Structure Formation in Quantum-wave Dark Matter Cosmologies

Simon Benjamin May

München 2022

Structure Formation in Quantum-wave Dark Matter Cosmologies

Simon Benjamin May

Dissertation
an der Fakultät für Physik
der Ludwig-Maximilians-Universität
München

vorgelegt von
Simon Benjamin May
aus Münster

München, den 26. August 2022

Erstgutachter: Prof. Dr. Volker Springel
Zweitgutachter: Prof. Dr. Eiichiro Komatsu
Tag der mündlichen Prüfung: 30. September 2022

Contents

List of Figures

List of Tables

Zusammenfassung

Obwohl das sogenannte kosmologische Standardmodell auf vielen physikalischen Längenskalen erfolgreiche Vorhersagen macht, liefert es keine Erklärung für seine beiden zentralen Bestandteile: Die (kalte) dunkle Materie und die dunkle Energie (in Form einer kosmologischen Konstante). Dieses „dunkle Universum", welches über 95 % des gesamten kosmischen Energiebudgets ausmacht, entzieht sich unserer Kenntnis: Es ist nicht bekannt, aus welchen elementaren physikalischen Komponenten es besteht.

Speziell bezüglich der dunklen Materie hat sich der Fokus in der näheren Vergangenheit verschoben, da nach Jahrzehnten intensiver teilchen- und astrophysikalischer Suche keine Spur von den bisher favorisierten Kandidaten zu finden ist. Ultraleichte Skalarteilchen stellen eine Alternative zu diesen Kandidaten dar, die durch ihre reichhaltige astrophysikalische Phänomenologie interessante Möglichkeiten für ihre potentielle Entdeckung liefern. Aufgrund ihrer extrem geringen Masse verhalten sie sich auf astrophysikalischen Skalen nicht als individuelle Teilchen, sondern kollektiv als Wellen. Dies resultiert in einer Vielzahl von Wellenphänomenen, wie etwa die Bildung von Solitonen und Interferenzmustern oder auf kurzen Zeitskalen oszillierende Dichtefluktuationen, die eher an quantenmechanische Effekte als an makroskopische Strukturen erinnern.

Im Rahmen dieser Dissertation befasse ich mich mit kosmologischen Modellen, in denen die dunkle Materie aus ebensolchen ultraleichten Bosonen besteht. Zu diesem Zweck setze ich umfangreiche numerische Simulationen der kosmischen Strukturbildung ein, die in der Lage sind, anhand der nicht-linearen Evolution von Strukturen im Universum entscheidende physikalische Unterschiede zwischen diesem Modell dunkler Materie und dem Standardmodell hervorzuheben. Als ein wichtiges Ziel und Werkzeug der Dissertation habe ich dementsprechend das Programm AxiREPO entwickelt, das die entsprechenden Bewegungsgleichungen ultraleichter dunkler Materie numerisch löst und so Simulationen der erwarteten kosmischen Strukturbildung berechnen kann.

Mithilfe dieses Programms habe ich große Simulationen ultraleichter und kalter

dunkler Materie geplant, durchgeführt und analysiert. Hierbei wurden insbesondere verschiedene Anfangsbedingungen verwendet, um sowohl den Einfluss von Unterschieden in den primordialen Dichtefluktuationen im Vergleich zu solchen, die von der Dynamik der zu lösenden Bewegungsgleichungen herrühren, als auch verschiedene Werte für die Masse der ultraleichten Bosonen sowie die Berücksichtigung von baryonischer Materie untersuchen und vergleichen zu können.

Die Dissertation ist wie folgt strukturiert: Kapitel 1 beginnt mit einer Einführung in das aktuelle Standardmodell der Kosmologie mit besonderem Augenmerk auf die empirischen Hinweise auf dunkle Materie und ihre verschiedenen Erklärungsansätze. Kapitel 2 konzentriert sich auf das Modell der ultraleichten dunklen Materie und stellt dieses detailliert vor, da es sich hierbei um das zentrale Thema der Arbeit handelt. Darauf folgt in Kapitel 3 die Erläuterung der angewandten numerischen Prinzipien und Methoden. Die wesentlichen Ergebnisse der Arbeit werden ab Kapitel 4 vorgestellt. In diesem Kapitel werden zunächst Simulationen thematisiert, die mit Anfangsbedingungen für kalte dunkle Materie berechnet wurden. Hierdurch kann ein direkter Vergleich zwischen dem Einfluss der Dynamik kalter und ultraleichter Materie gezogen werden, und auf Grundlage des resultierenden Katalogs von Strukturen aus dunkler Materie wird die Frage beantwortet, in welchem Zusammenhang die Massen eines Halos und des zugehörigen Solitons stehen. In Kapitel 5 wird mit vollständigen Simulationen ultraleichter dunkler Materie (inklusive der entsprechenden Anfangsbedingungen) fortgefahren und zusammen mit den Ergebnissen des vorherigen Kapitels die Unterschiede von Anfangsbedingungen und nicht-linearer Dynamik genauer beleuchtet. Darüber hinaus werden zum ersten Mal die Halo-Massenfunktion ultraleichter dunkler Materie in Simulationen der vollständigen Wellendynamik gemessen und die filamentären Strukturen, die sich in ihrer Beschaffenheit stark von denen des Modells kalter dunkler Materie unterscheiden, untersucht. Mit Kapitel 6 folgen schließlich Simulationen, die auch baryonische Materie berücksichtigen, und es werden erste Ergebnisse aus diesen präsentiert. Dabei werden sowohl Unterschiede zu den Simulationen, die ausschließlich dunkle Materie beinhalten, herausgestellt, als auch der Lyman-α-Wald untersucht. Schließlich wird in Kapitel 7 ein Fazit gezogen und ein Ausblick auf weitere Arbeiten auf Grundlage der Ergebnisse dieser Dissertation sowie die (mögliche) weitere Entwicklung des Feldes gegeben.

Summary

Although the so-called standard model of cosmology has been able to make successful predictions on many physical length scales, it does not provide an explanation for its two central components: (Cold) dark matter and dark energy (in the form of a cosmological constant). This "dark universe", which makes up more than 95 % of the total cosmic energy budget, still eludes our grasp: It is not known which elementary physical components it is composed of.

Particularly with regard to dark matter, the focus has shifted in the recent past, since there has been no trace of the candidates favored thus far even after decades of intensive experimental and observational search using particle- and astrophysical approaches. Ultra-light scalar particles represent an alternative to these candidates, offering intriguing possibilities for their potential detection due to their rich astrophysical phenomenology. Because of their extremely small masses, they do not behave as individual particles, but collectively as waves. This results in a multitude of wave phenomena, such as the formation of solitons and interference patterns or transient, oscillating density fluctuations which are rather reminiscent of quantum-mechanical effects than macroscopic structures.

In the course of this dissertation, I will consider cosmological models in which dark matter is composed of exactly such ultra-light bosons. To this end, I employ extensive numerical simulations of cosmic structure formation, which are capable of discerning key physical differences between this model of dark matter and the standard model by means of the non-linear evolution of structure in the universe. As an important goal and tool within the dissertation, I developed the AxiREPO code, which numerically solves the corresponding equations of motions for ultra-light dark matter and can thus compute simulations of the expected formation of cosmic structure.

Using this code, I designed, executed and analyzed large simulations of ultra-light and cold dark matter. In particular, different initial conditions were used in order to be able to study and compare both the influence of differences in the primordial density fluctuations as opposed to those which originate due to the dynamics of

the equations of motions, as well as different values for the masses of the ultra-light bosons and accounting for baryonic matter.

The dissertation is structured as follows: Chapter 1 begins with an introduction to the current standard model of cosmology, with special focus on the empirical evidence for dark matter and the different proposals for its explanation. Chapter 2 is focused on the ultra-light dark matter model and presents this in detail, since this is the central theme of this work. Subsequently following in chapter 3 is an explanation of the numerical principles and methods used. The main results of this work are presented starting with chapter 4. This chapter first discusses simulations which were carried out with initial conditions for cold dark matter. In this way, it is possible to draw a direct comparison between the impact of the dynamics of cold and ultra-light dark matter, and based on the resulting catalog of dark matter structure, the question about the relation between the masses of halos and their central solitons is answered. This continues in chapter 5 with full simulations of ultra-light dark matter (including the appropriate initial conditions), and together with the results from the previous chapter, the differences between initial conditions and non-linear dynamics is examined more closely. Further, the halo mass function of ultra-light dark matter is measured for the first time in simulations making use of the full wave dynamics, and the filamentary structures, which differ quite strikingly from those in the cold dark matter model, are investigated. With chapter 6, simulations follow which take baryonic matter into account, and first results from these are presented. Thereby, both differences to the dark-matter only simulations are highlighted, and the Lyman-α forest is studied. Finally, chapter 7 draws the conclusions, giving an outlook on future work based on the foundations of this dissertation as well as the (possible) future development of the field.

1 Introduction

The "standard model of cosmology" of cold dark matter and a cosmological constant (ΛCDM) has been extremely successful in describing a plethora of cosmological observations across a broad range of physical scales (e. g. Frenk and White 2012; Bull et al. 2016). However, it is in some sense a victim of its own success: Being an empirical, phenomenological model at its core, ΛCDM does not provide any answers to questions about the origins or nature of its eponymous components, the cosmological constant Λ and cold dark matter (CDM). Simply put, although dark energy and dark matter (Λ and CDM) constitute a combined $\approx 95\,\%$ of the universe's total energy budget, we do not actually know what they *are*. In other words, we a lack micro-physical description of ΛCDM's components – a more fundamental, lower-level model. The situation is comparable to classical thermodynamics in its phenomenological form, before its laws were recognized more fundamentally as a manifestation of predictions from statistical mechanics. Thus, while ΛCDM provides a very accurate and consistent description of what can be observed in the universe, the enduring lack of any serious discrepancies also means that there are no hints that could lead towards a true understanding of the nature of dark matter and dark energy.

Remarkably, there is a similar and related problem in particle physics, which features its own standard model. Like ΛCDM, the standard model of particle physics is in excellent agreement with countless experiments performed over decades. In fact, the problem here is arguably even more severe: Firstly, the standard model of particle physics has been tested to much higher levels of precision than ΛCDM, which is just now entering an "era of precision cosmology", and secondly, contrary to ΛCDM, it is known to be *incomplete*. Indeed, partially, it is none other than the establishment and confirmation of ΛCDM itself which spells disaster for the standard model of particle physics, which neither provides a suitable candidate for CDM, nor matches the required value for the cosmological constant.[1] Thus, we have another case of a

[1] Strictly speaking, there are possible dark matter candidates which do not require the existence of new elementary particles. However, there are more areas where the standard model of particle

model which does not tell the whole story, but whose wildly successful predictions increasingly preclude avenues towards its unavoidably necessary extension. For particle physics, this even goes so far as to elicit calls proclaiming a crisis of the entire field. Hence, the question of the nature of dark matter is one of the most pressing and crucial questions facing both cosmology and particle physics today, and its resolution would provide one of the deepest insights into the fundamental laws of the cosmos.

For a long time, weakly interacting massive particles (WIMPs) have been the favored dark matter candidates. Originally, these WIMPs were primarily in the form of particles postulated in supersymmetry (SUSY), a generalization of the standard model of particle physics, such as neutralinos, but more recently, other extensions to the standard model, with particles participating in the weak interaction, have gained prominence as well (e. g. Restrepo, Zapata, and Yaguna 2013; Fiaschi, Klasen, and May 2019). However, the justification for WIMPs as a whole is dwindling as dark matter detection experiments continue to constrain the viability of WIMP models. As a consequence, the focus is increasingly shifting towards a broader, more diverse range of dark matter candidates without any traces of detection.

On the astrophysical side, the perspective is rather different. CDM is a quite generic model, leaving a lot of room for more concrete implementations as a particular particle physics model, or even as a completely different substance not composed of unknown elementary particles (such as primordial black holes). The reason for this is that cosmological and astrophysical observations are often insensitive to the specific details of a particle dark matter model as long as it behaves similarly enough to CDM, with the only remnant of model-specific particle properties often being a combined annihilation or self-interaction cross section. Thus, the question here is more about the validity of the CDM model itself, or whether there are any conflicts with observations which other models (which would in turn inform the selection of viable particle candidates) would rectify. Although ΛCDM has been very successful in making predictions and matching observations, especially on large cosmological scales, there are a number of contentious small-scale tensions which have been a matter of continuous debate. Unfortunately, since these are also exactly the scales at which physical effects experienced by baryons become relevant, it has proven difficult to disentangle the apparent discrepancies from baryonic physics, leaving the implications of these tensions much less clear. Though it would be justified to consequently conclude that CDM is, at least for the moment, sufficient as

physics is incomplete, such as the existence of massive neutrinos (another observation from cosmology) and the coupling to gravity.

a cosmological and astrophysical model of dark matter, such an attitude can offer no directions that further the goal of understanding more properties, and ultimately, the micro-physical nature, of dark matter. The main driver behind the investigation of dark matter models alternative to CDM, common to the cosmological, astrophysical, and particle physics perspective, is thus the long-standing lack of detection of CDM, which, after decades, slowly seems to be turning the search for CDM into a dead end.

One such promising dark matter candidate is in the form of ultra-light particles, which in some sense form the opposite of the WIMP paradigm by positing that dark matter does not behave as individual, heavy particles, but rather as (quantum) waves. This dissertation specifically focuses on fuzzy dark matter (FDM), e. g. in the form of axion-like particles (ALPs), as an example of (quantum-)wave dark matter, which behaves similarly to CDM at large scales, but with different properties at galactic scales due to wave effects arising from the extremely light particles' large de Broglie wavelengths. The goal of the project has been to implement and perform large cosmological simulations of FDM (with and without baryons) using the moving-mesh code AREPO (Springel 2010; Weinberger, Springel, and Pakmor 2020) in order to study non-linear structure formation in such scenarios, and obtain predictions about the behavior of small-scale structure and how it differs from the CDM model.

1.1 The ΛCDM model of cosmology

As mentioned above, the ΛCDM model has become the standard model of cosmology due to its great success in the description of many independent observations probing a large range of length scales and times in cosmic history. Due to its status as the standard, details about the model and the general cosmological paradigm are presented in many review articles, such as Particle Data Group: Zyla et al. (2020), where most of the statements made in this section can be found.

1.1.1 Friedmann–Lemaître–Robertson–Walker cosmology

ΛCDM is a concrete model within the general framework of cosmology, following the *cosmological principle* which assumes that (at least on large enough scales), the distribution of matter in the universe is homogenous and isotropic. This assumption

directly implies that the space-time metric has the form of the FLRW metric

$$\mathrm{d}s^2 = -c^2\,\mathrm{d}t^2 + a(t)^2\left(\frac{1}{1-kr^2}\,\mathrm{d}r^2 + r^2\left(\mathrm{d}\theta^2 + \sin(\theta)^2\,\mathrm{d}\varphi^2\right)\right), \tag{1.1}$$

where t, r, θ and φ are the (spherical) space-time coordinates, c is the speed of light (in vacuum), k is the spatial curvature, and $a(t)$ is the scale factor, which determines the expansion or contraction of the universe over time. By convention, $a(t_0) = 1$, where t_0 is the present age of the universe. The scale factor $a(t)$ is related to the redshift z,

$$z = \frac{1}{a} - 1\,, \tag{1.2}$$

since the observed wavelength of light λ_{obs} is reduced (in the visible part of the electromagnetic spectrum: shifted towards red) compared to the wavelength λ_{i} at the time of emission, behaving as $\lambda_{\mathrm{obs}}/\lambda_{\mathrm{i}} = 1 + z$ as a distant emitter appears to recede from the observer due to the expansion of space. The sign of the curvature determines the topology of space-time:

- $k > 0$: closed (hyperspherical),
- $k = 0$: flat,
- $k < 0$: open (hyperbolical).

