (\text{A.6})$$

In the Bohr model, the energy of the ground state E_0 is proportional to the Rydberg constant R_H , which in turn can be related to the Bohr radius a_0 :

$$E_0 \sim R_H = \frac{1}{ma_0^2} = \frac{me^4}{16\pi\epsilon_0^2}. \quad (\text{A.7})$$

By analogy, we therefore conclude that the ground state of the Newton-Schrödinger system will have an energy $E_0 \sim G^2 M^2 m^3 \sim z^{5/4}$, where we have utilised the redshift scalings for m and M . Identifying this with the fermion binding energy ($m - \omega$), we obtain the relation:

$$(m - \omega) \sim z^{5/4} \quad \implies \quad \omega \sim z^{1/4}, \quad (\text{A.8})$$

where the final step follows since the binding energy relation forces m and ω to scale as the same power. Note that we have verified numerically that the various scalings presented here are indeed valid for low-redshift Einstein-Dirac solitons.

We note that the mass-radius relation $M \sim R^{-1/3}$ derived above differs from the known expressions for both astrophysical fermionic objects (neutron stars, white dwarfs etc.) and also boson stars. To see why, we shall briefly derive the relationships valid in those cases using the method outlined above. For astrophysical objects, the key difference is that the exclusion principle must now be incorporated, with the resulting degeneracy pressure effectively reducing the volume occupied by each fermion wavefunction to R^3/N_f .

Hence now $p \sim N_f^{1/3} R^{-1}$, and the balance equation (A.1) becomes:

$$\frac{M^{2/3}}{2m^{5/3}R^2} \sim \frac{GMm}{R} \quad \implies \quad M \sim \frac{1}{m^8 R^3}. \quad (\text{A.9})$$

Since the fermion mass is taken to be a constant, we therefore recover the often quoted $M \sim R^{-3}$ relationship for diffuse fermionic astrophysical objects.

For boson stars, there is again no degeneracy pressure, and the balance equation (A.1) is thus precisely the same as for the Einstein–Dirac case:

$$\frac{1}{2mR^2} \sim \frac{GMm}{R}. \quad (\text{A.10})$$

When considering boson stars, however, the lack of a normalisation condition results in the boson mass becoming a free parameter, while the number of bosons varies along the mass-radius curves, as previously discussed in section 1.3.1. Hence the relationship can be read as stated, i.e. $M \sim R^{-1}$, which should hold for boson stars with low central densities.

Appendix B

Data for Figures

In this appendix, we list the values of the various parameters and physical observables associated with the individual states presented in the figures throughout this thesis. These are divided into the relevant chapters, with the figure number indicated, followed by either an alphabetical or positional identifier to distinguish subfigures. Examples of these are (l): left, (m): middle and (r): right. Note that, for states that are included across multiple figures, only the first instance is recorded here.

Chapter 3:

Fig.	κ	n	z	m	ω	M	\bar{R}	R_{99}	$\tilde{\alpha}_1$
2.1 (l)	2	0	2.0041	0.50403	0.36939	1.0652	3.1766	10.711	0.17000
2.1 (m)	-2	1	2.0000	0.55054	0.43627	1.2467	3.8611	12.123	0.10875
2.1 (r)	2	8	2.0000	1.2406	1.0731	2.4741	14.197	25.543	0.15672
2.4 (l)	2	4	1.8194×10^6	1.1818	1.0822	2.3894	21.521	38.244	200.00
2.5	2	0	∞	0.41509	0.34523	1.0108	4.2815	16.737	–
3.1 (l)	90	0	0.091870	0.43925	0.41270	38.791	482.16	568.01	$1. \times 10^{-100}$
3.1 (m)	90	0	8.0166	0.34432	0.31473	43.725	172.46	315.24	1.2×10^{-30}
3.1 (r)	90	0	92.087	0.30048	0.28299	42.902	115.80	255.71	7.5×10^{15}
3.2 (l)	40	0	92.083	0.31086	0.28978	19.165	57.113	149.46	5.5094×10^{12}
3.2 (m)	20	0	92.113	0.31717	0.29206	9.5970	28.630	77.716	1.5601×10^8
3.2 (r)	6	0	92.104	0.34035	0.30286	2.8968	9.8418	33.031	322.14
3.4 (iv)	90	0	∞	0.30483	0.28611	43.025	121.05	293.25	–