What does this imply for the distribution of matter and energy in the universe? Homogeneity and isotropy in the energy–momentum tensor $(T^{\mu\nu})$ impose the form

$$(T^{\mu\nu}) = \begin{pmatrix} \rho_T(t)c^2 & 0 & 0 & 0 \\ 0 & p_T(t) & 0 & 0 \\ 0 & 0 & p_T(t) & 0 \\ 0 & 0 & 0 & p_T(t) \end{pmatrix}, \tag{1.3}$$

which describes a perfect fluid with mass density ρ_T and pressure p_T.

Now, this is clearly insufficient to describe the formation of structure as we can observe it – in this rendition, the density is constant throughout space! As hinted at above, this simplest description is only valid when taking a coarse view of the universe, on very large cosmological scales. However, it is nevertheless illustrative to understand global properties of the universe, such as its age, expansion history, or curvature. Any applications aiming to study structure formation must work with a

more general, perturbed form of eq. (1.1), which can take into account local variations in the density field.

Applying the FLRW metric and the energy–momentum for a perfect fluid, eqs. (1.1) and (1.3), in Einstein's field equations,

$$G_{\mu\nu} + \Lambda g_{\mu\nu} = \frac{8\pi G}{c^4} T_{\mu\nu} \,, \tag{1.4}$$

where $G_{\mu\nu} = R_{\mu\nu} - \frac{1}{2} g_{\mu\nu} R$ is the Einstein tensor (using the Ricci tensor ($R_{\mu\nu}$) and the Ricci scalar $R = g^{\mu\nu} R_{\mu\nu}$), Λ is a cosmological constant, ($g_{\mu\nu}$) is the metric tensor, and G is the gravitational constant, results in the *Friedmann equations*

$$\frac{\dot{a}(t)^2 + kc^2}{a(t)^2} = \frac{8\pi G}{3} \rho_T(t) + \frac{\Lambda c^2}{3} \tag{1.5}$$

$$\frac{\ddot{a}(t)}{a(t)} = -\frac{4\pi G}{3} \left(\rho_T(t) + \frac{3 p_T(t)}{c^2} \right) + \frac{\Lambda c^2}{3} \,. \tag{1.6}$$

Due to the assumption of homogeneity and isotropy, the solution only depends on the time variable t, while being constant across space. Introducing the Hubble parameter $H(t) = \dot{a}(t)/a(t)$ and subsuming the cosmological constant as an additional component of a "total" density ρ and pressure p,

$$\rho(t) = \rho_T(t) \underbrace{+ \frac{\Lambda c^2}{8\pi G}}_{=\rho_\Lambda} \tag{1.7}$$

$$p(t) = p_T(t) \underbrace{- \frac{\Lambda c^4}{8\pi G}}_{=p_\Lambda} \tag{1.8}$$

yields the simpler versions

$$H(t)^2 = \frac{8\pi G}{3} \rho(t) - \frac{kc^2}{a(t)^2} \tag{1.9}$$

$$\dot{H}(t) + H(t)^2 = -\frac{4\pi G}{3} \left(\rho(t) + \frac{3 p(t)}{c^2} \right) \tag{1.10}$$

of eqs. (1.5) and (1.6). Additionally, taking the time derivative of eq. (1.9) and using

eq. (1.10) to eliminate $\dot{H}(t)$ results in

$$\dot{\rho}(t) = -3H(t)\left(\rho(t) + \frac{p(t)}{c^2}\right), \tag{1.11}$$

which is also called the "third Friedmann equation", energy conservation equation, continuity equation, or fluid equation.[2]

Further, the density components ρ_i are commonly expressed using the density parameters Ω_i, which indicate their values as a fraction of the critical density

$$\rho_{\text{crit}}(t) = \frac{3H(t)^2}{8\pi G} \tag{1.12}$$

such that

$$\Omega_i(t) = \frac{\rho_i(t)}{\rho_{\text{crit}}(t)} = \frac{8\pi G\rho_i(t)}{3H(t)^2}\,. \tag{1.13}$$

The physical significance of the critical density ρ_{crit} is that a universe with this density, $\rho = \rho_{\text{crit}}$, is spatially flat, i. e. $k = 0$. Correspondingly, it is useful to consider the density parameter Ω of all components, since this is what determines the spatial curvature:

$$\Omega(t) = \Omega_{\text{r}}(t) + \Omega_{\text{m}}(t) + \Omega_\Lambda(t) \begin{cases} < 1 \Rightarrow k < 0 \\ = 1 \Rightarrow k = 0 \\ > 1 \Rightarrow k > 0 \end{cases} \tag{1.14}$$

The present values of these parameters are indicated as $\Omega_{i0} = \Omega_i(t_0)$. One can even go so far as to define a "density" and "pressure" associated with the curvature,

$$\rho_k(t) = -\frac{3kc^2}{8\pi G a(t)^2} \qquad \Omega_k(t) = \frac{\rho_k(t)}{\rho_{\text{crit}}} \tag{1.15}$$

$$p_k(t) = \frac{kc^4}{8\pi G a(t)^2}\,, \tag{1.16}$$

which formally look like ordinary contributions towards density and pressure in the Friedmann equations, although this can be misleading in the sense that the curvature is rather a consequence of the energy content of the universe, not a contribution towards it. Dividing by $\rho_{\text{crit}}(t)$ in eq. (1.9), another form of the first Friedmann

[2] This is a direct consequence of the fact that the energy–momentum tensor satisfies $\nabla^\mu T_{\mu\nu} = 0$, i. e. its covariant divergence vanishes.

equation is

$$1 = \Omega(t) + \Omega_k(t) = \Omega_\mathrm{r}(t) + \Omega_\mathrm{m}(t) + \Omega_k(t) + \Omega_\Lambda(t) \,, \tag{1.17}$$

i. e. the density parameters always sum to one![3]

Different components of the total density are characterized by their equations of state

$$p_i = w_i \rho_i c^2 \,, \tag{1.18}$$

which links pressure and density. The "original" density and pressure ρ_T and p_T (i. e. the contributions to the energy–momentum tensor $(T^{\mu\nu})$) can be split into (relativistic) radiation ρ_r and (non-relativistic) matter ρ_m, which have the following equations of state:

$$p_\mathrm{r} = \frac{1}{3}\rho_\mathrm{r} c^2 \qquad \left(w_\mathrm{r} = \frac{1}{3}\right) \tag{1.19}$$

$$p_\mathrm{m} = 0 \qquad \left(w_\mathrm{m} = 0\right) \tag{1.20}$$

A pressureless perfect fluid, which is used to model non-relativistic matter by approximating $p \ll \rho$ as $p = 0$, is also referred to as "dust". Analogously, the cosmological constant and curvature can also be viewed as fluid components using eqs. (1.7) and (1.8), with corresponding equations of state:

$$p_\Lambda = -\rho_\Lambda c^2 \qquad \left(w_\Lambda = -1\right) \tag{1.21}$$

$$p_k = -\frac{1}{3}\rho_k c^2 \qquad \left(w_k = -\frac{1}{3}\right) \tag{1.22}$$

Using eq. (1.11), the equation of state can instead be translated into the corre-

[3] Another way to view this is that the "curvature component" always combines with the total density to give the critical density, $\rho_\mathrm{crit} = \rho + \rho_k$. Similarly defining $p_\mathrm{crit} = p + p_k$, the Friedmann equations are

$$H(t)^2 = \frac{8\pi G}{3}\rho_\mathrm{crit}(t)$$

$$\dot{H}(t) + H(t)^2 = -\frac{4\pi G}{3}\left(\rho_\mathrm{crit}(t) + \frac{3p_\mathrm{crit}(t)}{c^2}\right).$$

sponding density component's scaling with the scale factor $a(t)$, giving[4]

$$\rho_i(t) = \rho_{i0}\, a(t)^{-3(1+w_i)} \qquad \Omega_i(t) = \Omega_{i0}\left(\frac{H_0}{H(t)}\right)^2 a(t)^{-3(1+w_i)} \tag{1.23}$$

and thus

$$\rho_\mathrm{r}(t) = \rho_{\mathrm{r}0} a(t)^{-4} \tag{1.24}$$

$$\rho_\mathrm{m}(t) = \rho_{\mathrm{m}0} a(t)^{-3} \tag{1.25}$$

$$\rho_k(t) = \rho_{k0} a(t)^{-2} \tag{1.26}$$

$$\rho_\Lambda(t) = \rho_\Lambda a(t)^0\,, \tag{1.27}$$

with $\rho_{i0} = \rho_i(t_0)$. Strictly speaking, eq. (1.11) only holds for the *total* energy–momentum tensor ($T^{\mu\nu}$), a priori making the relations above in eqs. (1.23) to (1.27) valid only for single-component universes (although eq. (1.26) is always true). If there are multiple components, such as $\rho_T = \rho_\mathrm{r} + \rho_\mathrm{m}$, this requires the additional assumption that a conservation equation like eq. (1.11) holds *separately* for each component, i. e. there is no interaction between components and energy–momentum is separately conserved for each:

$$\dot{\rho}_i(t) = -3H(t)\left(\rho_i(t) + \frac{p_i(t)}{c^2}\right), \tag{1.28}$$

With some slight foreshadowing, while this is valid in the late universe, where different kinds of matter and radiation have "decoupled" from each other, all known particles were in thermal equilibrium at the very high temperatures in the early universe. As a consequence, for the time of Big Bang nucleosynthesis (BBN) and earlier, interactions between and decoupling of different species must be taken into account throughout the evolution.

Finally, using eqs. (1.23) to (1.27) and dividing by the values at $t = t_0$ in eq. (1.9), the first Friedmann equation, can be expressed as

$$\frac{H(t)^2}{H_0^2} = \Omega_{\mathrm{r}0} a(t)^{-4} + \Omega_{\mathrm{m}0} a(t)^{-3} + \Omega_{k0} a(t)^{-2} + \Omega_{\Lambda 0} \tag{1.29}$$

with the present-day Hubble parameter ("Hubble constant") $H_0 = H(t_0)$. H_0 is

[4] Since ρ_Λ and ρ_k also obey an equation of state like eq. (1.18), their behavior is consistent with energy conservation as in eq. (1.11), although their evolution with $a(t)$ already follows directly from the definitions in eqs. (1.7) and (1.15).

often expressed in terms of the quantity h, with

$$H_0 = 100\,h\,\mathrm{km\,s^{-1}\,Mpc^{-1}}\ . \tag{1.30}$$

1.1.2 Definition and parameters of ΛCDM

As a concrete form of an FLRW cosmology, ΛCDM postulates that the universe and all of the matter and energy within it are composed of the following properties and constituents:

Flatness: ΛCDM sets $k = 0$, or, equivalently, $\Omega_0 = 1$, modeling a flat universe without spatial curvature. The effect of allowing for non-zero curvature has been investigated using observational data, finding that this assumption appears to be rather accurate.

Topology: Although not often discussed, the curvature only describes the local geometry of the universe. Globally, even for the same curvature, it could be finite or infinite, and have a non-trivial topology (periodic or multiply connected). ΛCDM assumes the simplest, trivial topology.

Cosmological constant Λ: As seen in eq. (1.27), the cosmological constant can be viewed as a contribution to the total energy density which does not change with the scale factor. Since all the other components in the model decay as the universe expands, it will be Λ-dominated ($\Omega_\Lambda \approx 1$) at late times, which leads to an exponentially accelerating expansion:

$$a(t) \propto \exp\left(H_0\sqrt{\Omega_{\Lambda 0}}\,t\right)\ . \tag{1.31}$$

As Λ contributes the majority to the universe's energy budget today according to observations ($\Omega_{\Lambda 0} \approx 0.7$), the universe has recently started to enter the Λ-dominated phase of accelerated expansion. Different variants of "dark energy" beyond the simplest, in the form of the cosmological constant, have been explored, e. g. in the form of a different equation of state with $w \neq -1$. However, so far, measurements show very good agreement with Λ and $w = -1$.

(Cold) dark matter: The motivation and evidence for dark matter is explained in more detail in section 1.2. Current measurements set the contribution of dark matter at $\Omega_\mathrm{m} \approx 0.25$, or roughly five times the amount of "known", baryonic matter. The fact that dark matter was non-relativistic (and thus behaving as

"matter") even relatively early during the evolution of the universe implies that its velocity dispersion is very low at later times, making it "cold". There is no parameter setting the velocity dispersion of dark matter in ΛCDM, leaving two potential interpretations for the meaning of "cold" in CDM:

1. ΛCDM does not specify the exact meaning of CDM, and encompasses any kind of dark matter which is "cold enough" to have been non-relativistic early on.
2. Assuming the extreme limiting case of zero velocity dispersion.

The first is a rather broad and vague definition, while the latter is an idealized assumption. In both cases, however, the definition is phenomenological: ΛCDM does not make statements about the nature of dark matter apart from its lack of interactions beyond gravity and its "cold" velocity dispersion.

For the purposes of structure formation on smaller scales, such as galaxies, the velocity dispersion does make a difference, though (see section 1.5.1), adding at least one parameter like the mass of a dark matter particle. Whether this is still considered to be an "implementation detail" (like the masses and interactions of standard model particles for baryons) in the larger ΛCDM framework or a departure from the simplest six-parameter model (see below) is ultimately a philosophical question.

Baryonic matter: This is the matter which (combined with photons, i. e. radiation) actually represents the known laws of physics, and the matter every directly observable object is made of: atoms, ions and molecules. (It should be noted that "baryonic matter" in the context of cosmology, somewhat sloppily, refers to all non-relativistic matter as described by the standard model of particle physics – excluding neutrinos –, including leptons such as the electron.) Although baryons are non-relativistic and, in that sense, behave the same as dark matter on large scales, their interactions with photons until the time of decoupling produce features in the form of baryon acoustic oscillations (BAO) which do not occur with matter not coupled to photons at this time. When it comes to the details of later structure formation, particularly galaxy formation, baryonic physics beyond mere gravitational attraction (as in the case of dark matter) also plays an extremely important role. Assuming independent conservation of components, as in eq. (1.28), and the dust approximation ($p = 0$) is insufficient in these cases.

Photons: Just as for baryonic matter, relativistic matter, mainly in the form of photons, is known to exist and can directly be measured. In fact, this is how almost all information about the universe's past and evolution reaches us, through either the cosmic microwave background (CMB) or observations of objects in the electromagnetic spectrum. Since the energy density of radiation has the strongest dependence on the scale factor, see eq. (1.24), the early universe was radiation-dominated, whereas it only contributes a negligible amount in the later phases.

Neutrinos: The neutrinos are an unusual case: Although it has been determined that they have non-zero masses (Nobel Prize 2015), the actual values must be very small, making them relativistic, and thus contributing to "radiation", in the early universe. However, given the bounds on the neutrino masses, at least some neutrino species have become non-relativistic ("matter") by the present day, making them a form of "hot dark matter". While the details depend on the values of the neutrino masses, these are not considered to be parameters of ΛCDM. Alternately, one could consider the neutrinos to be massless in ΛCDM (just as, technically, they are in the standard model of particle physics).

Density power spectrum: As explained above, a pure FLRW cosmology can only describe the behavior of space-time on the largest scales. However, such a homogenous and isotropic universe by definition has no structure. In order to account for structure, even on the largest observable scales in the CMB or the clustering of matter, local variations (anisotropies and inhomogeneities) in the density field must be introduced. This raises the question of how matter and energy were distributed in the very early universe. Although not an explicit assumption, ΛCDM assumes that these early density variations are consistent with inflation, which addresses a number of issues (such as the horizon, flatness, and magnetic monopole problems). Specifically, the primordial density fluctuations are assumed to be Gaussian and adiabatic, and can be described by a power spectrum of the form

$$P(k) = A_\mathrm{s} k^{n_\mathrm{s}} \,. \tag{1.32}$$

Here, A_s sets the amplitude of fluctuations, i. e. how strongly matter is clustered in general, and n_s determines the scale-dependence of this clustering. Measurements show that n_s is slightly smaller than 1, which is consistent with nearly scale-invariant potential fluctuations predicted by inflationary models.

In particular, the eponymous first two components (Λ and CDM) are the main ingredients resulting in the model's non-trivial predictions, whereas even the simplest observations demonstrate the presence of baryonic matter and radiation, which must thus be included in any cosmological model. However, the cosmic density of these components is not known a priori and can vary depending on the model, which also has an impact on the curvature and geometry (flat, open, or closed) of the universe.

Given the above ingredients of ΛCDM, the following question arises: What quantities need to be known to make predictions using this model? Since the Friedmann equations are a first-order differential equation for the Hubble parameter $H(t)$, with the evolution of the different fluid components, eq. (1.23), all determined once their values at a single point in time is known, the set of density parameters, together with the Hubble parameter, at one point in time is sufficient to solve the Friedmann equations. Two additional parameters come in to describe the initial density fluctuations. To summarize, this results in the following set of possible parameters:

- Curvature k;
- Density parameters for CDM, baryonic matter, photons, neutrinos, and the cosmological constant: $\Omega_{\mathrm{c}0}$, $\Omega_{\mathrm{b}0}$, $\Omega_{\gamma 0}$, $\Omega_{\nu 0}$, $\Omega_{\Lambda 0}$;
- Hubble "constant" H_0;
- Power spectrum normalization and shape: A_{s} and n_{s}.

Further, the optical depth at reionization τ (or, equivalently, the time of reionization z_{ion}) is usually included as an additional parameter, because it is needed to interpret CMB measurements. Additionally, the first Friedmann equation can be used to eliminate one of the parameters of the set $\{H_0, \Omega_{i0}\}$, and the curvature is fixed at $k = 0$, leaving us with $7 + 1$ parameters. However, the components of radiation, although technically parameters since their values must also be measured, are usually not included because $\Omega_{\gamma 0}$ can be measured directly from the temperature of the CMB, which also determines $\Omega_{\nu 0}$ from standard particle physics interactions assuming a value for the sum of the neutrino masses. This finally results in a set of six independent parameters, as used by e. g. Planck and the Wilkinson Microwave Anisotropy Probe (WMAP) (Planck Collaboration: Aghanim, Akrami, Ashdown, et al. 2020; Hinshaw et al. 2013) to determine cosmological parameters from the CMB.

The details of inflation and BBN, although consistent with ΛCDM, even to the extent that ΛCDM replicates some of their predictions as assumptions or is designed to accommodate them, are not treated as a part of ΛCDM itself. In that sense, ΛCDM

should be viewed more as a general framework, or, as noted before, a phenomenological description, rather than a complete model of the universe. The boundaries of what is considered part of ΛCDM and what is separate, including the question of which parameters are free and which are "fixed", to be determined externally within the frameworks of separate theories and simply adopted as given, can be fuzzy. Another point of view could be that ΛCDM does not make statements about the early universe at the time of BBN and earlier, providing only a description of the subsequent evolution.

In this early regime, the "agnostic" view on the nature of (cold) dark matter also becomes ill-defined. At extremely high temperatures in the very early universe, the micro-physical details of the nature of dark matter are wildly model-dependent. Even for a very common class of dark matter models, WIMPs with a thermal freeze-out, the fact that these particles would have been coupled to all other particle species at some point introduces a substantial dependence on the WIMP's mass and interactions, particularly concerning decoupling and the transition from relativistic "radiation" to non-relativistic "matter". Other models, which might not even depend on elementary particles, like primordial black holes (PBHs), have completely different implications at these times.

1.2 Observational evidence for dark matter

As explained above, dark matter is a very important part of the current cosmological model, being much more abundant than visible matter. It is worthwhile to take a step back and consider the definition of dark matter. The degree to which ΛCDM is well-defined can depend on the specific context and application of the model, as demonstrated in section 1.1. Most generally, "dark matter" is simply a term denoting a missing piece which is required to explain the formation of structures as they can be observed today. It is necessary because simply assuming the entire matter content to be baryonic is wildly inconsistent with observations due to the interactions with radiation, from features in the CMB (the BAO) to the formation of structure (which proceeds too slowly without dark matter) and the properties of galaxies.

As will be outlined below, observations have established the following properties of dark matter:

- Behaves like (non-relativistic) matter, with the characteristic density scaling of eq. (1.25), $\rho \propto a^{-3}$;

- (Reasonably) cold, i. e. it was also non-relativistic during the early universe, before the epoch of recombination, and allows structures such as dwarf galaxies to form;

- Negligible interactions with baryons beyond gravity.

However, while the obvious conclusion would be that there is in fact an additional, invisible matter component which simply does not noticeably interact with the matter we can observe (except gravitationally), it does not strictly imply that the physical mechanism responsible for the properties of this observed phenomenon must be a specific kind of object or particle, or even matter as we know it at all. As long as they are compatible with the effective properties above, and do not contradict observations, completely different explanations, such as a modification of established laws of physics (like general relativity; see section 1.5.4), are just as valid. In this sense, "dark matter" is merely a placeholder for an unknown, essential ingredient to cosmic structure formation.[5]

1.2.1 The cosmic microwave background

Measurements of the CMB, e. g. using the Planck and WMAP satellites (Planck Collaboration: Aghanim, Akrami, Ashdown, et al. 2020; Hinshaw et al. 2013), have been immensely successful in gaining insight about cosmology and the evolution of the universe. The most precise determinations of ΛCDM's cosmological parameters today are obtained using these measurements. Different parameters (such as the various density components) have different impacts on the angular temperature power spectrum of the CMB. In particular, changes in the dark matter and baryon densities have different effects on the BAO peaks, allowing one to disentangle the two matter components. Without a dark matter component, it would not be possible to interpret the CMB measurements.

Figure 1.1 shows several measurements of the angular power spectrum of both CMB temperature and polarization data, together with the fits to ΛCDM. While the temperature power spectrum is already enough to mandate the existence of

[5] Somewhat puzzlingly, the debate about whether dark matter is actually "matter" made up of some kind of particles or physical objects, tends to become quite heated at times. Faced with a phenomenon which behaves like matter, and which is of the kind for which there is even "precedent" among the known types of matter, such as neutrinos (which, in the form of hot dark matter, are a tiny component of the dark matter today), it appears only natural to infer that it *is* in fact a form of matter.

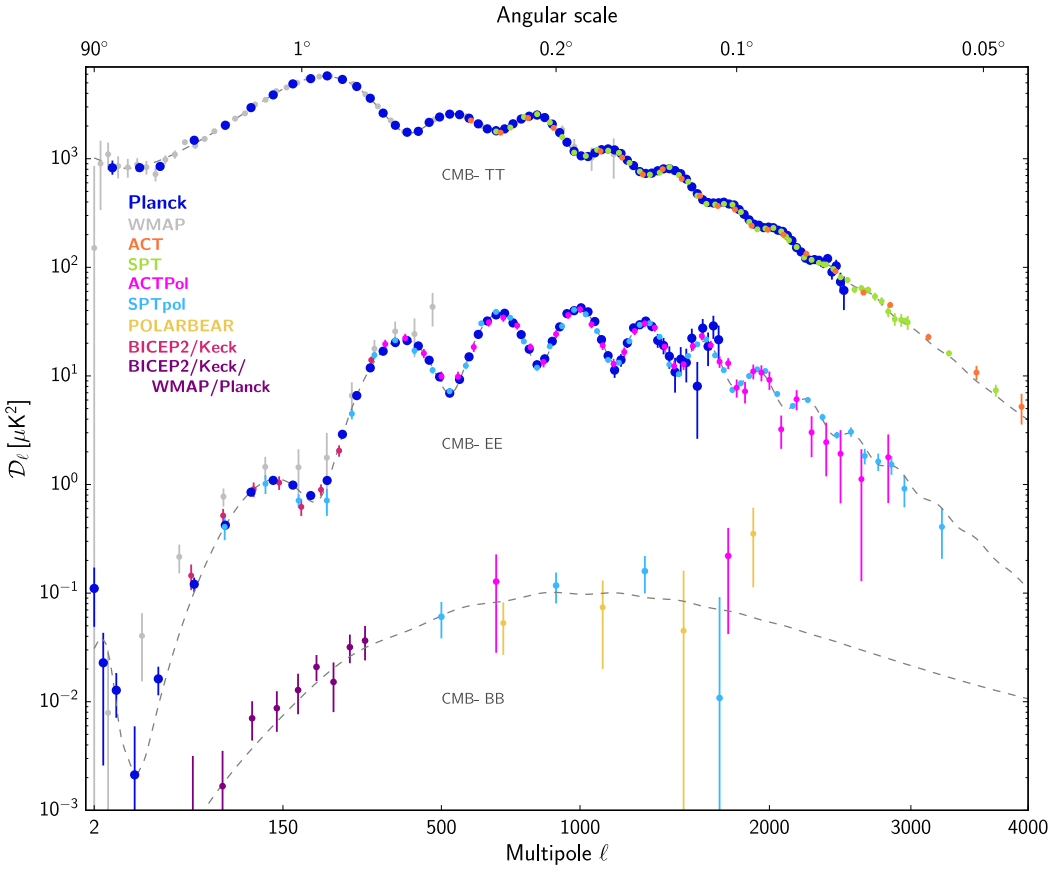

Figure 1.1: Angular power spectrum measurements for temperature (TT), *E*-mode (EE) and *B*-mode (BB) polarization signals compiled from various observations. Dashed lines show the ΛCDM fit to the Planck data (taken from Planck Collaboration: Aghanim, Akrami, Arroja, et al. 2020).

dark matter, fig. 1.1 demonstrates the precision of available data, and the excellent agreement of ΛCDM with the data even at this level of precision.

1.2.2 Large-scale structure and clustering of matter

The large-scale clustering of matter has been observed for a long time via galaxy surveys, such as the second Center for Astrophysics Redshift Survey (CfA2 survey) (Geller and Huchra 1989), 2-degree-Field Galaxy Redshift Survey (2dFGRS) (Peacock et al. 2001), and Sloan Digital Sky Survey (SDSS) (Blanton et al. 2017). The distribution of galaxies matches the "cosmic web" seen in numerical simulations of structure formation with collisionless CDM. Baryonic matter behaves differently in this regard because it is influenced by its interactions with photons until decoupling

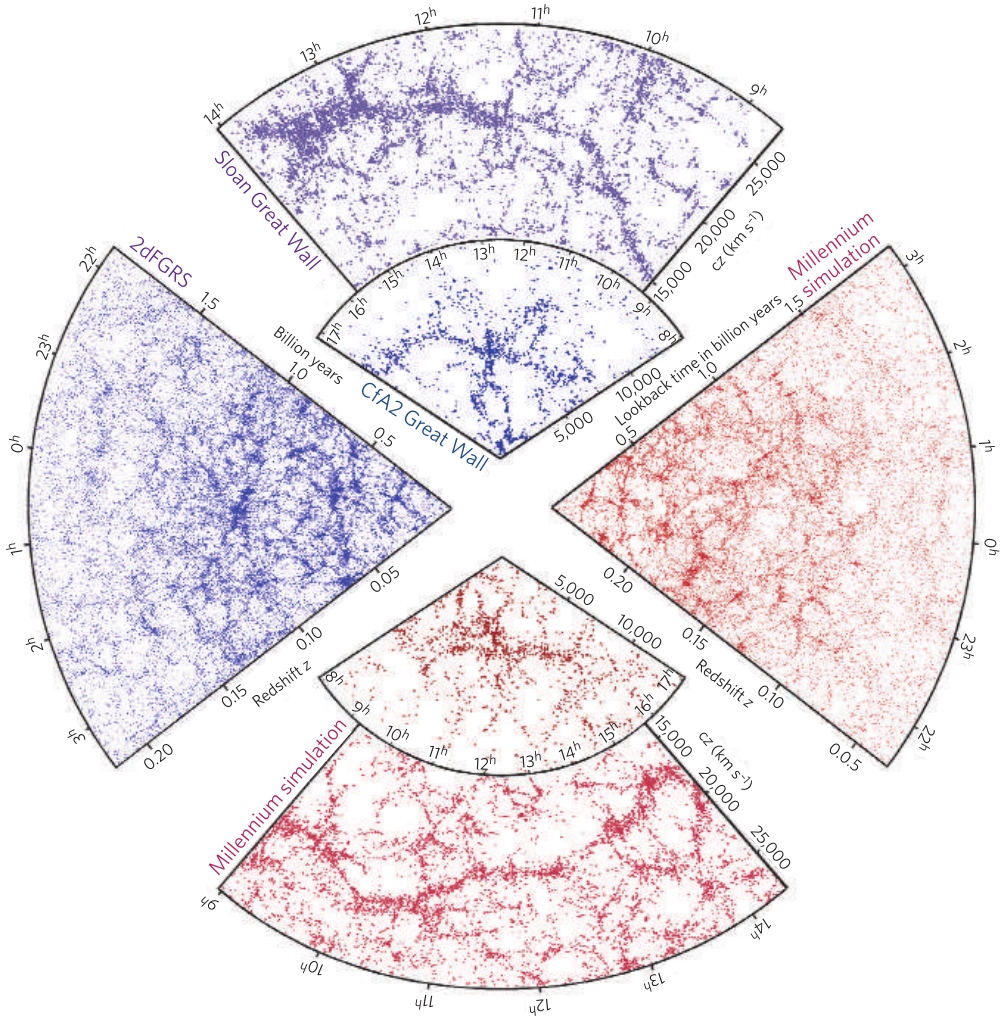

Figure 1.2: The distribution of galaxies in spectroscopic redshift surveys, such as the CfA2 survey, 2dFGRS, and SDSS, compared to mock catalogs from the numerical "Millennium simulation" (taken from Springel, Frenk, and White 2006).

from the CMB. Before decoupling, baryonic density perturbations cannot condense and instead oscillate due to these interactions (the BAO mentioned above). Dark matter, however, is unaffected by the presence of radiation and its perturbations can grow, providing a "potential well" for baryons to later fall into. This process of non-linear structure formation, which can be studied using numerical simulations, is consistent with the observed properties of structure, as illustrated in fig. 1.2.

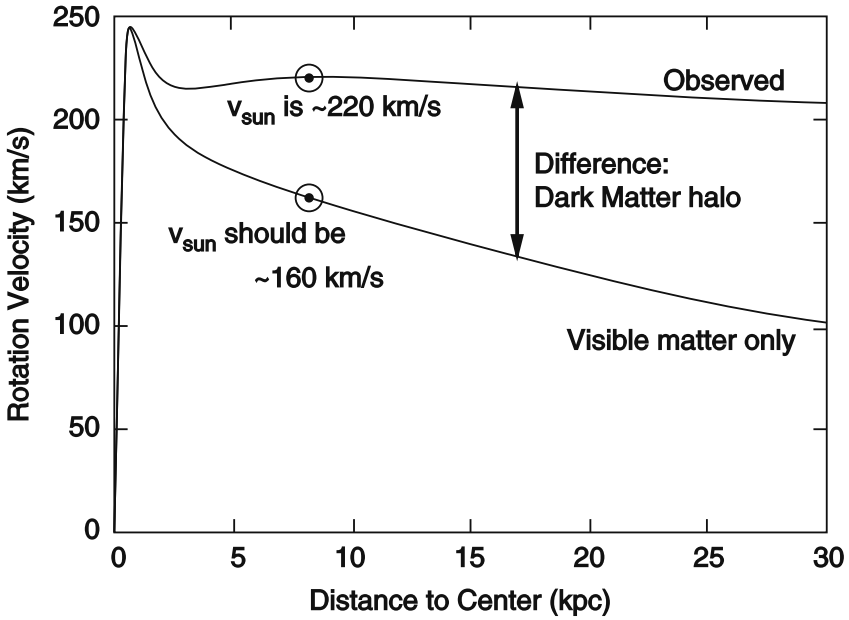

Figure 1.3: Expected and observed rotation curve of the Milky Way galaxy (taken from Schneider 2015).