Chapter 4:

Fig.	κ	n	z	m	ω	M	\bar{R}	R_{99}	$\tilde{\alpha}_1$
4.1 (A)	2	0	2.0000	0.50435	0.36961	1.0655	3.1778	10.744	0.16976
4.1 (B)	2	2	2.0001	0.71733	0.59135	1.4794	6.5071	16.100	0.16615
4.1 (C)	2	4	2.0002	0.91640	0.77680	1.8555	9.4184	19.674	0.16191
4.1 (D)	2	6	2.0003	1.0885	0.93448	2.1831	11.944	22.759	0.15894
4.5 (A)	12	2	1.8958	0.43709	0.32453	6.0672	14.609	25.952	0.00016581
4.5 (B)	12	2	3.9581	0.27822	0.26114	5.5711	13.149	42.091	0.021847
4.5 (C)	12	2	4.9498	0.29833	0.28181	5.6386	23.176	88.379	0.064947
4.5 (D)	12	2	5.1193	0.27964	0.26493	5.5761	14.969	65.819	0.078162
4.5 (E)	12	2	4.6535	0.33687	0.28819	5.7501	14.333	31.166	0.028376
4.5 (F)	12	2	3.0479	0.40829	0.35579	6.0531	24.403	51.046	0.0045503
4.5 (G)	12	2	2.3526	0.28674	0.27635	5.3308	48.092	147.81	0.0026826
4.5 (H)	12	2	2.2940	0.23612	0.23077	5.1683	11.405	19.579	0.0026248
4.5 (I)	12	2	2.2994	0.23473	0.23049	5.1661	11.785	19.938	0.0027268
4.5 (J)	12	2	2.3288	0.26697	0.26112	5.2551	72.098	248.76	0.0028097
4.5 (K)	12	2	2.7541	0.41758	0.38664	6.1931	45.823	96.279	0.0041705
4.5 (L)	12	2	5.1160	0.43367	0.36668	6.3370	22.115	45.766	0.042004
4.6 (A)	20	0	39.891	0.32270	0.29554	9.6314	29.173	75.938	47012
4.6 (B)	20	2	39.706	0.31187	0.28968	9.5616	29.499	87.679	47012
4.6 (C)	20	2	39.912	0.32191	0.29474	9.6261	28.784	74.598	47012
4.6 (D)	20	2	39.490	0.34589	0.31992	9.8406	44.932	115.15	47012
4.8 (A)	38	0	41.015	0.30735	0.28578	18.156	49.597	114.80	664250
4.8 (B)	38	2	40.975	0.30968	0.28750	18.186	50.586	118.52	559350
4.8 (C)	38	2	41.018	0.30733	0.28577	18.155	49.593	114.81	665470
4.8 (D)	38	2	41.038	0.31233	0.29103	18.220	55.869	148.64	696140
4.8 (E)	38	2	41.000	0.30017	0.28190	18.075	48.923	130.17	806480
4.8 (F)	38	2	40.943	0.33585	0.30899	18.555	67.679	159.94	536910
4.10 (A)	38	4	39.962	0.31424	0.29110	18.244	53.103	127.67	331130
4.10 (B)	38	4	39.973	0.31113	0.28845	18.203	51.004	118.85	349710
4.10 (C)	38	4	39.957	0.33370	0.30805	18.527	68.981	167.48	322620
4.10 (D)	38	4	39.969	0.33682	0.30972	18.570	67.977	159.97	344780
4.10 (E)	38	4	39.999	0.30174	0.28265	18.092	48.835	122.58	494560
4.10 (F)	38	4	39.989	0.31062	0.29000	18.199	55.806	152.65	464950
4.10 (G)	38	4	40.004	0.31126	0.29317	18.216	69.241	215.09	544850
4.10 (H)	38	4	39.997	0.32862	0.30433	18.451	67.863	169.51	394400
4.10 (I)	38	4	39.989	0.34797	0.32160	18.775	83.346	194.30	379700