1.2.3 Velocity dispersions and galactic rotation curves

Historically, dynamical properties of luminous matter, particularly the velocity dispersions and orbital velocities of stars, were the first cases prompting the proposition that many astronomical objects are much more massive than their visible components suggest, indicating the presence of an additional "dark" component (Kapteyn 1922; Lundmark 1930). On larger scales, but with the same underlying principle, observations of the velocities of galaxies within the Coma cluster (Zwicky 1933) led to similar results, finding galaxies to have much larger velocities than expected from the overall mass in the vicinity.

Later, more detailed and accurate spectroscopic measurements of the rotational velocities of matter in spiral galaxies (Rubin and Ford 1970), e. g. using gas emission line regions, allowed increasingly accurate mass determinations of galaxies, and thus of their "mass-to-light ratios", which indicated the presence of mass several times higher than inferred from luminous stars and gas. For example, the rotation curve of a galaxy may look as exemplified in fig. 1.3. The velocity in a circular orbit at a distance r from the center is

$$v = \sqrt{\frac{GM(r)}{r}}\,, \tag{1.33}$$

where $M(r)$ is the mass contained within a sphere of radius r. Since most of the visible mass is concentrated within a galaxy's center, one would expect $M(r)$ to be roughly constant for large distances, which would result in the "Keplerian" dependence of orbital velocities,

$$v \propto \frac{1}{\sqrt{r}} \tag{1.34}$$

for large r. However, as shown in fig. 1.3, the observed orbital velocity distribution as a function of distance (rotation curve) is rather different: Away from the center, it becomes flat, remaining constant even at large distances and towards the edge of the galaxy.

Since there is no luminous matter that can be detected in the outskirts of galaxies, this effect cannot be explained from observations in the electromagnetic spectrum under the assumption that all matter is baryonic. Instead, a large fraction of the total mass seems to be distributed in a halo throughout the galaxy, but in a form which does not emit light or any other kind of radiation.

1.2.4 Galaxy clusters

Galaxy clusters are a very important probe of dark matter because their masses can be estimated in several different ways, confirming the presence of dark matter through multiple independent approaches. As mentioned previously in section 1.2.3, one of the early signs of dark matter came from the measurement of galaxy cluster masses by applying the virial theorem to the velocity dispersions of their constituent galaxies (Zwicky 1933). A second method makes use of the hot gas in the intracluster medium (ICM). Due to the enormous masses of galaxy clusters, the matter contained within them reaches very high velocity dispersions, and thus, high temperatures. At these temperatures, the gas can actually be observed directly via its emission of X-rays. In turn, measuring the emitted X-ray spectrum makes it possible to recover the gas mass, temperature, and pressure, and finally, the cluster's total mass distribution. Once again, the baryon fraction of the total mass reveals that baryons do not make up the cluster alone, with the majority of matter "missing" from direct observations. Finally, measurements of gravitational lensing (discussed more in section 1.2.5) can be used as yet another estimate of galaxy cluster masses, with all three methods generally in reasonable agreement.

1.2.5 Gravitational lensing

Gravitational lensing is a very powerful tool since, in principle, it enables forming direct conclusions about the distribution of mass. Crucially, as a purely gravitational

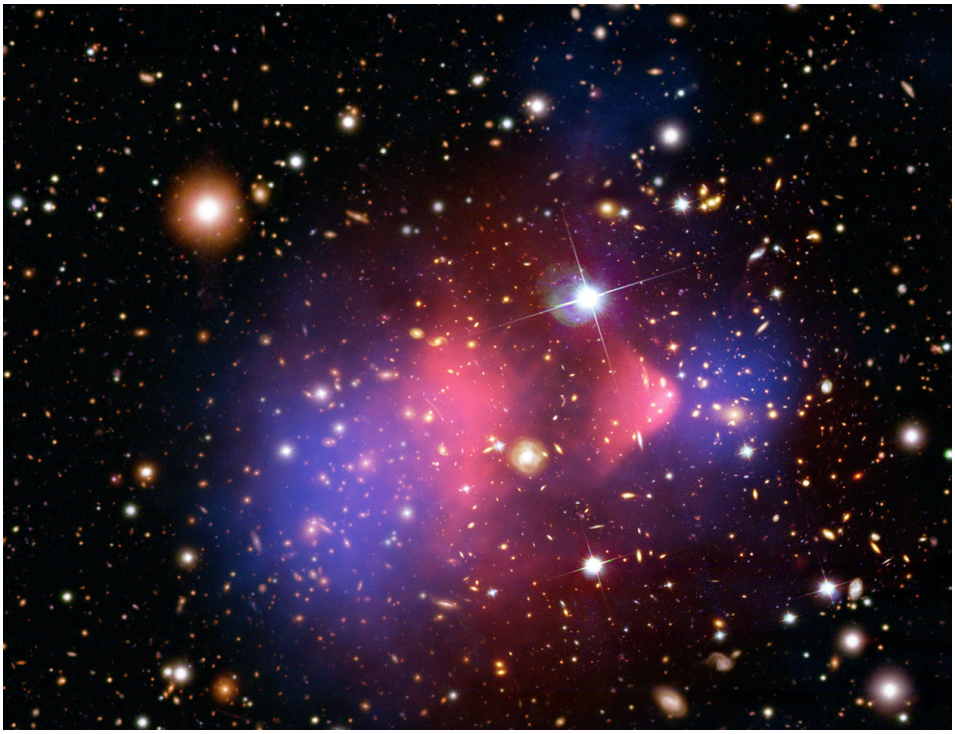

Figure 1.4: Combined gravitational lensing, X-ray and optical map of the Bullet Cluster (X-ray: NASA/CXC/CfA/M.Markevitch et al.; Optical: NASA/STScI; Magellan/U.Arizona/D.Clowe et al.; Lensing Map: NASA/STScI; ESO WFI; Magellan/U.Arizona/D.Clowe et al. 2016). Blue regions indicate the concentration of total mass as inferred from gravitational lensing, while red regions show the X-ray emissions of cluster gas.

effect, no matter can "hide" from it, so it will always include the full aggregation of all matter components, yet remain undisturbed by interference from other physical processes. A particularly striking example has been observed in the form of the so-called Bullet Cluster (see fig. 1.4), which is the result of a recent merger. The simultaneous measurement of X-ray flux and the impact of gravitational lensing provide the means to obtain separate maps of the baryonic and total mass distributions. Figure 1.4 demonstrates that the X-ray-emitting cluster gas (in red), which exhibits a bow shock as the gas from the progenitors collides, is lagging behind the bulk of the matter (in blue), which appears to have simply passed through each other unaffected. Indeed, the matter distribution coincides with the cluster galaxies, which have similarly passed each other in a collisionless manner (as expected). However, the galaxies only make up a tiny fraction of a cluster's total mass, implying the presence of another "invisible", pressure- and collisionless matter component.

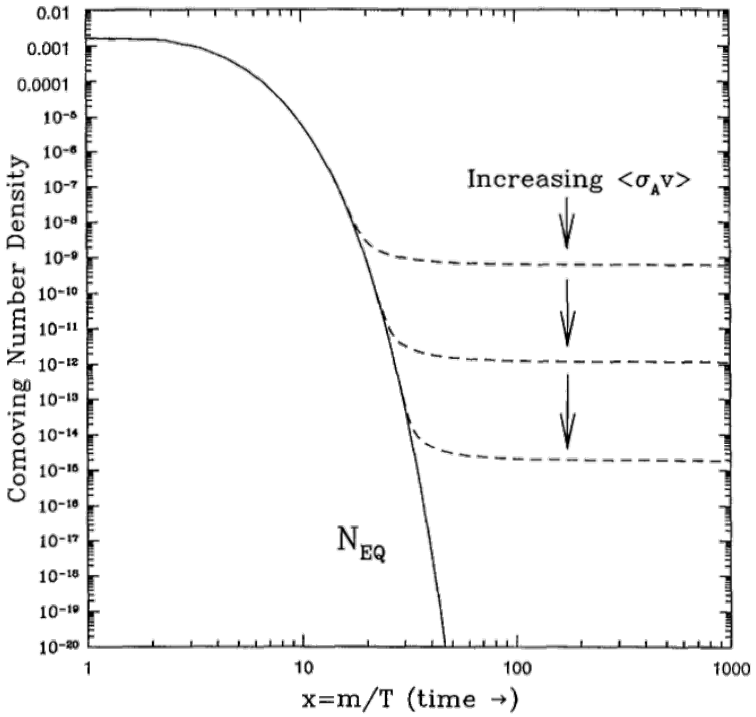

Figure 1.5: Freeze-out mechanism for WIMPs (taken from Jungman, Kamionkowski, and Griest 1996).

1.3 Realizations of cold dark matter

CDM is a model concerned with what can be determined from astronomical observations. Naturally, these observations are not sensitive to many of the micro-physical properties that dark matter itself may have. As these are the principal candidates which have been investigated using experiments or astronomical observations, I will give a brief overview of their motivations and properties.

1.3.1 Weakly interacting massive particles (WIMPs)

A central motivation for this class of dark matter candidates is the so-called "WIMP miracle": the observation that weakly interacting particles of masses around 1 TeV – which would be within reach to be detected at CERN's Large Hadron Collider (LHC) – surviving as a thermal relic from the early universe "automatically" result in the correct density of dark matter observed today in the universe.

However, while SUSY is appealing because it would address several problems with the standard model in addition to providing a dark matter candidate, no super-

symmetric particles have been detected in experiments so far, and data from particle accelerators such as the LHC at CERN are increasingly constraining the parameter space for SUSY. This is a significant detriment to the viability of SUSY, since some of the problems it was hoped to address are cases of parameter fine-tuning in the standard model, which is rendered moot if the parameters of SUSY must themselves be fine-tuned, particularly for supersymmetric particle masses > 1 TeV.

1.3.2 Axions

Axions originate from a problem purely in the realm of particle physics, the *strong CP problem*. This problem can be stated as follows: The strong interaction, in the form of quantum chromodynamics (QCD), in principle allows for a term which violates symmetry under the CP operation, which is a combined conjugation of charge and spatial parity. However, in practice, experiments have determined that QCD does not violate CP symmetry, or if it does, that the degree to which it does and the associated parameter θ_{QCD} in the CP-violating term are extremely tiny. This is puzzling, because an "accidental" symmetry or "fine-tuning" like this usually betrays the presence of a deeper physical mechanism responsible for its origin.

Peccei and Quinn (1977) proposed a solution in the form of the Peccei–Quinn mechanism. This mechanism effectively promotes the parameter θ_{QCD} to a new pseudo-scalar particle, the axion, which is the pseudo-Nambu–Goldstone boson of a spontaneously broken "Peccei–Quinn" $\mathrm{U}(1)$ symmetry. The presence of the axion and the specific form of its interactions arising from the Peccei–Quinn mechanism naturally cancel the CP-violating term of QCD.

While postulated to solve a problem in particle physics, the motivation for axions is strengthened significantly by the fact that they also serve as a natural candidate for CDM. Through the so-called misalignment mechanism, cold axion dark matter could be produced in the early universe, which would have a particle mass in the μeV to meV range (Marsh 2016).

1.3.3 Primordial black holes and other compact objects

Although dark matter does not seem to interact with visible matter (except gravitationally), one could in principle imagine that it is still composed of "ordinary", baryonic matter, as long as this matter has assembled in a form which would make it very difficult to detect in an astrophysical context. Neutron stars, "rogue" planets or even just very faint stars, such as white or red dwarfs, all consist of baryonic matter,

but do not emit any or very little light by themselves. Black holes, while not, strictly speaking, baryonic, are nevertheless objects which are known to form (even as the result of baryonic processes) and have been observed, similarly eliding the need for a new, "exotic" kind of matter. Such massive astrophysical compact halo objects (MACHOs), while still in principle interacting with other baryons, e. g. by absorbing incoming light, would prove elusive due to their small sizes.

PBHs are a particular example of such MACHOs. They are black holes which would have formed in the early universe, where sufficiently overdense regions could collapse into black holes. This allows for potentially much smaller masses than black holes which form as the result of stellar evolution.

Although evading detection via the electromagnetic spectrum, all MACHOs have in common that they are susceptible to discovery by gravitational microlensing. As MACHOs would have to be numerous to make up the missing dark matter, they would often pass near observable stars, causing a transient change in their brightness through lensing induced by their gravitational influence. This has been used to establish strict bounds on the viability of these dark matter candidates. PBHs additionally have a lower limit set on their available mass range because very light black holes would evaporate via Hawking radiation.

1.4 "Small-scale challenges" to ΛCDM

There are a number of issues related to structure formation with CDM at small scales, where the non-linear evolution of structure as predicted (e. g.) by numerical simulations, is difficult to reconcile with observations. Conspicuously, these problems arise at the same scales where baryonic physics become significant. While this could be a hint that these challenges are not actually problems with CDM, but rather the lack of a complete understanding of the complicated baryonic physics, it is at least rendering it very difficult to distinguish an incomplete description of baryons from inconsistencies caused by the nature of dark matter. The extent to which these challenges present a genuine threat to CDM and the cosmological paradigm as a whole is likewise a matter of enduring debate, with some going so far as to proclaim a "small-scale crisis" of ΛCDM (D'Onghia and Lake 2004; Viel et al. 2013; Nakama, Chluba, and Kamionkowski 2017).

There have been a number of contentious points in this category, some of which are considered to have been resolved over time. The ones that remain as the current

focus of inquiry are mostly the following (Weinberg et al. 2015; Del Popolo and Le Delliou 2017; Bullock and Boylan-Kolchin 2017):

The cusp–core problem: This problem deals with the subject of halo density profiles. Numerical simulations of structure formation with ΛCDM predict that dark matter halos follow the Navarro–Frenk–White (NFW) density profile, which is "cuspy", i. e. steeply rising, towards the center. Some observations of dwarf galaxies, however, seem to indicate a flat, constant-density "core" at the centers of their halos instead. Unfortunately, neither the actually predicted density profiles of halos *when including baryons* and baryonic physics such as feedback processes, nor the observed density profiles themselves are settled facts.

The missing satellites/dwarf galaxies problem: The equations of CDM structure are scale-free, with the result that CDM will collapse and fragment into halos even on the smallest scales. This brings about an "explosion" of low-mass halos, which become ever more abundant as their masses decrease. Since halos of a certain size are usually expected to form galaxies, the number smaller (dwarf) galaxies predicted in this manner vastly outstrips those which have been observed, e. g. as satellite galaxies of the Milky Way. While the most direct solution would be to suppress the formation of smaller halos, as is the case in some other dark matter models (such as warm dark matter (WDM)), it could also be that galaxy formation becomes less efficient in these smaller halos.

The "too-big-to-fail" problem: Somewhat related to the previous problem, there is another mass range where ΛCDM predicts more halos than can be matched to observed galaxies in the local universe. Halos with masses around $10^{10}\,M_\odot$ should be practically guaranteed to form stars (they are "too big to fail" to do so), but observations show an underabundance of galaxies which would reside in such halos. However, since the number of observed galaxies (e. g. the "classical dwarf" satellites of the Milky Way) is actually compatible with the number of the most massive sub-halos in simulations, another way to view the problem is that there is a mismatch between the central densities in the simulations and observations. This takes the issue in a similar direction as the core–cusp problem, once again raising the issue of the density distribution in the inner parts of halos.

The plane of satellites problem: While less clearly related to the model of dark matter on its own, it has been observed that the Milky Way's satellite galaxies

lie in a plane, which would be a rare occurrence and difficult to explain using structure formation with ΛCDM. On the other hand, this arrangement of the galaxies seems to be a transient "coincidence" occurring only for a short time in their orbits, and with the countless possible histories every halo and galaxy can undergo, it might simply be a statistical outlier. The problem then turns into the question of how to quantify the "rareness" of what is seen in the local universe, and what degree of statistical deviation would be cause for concern.

It is evident that for each of these claimed problems, legitimate counterpoints have been raised that have the potential to resolve the problem without appealing to an overhaul of the cosmological paradigm. As it stands, individually, none of them appear to unveil any fundamental, irreconcilable flaws with CDM. Yet, they should not be dismissed prematurely, either, and remain open questions.

1.5 "Alternative" dark matter models

In part due to the small-scale challenges confronting ΛCDM, but also increasingly because of the persistent lack of CDM detection (particularly WIMPs), a number of alternative models with different implications at smaller scales have arisen. As for CDM, these models are really broader categories which operate on the level of what can be observed astronomically, but which behave in the same manner on such astronomical scales. For each of these classes of dark matter, there exist again several possible micro-physical implementations. However, the origin of these models can also be viewed to proceed in a "bottom-up" fashion: Given a newly-proposed (e. g.) dark matter particle, if this is found to behave differently from CDM (but coinciding on large scales to preserve agreement with established observational facts), it might create a new class in the form of an "astrophysical" dark matter model. In fact, usually a concrete dark matter model would be investigated first, and only then be classified depending on its behavior on astronomical scales.

1.5.1 Warm dark matter

The term WDM refers to matter which is a "middle ground" between CDM, such as WIMPs, and hot dark matter, such as neutrinos. Although WDM is also non-relativistic matter, it has a non-zero velocity dispersion (contrary to CDM). This leads to a cut-off in the matter power spectrum on small scales, preventing small-scale structure formation. In particular, this suppresses the formation of dark matter halos

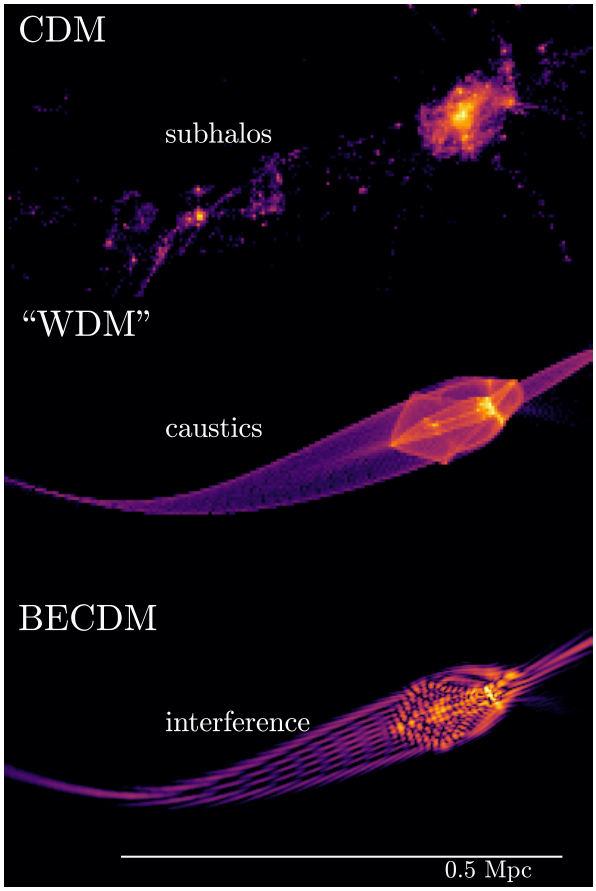

Figure 1.6: Density slice through the same dark matter filament at $z = 7$ in numerical simulations with different types of dark matter: CDM, a version similar to WDM, and FDM. CDM displays fragmentation into halos down to the smallest scales, while WDM and FDM (here labeled "BECDM") show caustics. FDM additionally has coherent interference patterns modulating the caustic structures (taken from Mocz et al. 2020).

below a certain mass threshold, potentially addressing the missing satellites problem. Figure 1.6 is a visualization of the effect this has on the formation of halos in the "cosmic web".

The archetypal example of a specific WDM model are sterile neutrinos, whose masses are in the keV range (Boyarsky et al. 2019). Apart from their WDM properties, they have strong motivation from particle physics as an explanation of neutrino masses. Since neutrinos are massless in the standard model of particle physics, the confirmation of non-zero (albeit small) neutrino masses requires an extension of the standard model, and sterile neutrinos in the keV mass range are one of the mechanisms which would result in massive neutrinos compatible with experimental data.

1.5.2 Self-interacting dark matter

From the particle physics perspective, it is very difficult to construct a model whose particles do not have *some* degree of self-interaction. However, the question is whether these interactions are strong enough to be relevant in the dark matter regime. Self-interacting dark matter (SIDM) is a dark matter model which assumes the presence of dark matter self-interactions to the extent that they have observable consequences. Often, this is in a very simple form, e. g. assuming a constant self-interaction cross section, but such models are heavily constrained by measurements from systems like the Bullet Cluster.

SIDM has less straightforward implications for structure formation than e. g. WDM, whose impact can be described with a simple change in the matter power spectrum (Tulin and Yu 2018). Generally, the impact will be strongest in the densest regions, i. e. the inner parts of dark matter halos, simply because a higher concentration of matter increases the amount of (self-)interactions occurring. Similar to baryons, the departure from the collisionless regime in favor of increased scattering between dark matter particles prevents their collapse to some extent, instead forming isothermal cores at the centers of halos. This posits SIDM as a potential solution to the cusp–core problem.

1.5.3 Ultra-light dark matter

Most of the specific dark matter models thus far have been in the form of elementary particles. While ultra-light dark matter models also postulate the presence of new, heretofore-undetected particles, their masses are so incredibly light that they behave collectively as a wave instead of as individual particles – e. g. in the form of a Bose–Einstein condensate –, similar to how photons collectively form electromagnetic waves. Correspondingly, another term for this type of dark matter is wave dark matter, which has very different dynamics than most other candidates based on collisionless particles, with quantum matter wave effects appearing on macroscopic, even galactic scales, forming, among other things, wave interference patterns and transient fluctuations in the density field. Generally, this means that these new particles must be bosons (e. g. scalar particles), as fermions are restricted from occupying the same states by the Pauli exclusion principle.

The simplest form of ultra-light dark matter, FDM, is the main subject of this dissertation. However, there are several more involved scenarios, e. g. including a

self-interaction term, which results in superfluid dark matter, or higher-spin wave dark matter (see also the review articles cited in chapter 2).

Figure 1.6 shows how FDM suppresses small-scale structure similar to WDM, but introduces additional effects in the form of wave interference patterns. More details concerning the motivation and phenomenology of FDM are laid out in chapter 2.

1.5.4 Modified gravity

As mentioned before, it is not strictly necessary for dark matter to actually be "matter" as would commonly be understood, as long as whatever physical mechanism used to replace it effectively behaves as such (Mannheim 2006). One of the first attempts to explain the effects of dark matter by modifying the laws of gravity was in the form of Modified Newtonian dynamics (MOND), which provides an empirical modification of the Newtonian inverse-square law of gravity in order to reproduce observed galactic rotation curves (section 1.2.3) without invoking a "missing" component of matter. However, MOND itself is an ad-hoc non-relativistic solution for a subset of the observational evidence, which is unable to explain the full range of observed gravitational effects, such as gravitational lensing (section 1.2.5). More general models, which modify the laws of general relativity in a consistent manner, are a subject of active research (Clifton et al. 2012).

2 The fuzzy dark matter model

After the introduction of the general cosmological paradigm in chapter 1, I will now introduce the main subject of this thesis, the FDM model, in more detail. For even more in-depth information, there are also a number of recent review articles which summarize many aspects of the field (Marsh 2016; Hui et al. 2017; Niemeyer 2020; Hui 2021; Ferreira 2021).

2.1 Theoretical background

Fundamentally, FDM is identical to the well-studied scalar field dark matter model, which is perhaps also the simplest particle-based dark matter model at the theoretical level. Such a model is described by the simple scalar field action

$$S = \frac{1}{\hbar c^2}\int \mathrm{d}^4x\,\sqrt{-g}\left(\frac{1}{2}g^{\mu\nu}\big(\partial_\mu\phi(x)\big)\big(\partial_\nu\phi(x)\big) - \frac{1}{2}\frac{m^2c^2}{\hbar^2}\phi(x)^2 - \frac{\lambda}{\hbar^2c^2}\phi(x)^4\right) \quad (2.1)$$

with the metric $g^{\mu\nu}$ and its determinant g, a real scalar field ϕ, its mass m, and a self-interaction coupling strength λ.[1] In this most fundamental description of quantum field theory in a curved spacetime, ϕ is a quantum field. This action is to be understood in the context of quantum field theory in a curved spacetime, i. e. ϕ is a ("second-quantized") quantum field, although for the purposes of numerical calculations, I will only consider the "classical" and non-relativistic limits.

The difference to most earlier studies of scalar field dark matter lies in the chosen parameters, specifically the mass m, which is commonly assumed to be in the range of 100 GeV–TeV, in line with the concept of CDM. In contrast, an *ultra-light* scalar field with a mass around $mc^2 \approx 10^{-22}$ eV, making fuzzy dark matter a form of scalar field dark matter with an "unusual" choice of the scalar field particle mass m, is under consideration here. This vast difference in the considered mass range compared to

[1] c and $\hbar$ are the speed of light in vacuum and the reduced Planck constant, and are explicitly included in all equations.

heavy scalar particles has drastic phenomenological consequences. Additionally, the ultra-light particles require a *non-thermal* production mechanism so that the resulting dark matter is not ultra-hot, but still resembles CDM.

The term "fuzzy dark matter" is typically used for an ultra-light scalar field *without* self-interactions, i. e. $\lambda = 0$. Thus, it corresponds to a limit ($\lambda \to 0$, or zero temperature, $T \to 0$) of more general ultra-light scalar field models (see Ferreira 2021, for a review), such as superfluid dark matter. Such models are often called ALP models due to the phenomenological similarity to the axion of QCD, which yields the same action as eq. (2.1) originating from a periodic potential $V(\phi) \propto \Lambda^4(1 - \cos(\phi/f_a))$ for $\phi \ll f_a$ (see e. g. Marsh 2016; Hui et al. 2017).

As mentioned above, the distinguishing feature of the FDM model is the value of the mass m, which is around $mc^2 \approx 10^{-22}\,\mathrm{eV}$, making this an *ultra-light* scalar field. This is in stark contrast to most implementations of CDM, such as WIMP models (which include scalar field dark matter models), which feature particle masses in the range of 100 GeV–TeV. The consequence of this wildly different mass regime is that the scalar field is better described as waves instead of individual particles, making FDM an example of a *wave dark matter* model, with striking differences in phenomenology. The reason is that, for such ultra-light values of the particle mass, the resulting particle number densities are extremely large, such that the quantum-mechanical de Broglie wavelength will be much larger than the inter-particle separation. Correspondingly, a collective wave description becomes much more appropriate than dealing with individual particles. Such dark matter models imply a non-thermal production mechanism to remain "cold", in which case bosons (like the scalar particles of FDM) can form a Bose–Einstein condensate.

I will assume non-relativistic approximations appropriate for simulations of cosmic structure formation, where all velocities are $\ll c$ and the considered scales are smaller than the Hubble horizon c/H_0. Rewriting the real scalar field ϕ in terms of a complex variable ψ,

$$\phi = \frac{1}{2}\sqrt{\frac{\hbar^3 c}{2m}}\,\mathrm{Re}\left(\psi e^{-i\frac{mc^2}{\hbar}t}\right) = \sqrt{\frac{\hbar^3 c}{2m}}\left(\psi e^{-i\frac{mc^2}{\hbar}t} + \psi^* e^{i\frac{mc^2}{\hbar}t}\right), \tag{2.2}$$

and taking the non-relativistic limit in Newtonian gauge yields the equations of

motions for fuzzy dark matter, the Schrödinger–Poisson (SP) equations:[2]

$$i\hbar\partial_t\psi(t,\vec{x}) = -\frac{\hbar^2}{2ma(t)^2}\nabla^2\psi(t,\vec{x}) + \frac{m}{a(t)}\Phi\psi(t,\vec{x}) \tag{2.3}$$

$$\nabla^2\Phi(t,\vec{x}) = 4\pi G m\big(|\psi(t,\vec{x})|^2 - \langle|\psi|^2\rangle(t)\big), \tag{2.4}$$

where a is the cosmological scale factor, Φ is the Newtonian gravitational potential, and the angle brackets in $\langle|\psi|^2\rangle$ indicate the spatial average. These are the equations of motions whose non-linear time evolution is solved in my numerical simulations. In the above, all quantities and coordinates are given in "comoving" form, to factor out the dependence on the scale factor $a(t)$ as much as possible.[3] Explicitly, comoving and "physical" quantities are related as follows:

$$\vec{x} = a^{-1}\vec{x}_{\text{phys}}, \quad \nabla = a\nabla_{\text{phys}}, \quad \psi = a^{3/2}\psi_{\text{phys}}, \quad \Phi = a\Phi_{\text{phys}}. \tag{2.5}$$

As I have assumed a production mechanism that can be described as a Bose–Einstein condensate, with almost all of the scalar particles in their quantum-mechanical single-particle ground states and large number densities, I can make use of the mean field approximation, also called the "classical" limit or wave limit, where the particles behave collectively and coherently. The "wave function" ψ in eq. (2.3), which is now simply a complex number-valued function (instead of a quantum field), is then understood as the single macroscopic amplitude of the FDM waves, with the mass density given by

$$\rho = m|\psi|^2. \tag{2.6}$$

This is why FDM is often called a "classical theory".[4]

As in quantum mechanics, or generally for diffusion equations, eq. (2.3) obeys the continuity equation

$$\partial_t\rho + \nabla\cdot\rho\vec{v} = 0, \tag{2.7}$$

[2]Since eq. (2.3), despite being formally identical to the Schrödinger equation from single-particle quantum mechanics, has a very different physical meaning and is in fact a Gross–Pitaevskii equation (Gross 1961; Pitaevskii 1961), eqs. (2.3) and (2.4) are technically the Gross–Pitaevskii–Poisson system of equations. However, I will use the more common term "Schrödinger–Poisson equations" here.

[3]All quantities will be given in terms comoving lengths unless specified otherwise.

[4]The distinction is irrelevant for the physical and phenomenological consequences of FDM, but often seems to spark debate. I would like to note that, although FDM may be described by what is formally a classical field theory, the underlying *physical* phenomena, such as matter waves and Bose–Einstein condensation, were not known in classical physics.

with the momentum density

$$\rho\vec{v} = \frac{\hbar}{2i}(\psi^*\nabla\psi - \psi\nabla\psi^*) \tag{2.8}$$

The momentum density $\rho\vec{v}$ can also be used to define the bulk peculiar velocity field $\vec{v}$, although this is problematic in the case $\rho = 0$. Separating the complex wave amplitude ψ into its absolute value $\sqrt{\rho/m}$ and complex phase θ,

$$\psi = \sqrt{\frac{\rho}{m}}e^{i\theta}, \tag{2.9}$$

yields a simpler form of eq. (2.8):

$$\vec{v} = \frac{\hbar}{m}\nabla\theta, \tag{2.10}$$

i. e. the velocity is given by the gradient of the wave amplitude's phase. Again, however, this is ill-defined when $\psi = \rho = 0$, in which case the phase θ is undefined. It should be noted that due to the presence of wave interference, the case $\psi = 0$ is actually quite common (occurring wherever destructive interference takes place) and cannot be neglected as perhaps for non-wave dark matter models.