Chapter 5:

Fig.	κ	n	ξ	ν	z	m_f	m_H	ω	M	\bar{R}	R_{99}	$\tilde{\alpha}_1$
5.2 (l)	2	0	0.6	0.1	0.1338	0.3761	0.2256	0.3344	0.7415	9.840	28.96	0.025
5.2 (m)	2	0	0.6	0.1	1.881	0.7408	0.4445	0.2638	0.9733	2.540	6.876	0.08
5.2 (r)	2	8	0.2828	0.1	1.172	1.161	0.3283	0.9859	2.302	13.84	24.35	0.075
5.7 (a)	2	0	0.02828	0.1	2.151	3.665	0.1037	0.2176	0.8952	1.937	10.03	0.01039
5.7 (b)	2	0	0.2828	0.1	2.167	1.155	0.3267	0.2413	0.9541	2.261	7.522	0.05377
5.7 (c)	2	0	2.263	0.1	2.147	0.4569	1.034	0.3105	0.9894	2.881	9.781	0.1319
5.11	2	0	0.2828	0.1	34180	0.4724	0.1336	0.2591	0.9421	2.901	13.6	15.00
5.12	2	0	0.2828	0.1	∞	0.4677	0.1323	0.2596	0.9417	2.908	12.45	–
5.13	20	0	0.1414	0.05	28.22	0.2648	0.03745	0.2324	8.974	20.65	37.11	1.00

Chapter 6:

Fig.	κ	n	m	ω	M	\bar{R}	Parameter values
6.1	1	0	0.51273	0.41663	0.55386	2.9351	–
6.3 (a)	-2	0	467.92	2.5200	0.038176	0.073515	$p_0 = 0.05$
6.3 (b)	-2	0	8.4574	0.74590	0.11606	0.28676	$p_0 = 1$
6.3 (c)	-2	0	0.095187	0.031723	0.00013111	0.69391	$p_0 = 1000$
6.5 (a)	2	0	0.84245	0.0097571	0.073589	0.12352	$\tilde{\alpha}_1 = 10$
6.5 (b)	2	0	0.78891	0.14573	0.16093	0.56197	$\tilde{\alpha}_1 = 0.02$
6.7 (a)	-2	0	0.029827	0.0077174	-0.047010	2.2698	$p_0 = 4000, p_1 = 0.4$
6.7 (b)	-2	0	0.54505	-0.19594	-1.9180	2.2994	$p_0 = 10, p_1 = 0.4$
6.9 (a)	2	0	5.5920	0.66886	1.3243	2.2741	$p_0 = 500, p_1 = 0.1$
6.9 (b)	2	0	0.52163	0.49158	1.0246	13.734	$p_0 = 0.0001, p_1 = 0.1$
6.11 (a)	2	0	0.20936	-0.057728	-2.5011	18.462	$p_0 = 0.05, p_1 = 1, p_2 = 0.01$
6.11 (b)	2	0	0.41429	0.34518	1.0101	22.178	$p_0 = 1, p_1 = 1, p_2 = 0.01$