The wave function in polar components (eq. (2.9)) can be rewritten in hydrodynamical form using the Madelung (1927) transformation, which results in the continuity eq. (2.7) along with a modified Euler equation

$$\partial_t\vec{v}_\mathrm{c} + \frac{1}{a^2}\nabla_\mathrm{c}\vec{v}_\mathrm{c}^2 = -\frac{1}{a}\nabla_\mathrm{c}\Phi_\mathrm{c} + \frac{1}{a^2}\frac{\hbar^2}{2m^2}\nabla_\mathrm{c}\frac{\nabla_\mathrm{c}^2\sqrt{\rho_\mathrm{c}}}{\sqrt{\rho_\mathrm{c}}} \tag{2.11}$$

(cf. e. g. Mocz et al. 2020). This allows for a hydrodynamical interpretation of the density and velocity fields ρ and $\vec{v}$.

In addition to the mass (as implied by the continuity equation, eq. (2.7)), the total energy, given by

$$\begin{aligned} E &= \int \mathrm{d}^3x \left(\frac{\hbar^2}{2m^2}|\nabla\psi|^2 + \frac{1}{2}\Phi|\psi|^2\right) \\ &= \int \mathrm{d}^3x \frac{\hbar^2}{2m^2}(\nabla\sqrt{\rho})^2 + \int \mathrm{d}^3x \frac{1}{2}\rho v^2 + \int \mathrm{d}^3x \frac{1}{2}\rho\Phi \\ &= T_\rho + T_v + V \end{aligned} \tag{2.12}$$

is conserved (Mocz et al. 2017). The kinetic energy $T = T_\rho + T_v$ is made up of the "bulk" kinetic energy T_v and a gradient energy term T_ρ, while V is the standard gravitational potential energy.

The separation of the wave amplitude into two real variables and their connection with the density (2.6) and velocity (2.10) allows a hydrodynamical interpretation analogous to that of Madelung (1927) in quantum mechanics, although its validity is problematic for destructive interference. FDM's resistance against gravitational collapse on small scales, which can be interpreted as an analogue to the Heisenberg uncertainty principle in the SP formulation, then manifests itself as an explicit additional "quantum pressure" appearing in the analogue of the hydrodynamical Euler equation.

There are two important length scales which serve as indicators of FDM wave phenomena, both of which are determined by the constant $\hbar/m$, which is the only independent parameter in the equations of motion, eqs. (2.3) and (2.4). The first is the de Broglie wavelength

$$\lambda_{\mathrm{dB}} = \frac{2\pi\hbar}{mv}, \tag{2.13}$$

which has the same meaning as in quantum mechanics and indicates the length scale on which density fluctuations of order one occur. The second is the FDM analogue of the Jeans length. In linear perturbation theory, this is the length scale where the gravitational attraction balances the "quantum pressure" resisting collapse, such that perturbations on larger scales grow, while smaller ones oscillate. Expressed as a wave number k_{J}, this is (Hu, Barkana, and Gruzinov 2000)

$$k_{\mathrm{J}} = \left(\frac{6\Omega_{\mathrm{m}}}{1+z}\right)^{1/4}\left(\frac{mH_0}{\hbar}\right)^{1/2}. \tag{2.14}$$

as a function of the redshift z, with the Hubble parameter H_0 and the cosmic matter density parameter Ω_{m}.

There is a correspondence between the Vlasov–Poisson equations which describe the evolution of a CDM fluid and the Schrödinger–Poisson equations, with the latter reducing to the former for $\hbar/m \to 0$ (Widrow and Kaiser 1993; Mocz et al. 2018; Garny, Konstandin, and Rubira 2020). This explains why FDM behaves like CDM on large scales. A wave function $\psi(\vec{x})$ can be constructed from a phase space distribution function $f(\vec{x}, \vec{p})$ by means of a Wigner or Husimi transform. This relationship illustrates why FDM behaves like CDM on scales larger than λ_{dB}. At first, this was in fact considered as an alternate method to simulate CDM (Widrow and Kaiser

1993). However, one can also take the equations seriously and actually consider the Schrödinger–Poisson equations as the "true" description of dark matter. In this case, the correspondence allows one to "translate" a phase space distribution function to the wave function formalism. This is very useful because it enables direct comparisons between structure formation with CDM and FDM from the same initial conditions (ICs).

In general, a wave function on a discretized lattice can be constructed using the prescription

$$\psi(\vec{x}_{\vec{n}}) \propto \sum_{n_1'=0}^{N-1} \dots \sum_{n_3'=0}^{N-1} \sqrt{f(\vec{x}_{\vec{n}'}, \vec{v}_{\vec{n}})}\, e^{i\frac{m}{\hbar}\vec{x}_{\vec{n}'}\cdot\vec{v}_{\vec{n}}+R_{\vec{n}}}, \tag{2.15}$$

where $\vec{n}, \vec{n}' = (n_1', n_2', n_3')$ are discrete grid indices, $\vec{x}_{\vec{n}} = \vec{n}\Delta x$ and $\vec{v}_{\vec{n}} = 2\pi\frac{\hbar}{m}\frac{\vec{n}}{N\Delta x}$ are discrete phase space grid points, and $R_{\vec{n}}$ is a random phase which is required to ensure that different velocity components are uncorrelated (cf. Widrow and Kaiser 1993; Mocz et al. 2018).

For the case of a "cold" or "single-stream" distribution function as in the case of CDM, where each point in space has a single well-defined value for the velocity, this value can be directly related to the wave function's phase as in eq. (2.10). The construction of the wave function then simply reduces to the polar decomposition (eq. (2.9)), with the absolute value and phase determined by the density and velocity at each point, respectively:

$$|\psi(\vec{x})| = \sqrt{\frac{\rho(\vec{x})}{m}}, \tag{2.16}$$

$$\nabla \arg(\psi(\vec{x})) = \nabla\theta(\vec{x}) = \frac{m}{\hbar}\vec{v}(\vec{x}). \tag{2.17}$$

Equation (2.17) can be easily solved numerically by applying the spectral method to the equation

$$\nabla^2\theta(\vec{x}) = \frac{m}{\hbar}\nabla\cdot\vec{v}(\vec{x}). \tag{2.18}$$

2.2 Astrophysical phenomenology

2.2.1 Suppression of small-scale structure

Due to the analog of the Heisenberg uncertainty principle, or the "quantum pressure" in the fluid formulation, FDM resists collapse on small scales. In linear perturbation

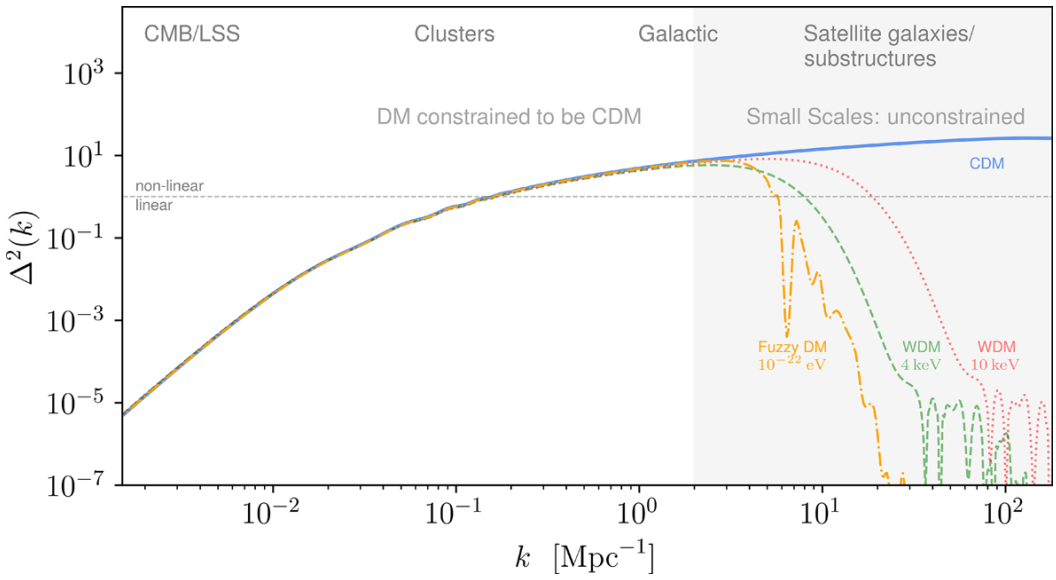

Figure 2.1: Matter clustering, as measured by the matter power spectrum, for different dark matter models. (Taken from Ferreira 2021).

theory, this results in a cut-off in the power spectrum on scales smaller than the Jeans scale eq. (2.14). For higher FDM particle masses, this cut-off moves to smaller scales.

Although the precise shape of the linear power spectrum and the cut-off is different, this behavior is very similar to the impact of WDM (cf. fig. 2.1). It similarly results in a suppression of small-scale structure, even without taking into account the characteristic wave dynamics of FDM. Due to this similarity, many attempts to put constraints on the FDM parameter space rely on translating the formalism for WDM to a slightly modified transfer function.

2.2.2 Soliton cores

Using numerical simulations of FDM halos, it was demonstrated that approximately spherical, constant-density cores called solitons form at the centers of FDM halos (Schive, Chiueh, and Broadhurst 2014). These solitons correspond to the spherically symmetric ground state solution of the SP equations (Edwards et al. 2018).

Firstly, since solitons are compact with a relatively sharp edge, they can disturb or tidally disrupt other surrounding matter, such as star clusters, especially as they move around the halo relative to other matter (Li, Hui, and Yavetz 2021). Secondly, as the soliton cores are present in a perturbed state characterized by interactions with matter from their host halo and its environment, they are actually not static, but

themselves fluctuate in their density amplitude (which is also called the "breathing mode"). Both effects can significantly influence the dynamics of visible matter at the centers of halos compared to a pure NFW profile.

2.2.3 Granules and interference patterns

Also first shown in detail in (Schive, Chiueh, and Broadhurst 2014) is that FDM halos feature transient density fluctuations order one (i. e. relative to the time-averaged density), which are often called granules, around the soliton core and throughout the halo. These arise from the superposition of waves, which results in constructive and destructive interference (outside of halos, this is also visible in the characteristic interference fringes which are reminiscent of interfering electromagnetic waves). Similar to the gravitational impact caused by the soliton, these fluctuations have the potential to disrupt other systems present within the halo, such as stellar populations, but on a much larger scale since they exist throughout the entirety of the halo (Dalal and Kravtsov 2022).

2.2.4 Vortices

Since the FDM velocity is a gradient field (see eq. (2.10)), it does not appear to have any vorticity. However, the complex phase is undefined at the sites of destructive interference, where $\rho = 0$. Consequently, these are the locations where vorticity can develop (Hui et al. 2021). The properties of these objects and their distribution present some unique opportunities for the potential detection of FDM, and have also been proposed as an explanation of observations of cosmic filament spin (Alexander et al. 2022).

2.3 Observational constraints

After FDM received increasing interest due to motivation from astrophysics and particle physics, constraints on the mass m of the scalar particles at values around $mc^2 \gtrsim 10^{-21}$ eV accumulated using many different contexts and observables for some time (Ferreira 2021). A recent summary is shown in fig. 2.2. FDM's soliton halo cores and their scaling relations have revealed themselves to be seemingly incompatible with observations if a single value is assumed for the particle mass m (Burkert 2020) (see also section 4.9). This by itself does not rule out FDM as a dark matter candidate

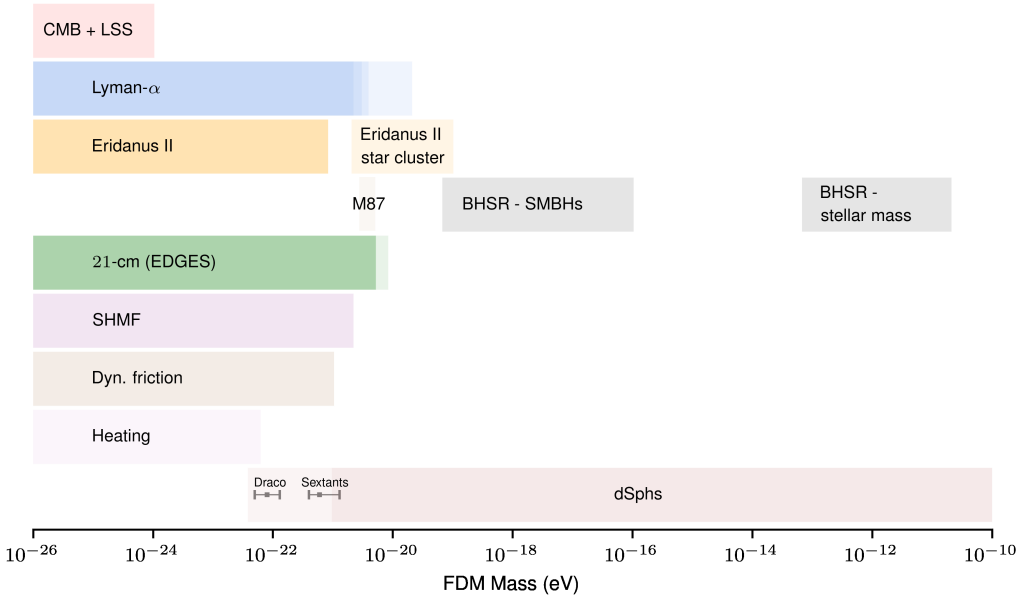

Figure 2.2: Compilation of observational constraints on m (taken from Ferreira 2021).

– after all, if the origin of the "cusp–core problem" does not lie in the nature of dark matter, this merely means that this problem cannot be used to constrain this nature. In this sense, fig. 2.2 must also be viewed with care, since not all bounds apply in general to the entire model, but sometimes make additional assumptions (such as those concerning dSph galaxies). However, this does weaken FDM's motivation as a possible solution to the cusp–core question. Further, it constrains the parameter space, since the cores must be relegated to regimes where they would not have a strong observational impact if it is determined that their observation would be incompatible with known results.

In a similar vein, dynamical modeling has recently been applied to ultra-faint dwarf (UFD) galaxies, which would feature the largest FDM cores, whose results have lead to claims of strong constraints on m, up to $\approx 10^{-19}$ eV (Hayashi, Ferreira, and Chan 2021; Zoutendijk et al. 2021). However, these analyses do not take the significant scatter in the FDM core–halo mass relation, discovered in cosmological simulations, into account (Chan, Ferreira, May, et al. 2022), which would weaken these bounds. A very recent result from Dalal and Kravtsov (2022), claiming $mc^2 > 3 \times 10^{-19}$ eV due to the heating effect of stellar orbits caused by potential fluctuations in FDM halos, has more serious implications, since it does not depend on the uncertain core–halo relation. However, most of these constraints (including the most recent) have in

common that they have not been verified by self-consistent numerical simulations, instead relying on approximate, idealized, or simplified numerical and analytic approaches.

Numerical methods and implementation

3

In this chapter, I will briefly describe the "state of the art" of current methods in computational cosmic structure formation in order to set the stage for the general framework in which the work in this dissertation was performed in, as well as present the methodology and implementation developed specifically for the simulation of structure formation within FDM cosmologies. For a more extensive recent review of simulation techniques and results with different dark matter models, including FDM, see Angulo and Hahn (2022).

3.1 Numerical simulations of cosmic structure formation

Due to the multitude of physical scales involved in cosmic structure formation, both in space and in time, the endeavor of performing numerical simulations of this non-linear process involves a whole "stack" of tools and ingredients, as well as many technical details. To summarize, a "cosmological simulation of structure formation" consists of the following components and inputs:

A "background cosmology": It is generally not feasible to follow the evolution of a fully self-consistent general-relativistic space-time. Consequently, simulations usually work with the non-relativistic, weak-field Newtonian approximation and assume a static FLRW "background cosmology". This background is static because it does not take into account how inhomogeneities and changes in the matter distribution would affect the space-time metric: there is no "back-reaction" of matter on the metric (Buchert and Räsänen 2012). Instead, there is only one global, position-independent expansion rate throughout the entire simulation volume, corresponding to the approximation that the universe is homogenous when smoothed on large scales. On smaller scales, peculiar velocities and accelerations are expected to dominate over the "Hubble flow"

of cosmic expansion, justifying the static background approximation. While this is the framework predominantly applied in large simulations of structure formation, approaches investigating the impact of back-reaction using a general relativity (GR) treatment have been developed as well (Barrera-Hinojosa and Li 2020).

The assumption of a background cosmology ultimately boils down to the requirement of taking spatial expansion into account in the equations of motion. Usually, when expressed in terms of suitably-defined "co-moving" quantities, this only results in some additional factors of the scale factor $a(t)$ appearing in the equations compared to the corresponding "non-cosmological" versions. Since the evolution of expansion in the form of the scale factor must be followed, cosmological simulations include the FLRW parameters – usually the density parameters Ω_{c0}, Ω_{b0} and $\Omega_{\Lambda 0}$, and the Hubble parameter H_0 – as external inputs. In the case of ΛCDM, no cosmological parameters are required beyond this; however, for other dark matter models, more parameters, such as the particle mass for WDM and FDM, or the self-interaction cross section for SIDM, enter as additional inputs.

"Low"-redshift ICs: For a number of reasons, it is neither feasible nor efficient to attempt to follow the evolution of the universe (or some region within it) throughout all of time – from the Big Bang to the present day – within a single framework. The early, hot, relativistic universe is characterized by several phase transitions and countless physical processes which are not only difficult and expensive to treat comprehensively, but also completely irrelevant today. Further, it is not the precise micro-arrangement of matter and energy (which would be hopeless to keep track of in detail in any case), but rather their global and statistical properties which are most relevant for the following formation of structure. At slightly later times, the evolution can be very accurately described analytically by perturbation theory, making a fully non-linear numerical treatment unnecessary.

A numerical method: Given ICs at some time, the appropriate equations of motion for the relevant physics to be included in the simulation generally allow to solve for the evolution of state over time. Since the resulting non-linear equations cannot be solved analytically, a numerical scheme must be devised which allows for the computational treatment of the equations.

3.2 Simulations of structure formation with ΛCDM

For ΛCDM, the fundamental non-relativistic equations describing the evolution of the dark matter fluid are the Vlasov–Poisson equations. The fact that CDM is "cold" with negligible velocity dispersion implies that the fluid which it describes is initially "single-streaming", i. e. there is only a single value of the velocity associated with each point in space with which the fluid is coherently moving:[1]

$$f(\vec{x}, \vec{v}) = n(\vec{x})\delta(\vec{v} - \nabla S(\vec{x})) \,, \tag{3.1}$$

where $n(\vec{x})$ is the number density and $S(\vec{x})$ a velocity potential. This allows for an accurate description of phase space by sampling discrete points, each having a position, velocity and mass associated with them, resulting in a "coarse-grained" version of the fluid and its phase space. These discrete sampling points are usually called "simulation particles", "macro-particles" or simply "particles" for short. However, it should always be clear that these do not correspond in any way to physical particles (such as elementary particles – their masses are many orders of magnitude larger, typically, than even the mass of the sun, $M_\odot$), but are simply part of a procedure which makes the underlying equations tractable for a numerical treatment. Since it is impossible to handle a continuous field (such as density or phase space) without an infinite amount of (discrete) memory storage and computational power, any numerical approach must necessarily involve some kind of discretization step.

Except for specialized applications, this N-body approach is the vastly predominant framework which is used to dynamically evolve the distribution of CDM. It is implemented by the vast majority of simulation software in common use, such as GADGET-4 (Springel et al. 2021) and its predecessors, AREPO (Springel 2010; Weinberger, Springel, and Pakmor 2020), SWIFT (Schaller et al. 2016), GIZMO (Hopkins 2017), and many others. Concerning the specific numerical approaches to solving the resulting equations of motions after discretization via the N-body prescription, there is slightly more variation among implementations. Through a number of steps, the original problem has returned to the familiar N-body problem of Newtonian gravity (although N is vastly lower than if all physical particles were considered directly).

[1] As the CDM distribution evolves, it will not remain perfectly cold and eq. (3.1) eventually becomes invalid. As structures collapse, they develop "multi-streaming", with multiple values of the velocity associated with a single point in space (a non-zero velocity dispersion). These events are referred to as "shell crossing".

Even so, the simplest approach of determining the evolution of an N-body system in Newtonian gravity, direct summation, incurs a computational cost of order $O(N^2)$, which quickly becomes prohibitive even for relatively low particle numbers (at least when comparing the effective spatial resolution to the sizes of the simulated cosmological and astrophysical systems). Some of the most common state-of-the-art approaches to this problem are tree methods, the particle–mesh (PM) method, their combination TreePM, and fast multipole methods (FMM).

3.3 Simulations of structure formation with fuzzy dark matter

For FDM, the origin and framework of the equations of motions is very different than for CDM. While the treatment of CDM is based on a phase space as in classical statistical mechanics, FDM is a wave theory with completely different fundamental variables. The difference is akin to the contrast between classical (e. g. Lagrangian or Hamiltonian) mechanics and quantum mechanics: While the dynamics in the former are described using a six-dimensional phase space consisting of positions $\vec{x}$ and velocities $\vec{v}$ (or momenta $\vec{p}$) as input variables, the latter contains all information in a function of only the three-dimensional positions, at the "cost" of the function being a complex-valued "wave function" instead of a real-valued distribution function.

These incompatible frameworks raise conceptual questions about how quantities correspond to each other and how they should be compared in both cases. While there are prescriptions relating wave functions and six-dimensional phase space, these are not unique and do not necessarily have the properties expected of "true" phase space distribution functions. Another approach, making use of the Madelung fluid correspondence between the Schrödinger equations and a hydrodynamical analog, can serve as yet another source of intuition about the relation of wave and particle quantities.

On the level of equations of motions, approaches to numerical FDM simulations can be separated into these two categories, namely those making use of the Madelung fluid formulation and those tackling the SP equations directly. As with N-body CDM or hydrodynamics, implementations then further differ with respect to the specific numerical schemes used to solve either set of equations of motion.

3.3.1 Generating cosmological initial conditions

As detailed in section 3.2, numerical simulations of non-linear cosmic structure formation in the late universe generally start with ICs derived from perturbation theory (linear or higher order) at a time where this approach is still valid. In principle, it would be necessary to adopt a similar approach for FDM, using (e. g.) the linearized equations to obtain an initial, perturbed density amplitude $\psi(\vec{x})$ at the appropriate redshift. However, it is possible to make use of the correspondence between FDM and CDM (the Schrödinger–Poisson–Vlasov–Poisson correspondence; Mocz et al. 2018) to directly map ICs from the well-known ΛCDM machinery to FDM.

As long as the initial redshift is chosen to be before shell crossing occurs, this procedure is safe, since fluctuating granules and interferences patterns develop in FDM after this point, which would not be captured in CDM ICs. Fortunately, this is not a significant restriction, because perturbation theory is generally only valid up to shell crossing to begin with.

3.4 The AxiREPO code

One major result of the work on this dissertation is the AxiREPO code. It has been developed to perform highly performant and scalable numerical simulations of FDM, with and without baryons, and is the most fundamental instrument enabling all the scientific results presented in the following chapters. Initially starting out with small, idealized, non-cosmological test problems, such as those treated in Edwards et al. (2018), the code's functionality continuously grew with the requirements of the projects. Today, it is a state-of-the-art tool for any kind of numerical simulation of FDM, and has been used to carry out SP simulations of cosmic structure formation with FDM at unprecedented scales – by far the largest to date.

AxiREPO is implemented as a module in the widely-used and highly successful AREPO code (Springel 2010), which has also been publicly released (Weinberger, Springel, and Pakmor 2020). AxiREPO implements a pseudo-spectral method with a second-order split-step time integration to solve the SP equations, eqs. (2.3) and (2.4). It boasts the following features:

- Both cosmological and non-cosmological simulations of FDM, solving the SP equations with arbitrary ICs. This means that the code can be used both for idealized or simplified setups, as well as for full-blown cosmological boxes. Further, it can handle any problem that can be described using the SP equations,

even including "unconventional" simulations of large-scale structure with CDM.

- Support for both dark matter-only simulations as well as full hydrodynamical simulations with baryons, even including a state-of-the-art galaxy formation model such as IllustrisTNG (Weinberger et al. 2017; Pillepich et al. 2018), making use of AREPO's moving mesh implementation of hydrodynamics.

- High performance and scalability: The code is implemented in the C programming language, using the Message Passing Interface (MPI) standard for distributed-memory parallelism. It has been used to perform simulations with grid sizes up to 8640^3, which is several orders of magnitude larger than any other equivalent existing work on FDM.

- A novel halo finding algorithm which operates on the Cartesian grid used by AxiREPO. This includes capabilities to compute halo masses (such as M_{200}) using the spherical overdensity algorithm, as well as the halos' gravitational "self-potential" (a measure for the gravitational binding energy). The halo finder is described in more detail in sections 4.3.5 and 5.3.3.

- An implementation of higher-spin wave dark matter, such as vector dark matter. This not entails the ability to perform simplified simulations of the same kind as Amin et al. (2022), but also much more general cosmological simulations, even including baryons.

- In principle, it is even possible to conduct simulations featuring both N-body dark matter particles and the SP grid, allowing for mixed CDM and FDM simulations similar to Schwabe et al. (2020), although there may be minor technical details to resolve since this mode of operation has not been tested in practice thus far.

In order to ensure the reproducibility of numerical results, as well as in the hopes of providing a useful and versatile tool to the scientific community, AxiREPO will also be released in a publicly documented form in the near future. Of course, this is not to say that the code cannot be acquired right now: It can always be obtained by contacting the author, but the missing piece for a proper public release is the lack of user-facing technical documentation, without which the use of a complex piece of numerical software becomes quite a daunting challenge unless guided by the author.

While the details of the pseudo-spectral method and the implementation within AxiREPO are laid out in sections 4.3 and 5.3, I will here discuss in more detail the motivation for choosing this numerical approach and its merits compared to other techniques. an important question is of course which of these should be chosen when setting out to perform new simulations. The first choice is one between the Madelung fluid formulation and the SP equations. For the work in this dissertation, I decided to focus on the SP form because the fluid approach is generally unable to capture the defining wave nature of FDM, with its interference fluctuations and granules, which is the primary novel feature distinguishing FDM from other particle dark matter candidates. While FDM fluid simulations may have their place for specific applications in regimes where the details of the wave nature are not important, the question of the validity of numerical results obtained in this fashion is always looming in the background with this approach. In general, it has been shown that there are significant discrepancies between fluid solvers and a full wave treatment (Li, Hui, and Bryan 2019). Without a clear identification of the regimes where the fluid approach yields reliable results, these questions will always cast doubt on the outcomes of such simulations.

Having decided to numerically solve the SP equations, the question remained on the numerical method to use. While one option would have been to implement several different numerical schemes, which would also offer the technically interesting opportunity to directly compare and validate several approaches, it is generally already technically challenging enough to develop even a single implementation until it can efficiently and reliably run the kind of large-scale simulation which was the aim in this dissertation. Generally, the technical issues encountered when going to large scales are difficult to predict and investigate, because testing big datasets on a large amount of computing resources is necessarily slow and expensive, as access to these resources is limited. Further, there is always a trade-off between implementing and optimizing more code, and actually applying it to obtain scientific results.

In the end, the decision for the numerical approach in AxiREPO fell on the pseudo-spectral method mentioned above. This choice has several advantages:

- Spectral methods generally provide excellent spatial convergence. While finite-difference methods are classified based on the polynomial order of the error ε incurred for a given resolution $\Delta x = L/N$, where L is the size of the domain and N the number of grid points, behaving as

$$\varepsilon \sim O((L/N)^p) \,, \tag{3.2}$$

e. g. $p = 2$ for a second-order method, spectral methods behave as

$$\varepsilon \sim O\big((L/N)^N\big) \tag{3.3}$$

(see e. g. Boyd 2001). In other words, when improving the grid resolution L/N by increasing the grid size N, finite-difference methods converge polynomially with order p (the error decreases like a polynomial as the resolution is increased), whereas spectral methods converge exponentially, and are thus sometimes said to be of "infinite" polynomial order. Since FDM simulations generally choose the lowest resolution that they can get away with within the resolution constraints due to the cost stemming from the rapid $\propto \Delta x^2$ scaling of the time steps, the pseudo-spectral method is the most accurate for a given simulation box.

- The total mass is automatically and manifestly conserved for the pseudo-spectral method. Other methods, such as adaptive mesh refinement (AMR) implementations based finite-difference schemes, do not have necessarily the same conservation property (e. g. Mina, Mota, and Winther 2020).

- The implementation of a pseudo-spectral solver for FDM is directly compatible with existing implementations of the PM scheme, such as the one present in AREPO. This simplifies the implementation and integration with the rest of the code, given that AREPO was chosen as the general framework for development in order to have direct access to its moving mesh implementation of hydrodynamics to handle baryonic physics. Indeed, almost no code changes were necessary to implement the potential update step of the pseudo-spectral algorithm – the existing PM code could be reused "as-is" for this purpose. Relatedly, AREPO already features a high-performance parallel Fast Fourier Transform (FFT) implementation based on the FFTW library (Frigo and Johnson 2005), which could be adapted for the SP solver as well.

However, there are also drawbacks to the pseudo-spectral method. Chief among them is the global nature of the FFTs: While this is also the property that affords it the superior spatial convergence mentioned above, it seriously limits its versatility because the method can only operated at a single, globally fixed resolution in boxes with periodic boundary conditions. This limitation is less severe than it might sound at first – although it does imply that to reach a given desired resolution somewhere in the simulation box, it must be enforced *everywhere*, in practice, even adaptive

approaches to FDM simulations using AMR cannot afford to significantly refine when running in large volumes due to the quadratic scaling of the time steps. Further, since refinement must take into account density *and* velocity for FDM, rather high resolution is generally required throughout the simulation volume. And finally, while an AMR code may conceptually perform less work by omitting operations in regions where they are not needed, modern computing architectures incur severe penalties for branches in the control flow. Thus, paradoxically, it is increasingly becoming less efficient to be "smart" with the code's logic, and instead "brute force" methods using optimized algorithms (like FFTs) are actually the most effective on highly parallel hardware such as graphics processing units (GPUs).

Large simulations of cosmic structure formation with fuzzy dark matter

4

The main contents of this chapter have previously been published as May and Springel (2021), and are reproduced here with slight modifications. Section 4.9 presents some of the findings published in Chan, Ferreira, May, et al. (2022).