Bibliography

- [1] J. A. Wheeler, *Geons*, [Phys. Rev.](#) **97**, 511 (1955).
- [2] D. J. Kaup, *Klein-Gordon Geon*, [Phys. Rev.](#) **172**, 1331 (1968).
- [3] F. Finster, J. Smoller, and S.-T. Yau, *Particlelike solutions of the Einstein-Dirac equations*, [Phys. Rev. D](#) **59**, 104020 (1999).
- [4] D. Bakucz Canário, S. Lloyd, K. Horne, and C. A. Hooley, *Infinite-redshift localized states of Dirac fermions under Einsteinian gravity*, [Phys. Rev. D](#) **102**, 084049 (2020).
- [5] P. E. D. Leith, C. A. Hooley, K. Horne, and D. G. Dritschel, *Fermion self-trapping in the optical geometry of Einstein-Dirac solitons*, [Phys. Rev. D](#) **101**, 106012 (2020).
- [6] P. E. D. Leith, C. A. Hooley, K. Horne, and D. G. Dritschel, *Nonlinear effects in the excited states of many-fermion Einstein-Dirac solitons*, [Phys. Rev. D](#) **104**, 046024 (2021).
- [7] P. E. D. Leith, A. Dorkenoo Leggat, C. A. Hooley, K. Horne, and D. G. Dritschel, *Gravitationally bound states of two neutral fermions and a Higgs field can be parametrically lighter than their constituents*, [arXiv preprint arXiv:2202.03228](#) (2022).
- [8] M. H. Goroff and A. Sagnotti, *The ultraviolet behavior of Einstein gravity*, [Nucl. Phys. B](#) **266**, 709–736 (1986).
- [9] S. Mukhi, *String theory: a perspective over the last 25 years*, [Class. Quant. Grav.](#) **28**, 153001 (2011).
- [10] C. Rovelli, *Loop quantum gravity: the first 25 years*, [Class. Quant. Grav.](#) **28**, 153002 (2011).
- [11] S. Surya, *The causal set approach to quantum gravity*, [Living Rev. Relativ.](#) **22**, 1–75 (2019).
- [12] A. Eichhorn, *An asymptotically safe guide to quantum gravity and matter*, [Front. Astron. Space Sci.](#) **5**, 47 (2019).

- [13] P. D. Mannheim, *Making the case for conformal gravity*, *Found. Phys.* **42**, 388–420 (2012).
- [14] R. Penrose, *On the gravitization of quantum mechanics 1: Quantum state reduction*, *Foundations of Physics* **44**, 557–575 (2014).
- [15] J. A. Wheeler and K. Ford, *Geons, black holes, and quantum foam: A life in physics* (W. W. Norton & Company, 1998).
- [16] S. W. Hawking, *Particle creation by black holes*, *Commun. Math. Phys* **43**, 199–220 (1975).
- [17] R. H. Brandenberger, *Quantum field theory methods and inflationary universe models*, *Rev. Mod. Phys.* **57**, 1 (1985).
- [18] D. R. Brill and J. B. Hartle, *Method of the self-consistent field in general relativity and its application to the gravitational geon*, *Phys. Rev.* **135**, B271 (1964).
- [19] P. R. Anderson and D. R. Brill, *Gravitational geons revisited*, *Phys. Rev. D* **56**, 4824 (1997).
- [20] G. P. Perry and F. I. Cooperstock, *Stability of gravitational and electromagnetic geons*, *Class. Quant. Grav.* **16**, 1889 (1999).
- [21] O. J. Dias, G. T. Horowitz, D. Marolf, and J. E. Santos, *On the nonlinear stability of asymptotically anti-de Sitter solutions*, *Class. Quant. Grav.* **29**, 235019 (2012).
- [22] D. A. Feinblum and W. A. McKinley, *Stable states of a scalar particle in its own gravitational field*, *Phys. Rev.* **168**, 1445 (1968).
- [23] T. D. Lee and Y. Pang, *Stability of mini-boson stars*, *Nucl. Phys. B* **315**, 477–516 (1989).
- [24] M. Gleiser and R. Watkins, *Gravitational stability of scalar matter*, *Nucl. Phys. B* **319**, 733–746 (1989).
- [25] F. V. Kusmartsev, E. W. Mielke, and F. E. Schunck, *Stability of neutron and boson stars: a new approach based on catastrophe theory*, *Phys. Lett. A* **157**, 465–468 (1991).
- [26] F. E. Schunck and E. W. Mielke, *General relativistic boson stars*, *Class. Quant. Grav.* **20**, R301 (2003).
- [27] S. L. Liebling and C. Palenzuela, *Dynamical boson stars*, *Living Rev. Relativ.* **20**, 5 (2017).
- [28] J. Lee and I. Koh, *Galactic halos as boson stars*, *Phys. Rev. D* **53**, 2236 (1996).

- [29] J. Eby, C. Kouvaris, N. G. Nielsen, and L. C. R. Wijewardhana, *Boson stars from self-interacting dark matter*, *J. High Energy Phys.* **2016** (2), 1–19.
- [30] D. F. Torres, S. Capozziello, and G. Lambiase, *Supermassive boson star at the galactic center?*, *Phys. Rev. D* **62**, 104012 (2000).
- [31] N. Sennett, T. Hinderer, J. Steinhoff, A. Buonanno, and S. Ossokine, *Distinguishing boson stars from black holes and neutron stars from tidal interactions in inspiraling binary systems*, *Phys. Rev. D* **96**, 024002 (2017).
- [32] E. Braaten and H. Zhang, *Colloquium: The physics of axion stars*, *Rev. Mod. Phys.* **91**, 041002 (2019).
- [33] E. Seidel and W.-M. Suen, *Dynamical evolution of boson stars: Perturbing the ground state*, *Phys. Rev. D* **42**, 384 (1990).
- [34] F. S. Guzmán, *Evolving spherical boson stars on a 3d cartesian grid*, *Phys. Rev. D* **70**, 044033 (2004).
- [35] F. S. Guzmán, *The three dynamical fates of boson stars*, *Rev. Mex. Fis.* **55**, 321–326 (2009).
- [36] R. Brito, V. Cardoso, C. A. Herdeiro, and E. Radu, *Proca stars: gravitating Bose–Einstein condensates of massive spin 1 particles*, *Phys. Lett. B* **752**, 291–295 (2016).
- [37] A. B. Henriques, A. R. Liddle, and R. G. Moorhouse, *Combined boson-fermion stars*, *Phys. Lett. B* **233**, 99–106 (1989).
- [38] F. Di Giovanni, S. Fakhry, N. Sanchis-Gual, J. C. Degollado, and J. A. Font, *Dynamical formation and stability of fermion-boson stars*, *Phys. Rev. D* **102**, 084063 (2020).
- [39] R. Ruffini and S. Bonazzola, *Systems of self-gravitating particles in general relativity and the concept of an equation of state*, *Phys. Rev.* **187**, 1767 (1969).
- [40] T. D. Lee and Y. Pang, *Fermion soliton stars and black holes*, *Phys. Rev. D* **35**, 3678 (1987).
- [41] F. Finster, J. Smoller, and S.-T. Yau, *Particle-like solutions of the Einstein–Dirac–Maxwell equations*, *Phys. Lett. A* **259**, 431–436 (1999).
- [42] F. Finster, J. Smoller, and S.-T. Yau, *The interaction of Dirac particles with non-Abelian gauge fields and gravity-bound states*, *Nucl. Phys. B* **584**, 387–414 (2000).
- [43] F. Finster, J. Smoller, and S.-T. Yau, *Non-existence of black hole solutions for a*

- spherically symmetric, static Einstein–Dirac–Maxwell system*, *Commun. Math. Phys.* **205**, 249–262 (1999).
- [44] F. Finster, S.-T. Yau, and J. Smoller, *The interaction of Dirac particles with non-Abelian gauge fields and gravity—black holes*, *Mich. Math. J.* **47**, 199–208 (1999).
- [45] F. Finster, J. Smoller, and S.-T. Yau, *Non-existence of time-periodic solutions of the Dirac equation in a Reissner–Nordström black hole background*, .
- [46] F. Finster, J. Smoller, and S.-T. Yau, *Absence of static, spherically symmetric black hole solutions for Einstein–Dirac–Yang/Mills equations with complete fermion shells*, *Adv. Theor. Math. Phys.* **4**, 1231–1257 (2000).
- [47] Y. Bernard, *Non-existence of black-hole solutions for the electroweak Einstein–Dirac–Yang/Mills equations*, *Class. Quant. Grav.* **23**, 4433 (2006).
- [48] E. J. Bird, *A proof of existence of particle-like solutions of Einstein Dirac equations*, *Ph.D. thesis*, University of Michigan (2005).
- [49] S. R. Nodari, *Perturbation method for particle-like solutions of the Einstein–Dirac equations*, *Ann. Henri Poincaré* **10**, 1377–1393 (2010).
- [50] S. R. Nodari, *Perturbation method for particle-like solutions of the Einstein–Dirac–Maxwell equations*, *C. R. Math.* **348**, 791–794 (2010).
- [51] D. Stuart, *Existence and Newtonian limit of nonlinear bound states in the Einstein–Dirac system*, .
- [52] D. Giulini and A. Großardt, *The Schrödinger–Newton equation as a non-relativistic limit of self-gravitating Klein–Gordon and Dirac fields*, *Class. Quant. Grav.* **29**, 215010 (2012).
- [53] A. Dorkenoo Leggat, *Dirac solitons in general relativity and conformal gravity*, *Ph.D. thesis*, University of St Andrews (2017).
- [54] J. L. Blázquez-Salcedo and C. Knoll, *Constructing spherically symmetric Einstein–Dirac systems with multiple spinors: Ansatz, wormholes and other analytical solutions*, *Eur. Phys. J. C.* **80**, 1–17 (2020).
- [55] J. L. Blázquez-Salcedo, C. Knoll, and E. Radu, *Boson and Dirac stars in $D \geq 4$ dimensions*, *Phys. Lett. B* **793**, 161–168 (2019).
- [56] C. Herdeiro, P. Perapechka, E. Radu, and Y. Shnir, *Asymptotically flat spinning scalar, Dirac and Proca stars*, *Phys. Lett. B* **797**, 134845 (2019).

- [57] C. A. R. Herdeiro, A. M. Pombo, and E. Radu, *Asymptotically flat scalar, Dirac and Proca stars: Discrete vs. continuous families of solutions*, *Phys. Lett. B* **773**, 654–662 (2017).
- [58] V. Dzhunushaliev and V. Folomeev, *Dirac stars supported by nonlinear spinor fields*, *Phys. Rev. D* **99**, 084030 (2019).
- [59] V. Dzhunushaliev and V. Folomeev, *Dirac star in the presence of Maxwell and Proca fields*, *Phys. Rev. D* **99**, 104066 (2019).
- [60] V. Dzhunushaliev and V. Folomeev, *Dirac star with $SU(2)$ Yang-Mills and Proca fields*, *Phys. Rev. D* **101**, 024023 (2020).
- [61] E. Daka, N. N. Phan, and B. Kain, *Perturbing the ground state of Dirac stars*, *Phys. Rev. D* **100**, 084042 (2019).
- [62] C. W. K. Lai, *A numerical study of boson stars*, *Ph.D. thesis*, University of British Columbia (2004).
- [63] M. W. Choptuik and F. Pretorius, *Ultrarelativistic particle collisions*, *Phys. Rev. Lett.* **104**, 111101 (2010).
- [64] M. Alcubierre, J. Barranco, A. Bernal, J. C. Degollado, A. Diez-Tejedor, M. Megevand, D. Núñez, and O. Sarbach, *ℓ -boson stars*, *Class. Quant. Grav.* **35**, 19LT01 (2018).
- [65] V. Jaramillo, N. Sanchis-Gual, J. Barranco, A. Bernal, J. C. Degollado, C. Herdeiro, M. Megevand, and D. Núñez, *Head-on collisions of ℓ -boson stars*, *Phys. Rev. D* **105**, 104057 (2022).
- [66] R. M. Wald, *General relativity* (University of Chicago press, 2010).
- [67] S. Weinberg, *Gravitation and cosmology: Principles and applications of the general theory of relativity* (John Wiley and Sons, 1972).
- [68] J. Yopez, *Einstein's vierbein field theory of curved space*, *arXiv preprint arXiv:1106.2037* (2011).
- [69] F. Finster, *Local $U(2, 2)$ symmetry in relativistic quantum mechanics*, *J. Math. Phys.* **39**, 6276–6290 (1998).
- [70] J. J. Sakurai, *Advanced Quantum Mechanics* (Addison-Wesley, 1982).
- [71] J. S. Dowker and G. Kennedy, *Finite temperature and boundary effects in static space-times*, *J. Phys. A* **11**, 895 (1978).

- [72] G. W. Gibbons and M. J. Perry, *Black holes and thermal green functions*, *Proc. R. Soc. A* **358**, 467–494 (1978).
- [73] M. A. Abramowicz, B. Carter, and J.-P. Lasota, *Optical reference geometry for stationary and static dynamics*, *Gen. Relativ. Gravit.* **20**, 1173–1183 (1988).
- [74] S. Kristiansson, S. Sonogo, and M. A. Abramowicz, *Optical space of the Reissner-Nordström solutions*, *Gen. Relativ. Gravit.* **30**, 275–288 (1998).
- [75] Z. Stuchlík, S. Hledík, and J. Jurán, *Optical reference geometry of Kerr-Newman spacetimes*, *Class. Quant. Grav.* **17**, 2691 (2000).
- [76] J. Kovář and Z. Stuchlík, *Optical reference geometry and inertial forces in Kerr–de Sitter spacetimes*, *Class. Quant. Grav.* **24**, 565 (2006).
- [77] S. Hledík, Z. Stuchlík, and A. Cipko, in *AIP Conf. Proc.*, Vol. 861 (American Institute of Physics, 2006) pp. 883–890.
- [78] M. A. Abramowicz, N. Andersson, M. Bruni, P. Ghosh, and S. Sonogo, *Gravitational waves from ultracompact stars: the optical geometry view of trapped modes*, *Class. Quant. Grav.* **14**, L189 (1997).
- [79] S. Chandrasekhar and V. Ferrari, *On the non-radial oscillations of a star. III. a reconsideration of the axial modes*, *Proc. R. Soc. Lond. A* **434**, 449–457 (1991).
- [80] M. A. Abramowicz and A. R. Prasanna, *Centrifugal force reversal near a Schwarzschild black-hole*, *Mon. Not. R. Astron. Soc.* **245**, 720 (1990).
- [81] M. A. Abramowicz, *Black holes and the centrifugal force paradox*, *Sci. Am.* **268**, 74–81 (1993).
- [82] M. A. Abramowicz, J. C. Miller, and Z. Stuchlík, *Concept of radius of gyration in general relativity*, *Phys. Rev. D* **47**, 1440 (1993).
- [83] S. Sonogo and A. Lanza, *Relativistic perihelion advance as a centrifugal effect*, *Mon. Not. R. Astron. Soc.* **279**, L65–L66 (1996).
- [84] C. Cederbaum and G. J. Galloway, *Uniqueness of photon spheres in electro-vacuum spacetimes*, *Class. Quant. Grav.* **33**, 075006 (2016).
- [85] C. Cederbaum and G. J. Galloway, *Uniqueness of photon spheres via positive mass rigidity*, *Commun. Anal. Geom.* **25**, 303–320 (2017).
- [86] Q. Gan, P. Wang, H. Wu, and H. Yang, *Photon spheres and spherical accretion image of a hairy black hole*, *Phys. Rev. D* **104**, 024003 (2021).

-
- [87] G. Guo, P. Wang, H. Wu, and H. Yang, *Quasinormal modes of black holes with multiple photon spheres*, *J. High Energy Phys.* **2022** (6), 1–24.
- [88] M. Karlovini, K. Rosquist, and L. Samuelsson, *Constructing stellar objects with multiple necks*, *Class. Quant. Grav.* **18**, 817 (2001).
- [89] M. Karlovini, K. Rosquist, and L. Samuelsson, *Ultracompact stars with multiple necks*, *Mod. Phys. Lett. A* **17**, 197–203 (2002).
- [90] M. Alcubierre, J. Barranco, A. Bernal, J. C. Degollado, A. Diez-Tejedor, V. Jaramillo, M. Megevand, D. Núñez, and O. Sarbach, *Extreme ℓ -boson stars*, *Class. Quant. Grav.* **39**, 094001 (2022).
- [91] L. G. Collodel, B. Kleihaus, and J. Kunz, *Excited boson stars*, *Phys. Rev. D* **96**, 084066 (2017).
- [92] S. Schlögel, M. Rinaldi, F. Staelens, and A. Füzfa, *Particlelike solutions in modified gravity: the Higgs monopole*, *Phys. Rev. D* **90**, 044056 (2014).
- [93] P. S. Joshi and D. Malafarina, *Recent developments in gravitational collapse and spacetime singularities*, *Int. J. Mod. Phys. D* **20**, 2641–2729 (2011).
- [94] F. Finster, J. Smoller, and S.-T. Yau, *The coupling of gravity to spin and electromagnetism*, *Mod. Phys. Lett. A* **14**, 1053–1057 (1999).