Abstract

An ultra-light bosonic particle of mass around $10^{-22}\,\mathrm{eV}/c^2$ is of special interest as a dark matter candidate, as it both has particle physics motivations, and may give rise to notable differences in the structures on highly non-linear scales due to the manifestation of quantum-physical wave effects on macroscopic scales, which could address a number of contentious small-scale tensions in the standard cosmological model, ΛCDM. Using a spectral technique, I here discuss simulations of such fuzzy dark matter (FDM), including the full non-linear wave dynamics, with a comparatively large dynamic range and for larger box sizes than considered previously. While the impact of suppressed small-scale power in the initial conditions associated with FDM has been studied before, the characteristic FDM dynamics are often neglected; in these simulations, I instead show the impact of the full non-linear dynamics on physical observables. I focus on the evolution of the matter power spectrum, give first results for the FDM halo mass function directly based on full FDM simulations, and discuss the computational challenges associated with the FDM equations. FDM shows a pronounced suppression of power on small scales relative to cold dark matter (CDM), which can be understood as a damping effect due to "quantum pressure". In certain regimes, however, the FDM power can exceed that of CDM, which may be interpreted as a reflection of order-unity density fluctuations occurring in FDM. In the halo mass function, FDM shows a significant abundance reduction

below a characteristic mass scale only. This could in principle alleviate the need to invoke very strong feedback processes in small galaxies to reconcile ΛCDM with the observed galaxy luminosity function, but detailed studies that also include baryons will be needed to ultimately judge the viability of FDM.

4.1 Introduction

The "standard cosmological model" with ΛCDM has been extremely successful in describing a wide variety of cosmological observations across a broad range of physical scales (e. g. Frenk and White 2012; Bull et al. 2016). On small cosmological scales, however, challenges such as the "cusp–core problem", the "missing satellite problem", or the "too-big-to-fail problem" have sometimes raised questions about the validity of ΛCDM (Weinberg et al. 2015; Del Popolo and Le Delliou 2017; Boylan-Kolchin, Bullock, and Kaplinghat 2011). Since physical effects experienced by baryons can become relevant at these scales, it has proven difficult to disentangle the apparent discrepancies from baryonic physics. In addition, even though the cosmological properties of cold dark matter are constrained extremely well on large scales, its micro-physical nature is still completely unknown thus far.

In light of these small-scale questions and the enduring lack of any direct detection of the most well-studied dark matter candidates, in particular WIMPs, models based on ultra-light ("axion-like") scalar particles have been gaining interest as alternative dark matter models. Due to their small masses, quantum effects are expected to cause interesting wave-like behavior at small (i. e. galactic) scales compared to heavy particles like WIMPs or compact objects. These effects have also been proposed to elucidate some of the "small-scale problems" of CDM; for example, early numerical simulations showed that ultra-light scalars form cores in the centers of dark matter halos (Schive, Chiueh, and Broadhurst 2014). Furthermore, light (pseudo-)scalar particles are a common feature of theories in particle physics, from the original axion in QCD to a plethora of axion-like particles predicted by unified and early-universe theories such as string theories (Marsh 2016).

The amount of existing work studying structure formation with CDM utterly dwarfs that of such FDM, particularly concerning numerical simulations that reliably probe the non-linear regime. Computations have been performed using a variety of approaches and numerical methods, although many attempts were limited in scope (Zhang, Liu, and Chu 2018; Laguë et al. 2021, table 1). Correspondingly, our understanding of structure formation in FDM cosmologies is still comparatively

spotty compared to CDM, where decades of experience have led to extremely detailed insights. Apart from the lower level of research attention, an important reason impeding insight into FDM has been that numerically solving the corresponding equations of motion incurs very large computational costs – much higher than those associated with corresponding calculations of ΛCDM. In particular, while it has been established that, in the limit of large scales (or large particle mass), the FDM equations are equivalent to the Vlasov–Poisson equations of CDM (Mocz et al. 2018; Widrow and Kaiser 1993), the effects of FDM in (mildly) non-linear regimes of structure formation are still poorly understood when compared to CDM.

Due to the computational requirements, the cosmological volumes studied in simulations with full FDM dynamics have been especially limited (Woo and Chiueh 2009; Schive, Chiueh, and Broadhurst 2014; Veltmaat, Niemeyer, and Schwabe 2018; Mocz et al. 2020), which is an issue that this chapter seeks to improve upon. In particular, I would like to carry out simulations that smoothly connect the non-linear state reached in isolated FDM halos to the still linear large-scale structure, thereby bridging, in particular, the regime of mildly non-linear evolution where differences in the temporal evolution compared to CDM can be expected. To this end, I carry out very large uni-grid FDM simulations with a spectral method, which fully retains the quantum-mechanical effects. Because there is still a dearth of precision studies of how FDM compares to CDM for traditional measures of large-scale structure, I focus my analysis on central measures of matter clustering, namely the power spectrum and the halo mass function, and compare to ordinary ΛCDM where appropriate.

While methods which forego a treatment of the full wave dynamics have been able to conduct simulations with volumes much closer to those attainable using traditional N-body and smoothed-particle hydrodynamics (SPH) approaches for CDM (Schive et al. 2016; Veltmaat and Niemeyer 2016; Zhang et al. 2018; Nori and Baldi 2018; Nori et al. 2019), these do not capture inherent wave phenomena such as interference effects, which can have a significant impact on the overall evolution at least on small scales (Li, Hui, and Bryan 2019), leaving the validity of results obtained in this way unclear in the absence of similar computations solving the fundamental wave equations. In particular, while all simulations can easily incorporate the impact of the suppressed small-scale power spectrum present with FDM in the initial conditions, such methods either lack the wave nature of FDM entirely or only approximate it. Using the simulations presented in this chapter, I aim to clarify the reliability of such approximative results by explicitly omitting the suppression in the FDM initial conditions, starting instead from "standard" ΛCDM initial conditions. In this way, I

disentangle the two essential physical differences distinguishing FDM from CDM in cosmological numerical simulations: the initial conditions and the non-linear dynamics.

This chapter is structured as follows. I describe my numerical approaches for cosmological simulations of FDM in section 4.3. In section 4.4, I turn to an analysis of the results for the matter power spectrum, while in section 4.5 I report the findings for the halo mass function. Section 4.6 discusses halo profiles, and section 4.7 the differences expected in FDM due to modifications of the initial linear theory spectrum relative to ΛCDM if this is self-consistently computed. Finally, I give my conclusions in section 4.8.

4.2 Theoretical background

Since the theoretical background and fundamental equations relevant for FDM have already been covered in section 2.1, the corresponding commentary in May and Springel (2021) will not be reiterated here and the reader is referred to section 2.1.

4.3 Numerical methodology

My simulations are performed within the same framework as ordinary cosmological ΛCDM simulations. The simulation volume consists of a cubic box of side length L with periodic boundary conditions, which is intended to sample the matter distribution in the universe. The volume is filled with matter whose average comoving density is the mean background density

$$\langle \rho_\mathrm{c} \rangle = \Omega_\mathrm{m} \rho_\mathrm{crit} = \Omega_\mathrm{m} \frac{3H_0^2}{8\pi G}. \tag{4.1}$$

In order to solve eqs. (2.3) and (2.4), a 2nd-order symmetrized split-step pseudo-spectral Fourier method ("kick–drift–kick") is employed. For a small time step Δt, the time evolution can be approximated as follows (Woo and Chiueh 2009; Edwards et al. 2018):

$$\begin{aligned} &\psi_\mathrm{c}(t + \Delta t, \vec{x}) \\ &= \mathcal{T}\, e^{-i \int_t^{t+\Delta t} \left(-\frac{\hbar}{2m} \frac{1}{a(t')^2} \nabla_\mathrm{c}^2 + \frac{m}{\hbar} \frac{1}{a(t')} \Phi_\mathrm{c}(t', \vec{x}) \right) \mathrm{d}t'} \psi_\mathrm{c}(t, \vec{x}) \end{aligned} \tag{4.2}$$

$$\approx e^{i\frac{\Delta t}{2}\left(\frac{\hbar}{m}\frac{1}{a(t)^2}\nabla_c^2-\frac{m}{\hbar}\frac{1}{a(t)}\Phi_c(t+\Delta t,\vec{x})-\frac{m}{\hbar}\frac{1}{a(t)}\Phi_c(t,\vec{x})\right)}\psi_c(t,\vec{x}) \tag{4.3}$$

$$\approx \underbrace{e^{-i\frac{m}{\hbar}\frac{1}{a(t)}\frac{\Delta t}{2}\Phi_c(t+\Delta t,\vec{x})}}_{\text{“kick”}}\ \underbrace{e^{i\frac{\hbar}{m}\frac{1}{a(t)^2}\frac{\Delta t}{2}\nabla_c^2}}_{\text{“drift”}}\ \underbrace{e^{-i\frac{m}{\hbar}\frac{1}{a(t)}\frac{\Delta t}{2}\Phi_c(t,\vec{x})}}_{\text{“kick”}}\ \psi_c(t,\vec{x}) \tag{4.4}$$

where $\mathcal{T}$ is the time ordering operator and, using the Baker–Campbell–Hausdorff formula, the time evolution operator has been split into three parts which do not mix functions of the position and derivative operators.

The fields ψ and Φ are discretized on a uniform Cartesian grid with N^3 grid points to allow for numerical computations using the FFT. Accordingly, the numerical algorithm of the pseudo-spectral method performs the following operations:

- $\psi_c \leftarrow e^{-i\frac{m}{\hbar}\frac{1}{a}\frac{\Delta t}{2}\Phi_c}\psi_c$ (kick) (4.5a)
- $\psi_c \leftarrow \mathrm{FFT}^{-1}\left(e^{-i\frac{\hbar}{m}\frac{1}{a^2}\frac{\Delta t}{2}k^2}\,\mathrm{FFT}(\psi_c)\right)$ (drift) (4.5b)
- $\Phi_c \leftarrow \mathrm{FFT}^{-1}\left(-\frac{1}{k^2}\,\mathrm{FFT}\left(4\pi G m\left(|\psi_c|^2-\langle|\psi_c|^2\rangle\right)\right)\right)$ (update potential) (4.5c)
- $\psi_c \leftarrow e^{-i\frac{m}{\hbar}\frac{1}{a}\frac{\Delta t}{2}\Phi_c}\psi_c$ (kick) (4.5d)
- Go to algorithm (4.5a) (4.5e)

Consecutive executions of algorithms (4.5a) and (4.5d) (i. e. except for the initial and final time steps) can be combined into a single operation $\psi_c \leftarrow e^{-i\frac{m}{\hbar}\Delta t\Phi_c}\psi_c$ to improve performance.

The choice of the time step Δt in algorithms (4.5a) to (4.5e) is determined by the requirement that the phase difference in the exponentials must not exceed 2π, at which point the time step would be incorrectly “aliased” to a smaller time step corresponding to the phase difference subtracted by a multiple of 2π due to the periodicity of the exponential function. Both the kicks (algorithms (4.5a) and (4.5d)) and the drift (algorithm (4.5b)) yield separate constraints for Δt, which must be simultaneously fulfilled. The resulting time step criterion is

$$\Delta t < \min\left(\frac{4}{3\pi}\frac{m}{\hbar}a^2\Delta x^2,\ 2\pi\frac{\hbar}{m}a\frac{1}{|\Phi_{c,\max}|}\right), \tag{4.6}$$

where $\Delta x = L/N$ is the spatial resolution and $\Phi_{c,\max}$ is the maximum value of the potential. The constraint involving the resolution Δx stems from the drift operation, while the constraint with the potential $\Phi_{c,\max}$ results from the kick operation.

The dependence $\Delta t \propto \Delta x^2$ seems to be typical for all numerical approaches to the Schrödinger–Poisson system of eqs. (2.3) and (2.4), and can be viewed as a reflection of the relation of the Schrödinger equation to diffusion problems.

Another constraint on the validity of the discretization becomes apparent when considering the velocity field (eq. (2.10)). Since this is given by the *gradient* of the wave function's phase, whose difference between two points can be at most 2π, it follows that the discretized velocity field cannot exceed a maximum value (depending on the concrete form of the discretized gradient operator) of about

$$v_{\mathrm{max}} = \frac{\hbar}{m}\frac{\pi}{\Delta x}. \tag{4.7}$$

Velocities $v \geq v_{\mathrm{max}}$ cannot be represented in a simulation with resolution Δx. Comparing to eq. (2.13), it becomes apparent that this statement is equivalent to a constraint on the resolution, which should be good enough to resolve the de Broglie wavelength of the largest velocities:

$$\Delta x < \frac{\pi\hbar}{m v_{\mathrm{max}}} = \frac{1}{2}\lambda_{\mathrm{dB}}(v_{\mathrm{max}}). \tag{4.8}$$

The requirement to resolve structures on the scale of the de Broglie wavelength, given by eq. (4.8), in combination with the time step requirement $\Delta t \propto \Delta x^2$ (eq. (4.6)), is one of the main reasons why FDM simulations require so many more computational resources than traditional particle-based CDM simulations.

To enable my simulation work, I have implemented the pseudo-spectral algorithm (algorithms (4.5a) to (4.5e)) as a module in the AREPO code (Springel 2010; Weinberger, Springel, and Pakmor 2020), which has been dubbed AxiREPO, in a similar approach to Mocz et al. (2017) and Mocz et al. (2020). Algorithm (4.5c) is performed using AREPO's existing Poisson solver algorithm, while the split-step solution to the Schrödinger equation is solved using newly developed, highly parallel code inspired by the same algorithm. It, too, makes use of the "Fastest Fourier Transform in the West (FFTW)" library (Frigo and Johnson 2005) to perform the FFT.

This new implementation enables simulations of grids which are several orders of magnitude larger than any previous work. The split-step scheme and the integration with AREPO also make it suitable for future simulations including baryons. All FDM simulations reported in this chapter were performed using AxiREPO, whereas the CDM simulations were done using AREPO's standard TreePM method.

4.3.1 Initial conditions

As mentioned previously, there is in fact a correspondence between the Vlasov–Poisson equations which describe the evolution of a CDM fluid and the Schrödinger–Poisson equations, with the latter reducing to the former for $\hbar/m \to 0$ (Widrow and Kaiser 1993; Mocz et al. 2018; Garny, Konstandin, and Rubira 2020). A wave function $\psi(\vec{x})$ can be constructed from a phase space distribution function $f(\vec{x}, \vec{p})$ by means of a Wigner or Husimi transform. In this case, the correspondence allows one to "translate" a phase space distribution function to the wave function formalism. This is very useful because it enables direct comparisons between structure formation with CDM and FDM from the same ICs.

In general, a wave function on a discretized lattice can be constructed using the prescription eq. (2.15) (cf. Widrow and Kaiser 1993; Mocz et al. 2018). For the case of a "cold" or "single-stream" distribution function as in the case of CDM, where each point in space has a single well-defined value for the velocity, the construction of the wave function simply reduces to the polar decomposition (eq. (2.9)), with the absolute value and phase determined by the density and velocity at each point, respectively,via eqs. (2.16) and (2.17). Equation (2.17) can be easily solved numerically by applying the spectral method to eq. (2.18).

The ICs were generated using the N-GenIC code (Springel 2015), which employs the Zel'dovich approximation to generate a perturbed particle distribution, for an ordinary ΛCDM cosmological simulation at the starting redshift $z = 1/a - 1 = 127$ with an input power spectrum following Efstathiou, Sutherland, and Maddox (1990) and Efstathiou, Bond, and White (1992), i. e. of the form

$$P(k) \propto k\left(1 + \left(ak + (bk)^{\frac{3}{2}} + c^2k^2\right)^{\nu}\right)^{-\frac{2}{\nu}} \tag{4.9}$$

with $a = 6.4/\Gamma\, h^{-1}\,\mathrm{Mpc}$, $b = 3.0/\Gamma\, h^{-1}\,\mathrm{Mpc}$, $c = 1.7/\Gamma\, h^{-1}\,\mathrm{Mpc}$, $\Gamma = \Omega_\mathrm{m} h = 0.21$,[1] and $\nu = 1.13$. A wave function was constructed from the same ICs using the procedure in eqs. (2.16) and (2.17).

4.3.2 Resolution and convergence tests

Before using the newly-developed FDM AREPO module AxiREPO for simulations of larger cosmological boxes, extensive tests were performed on boxes of comoving size

[1] See section 4.3.4.

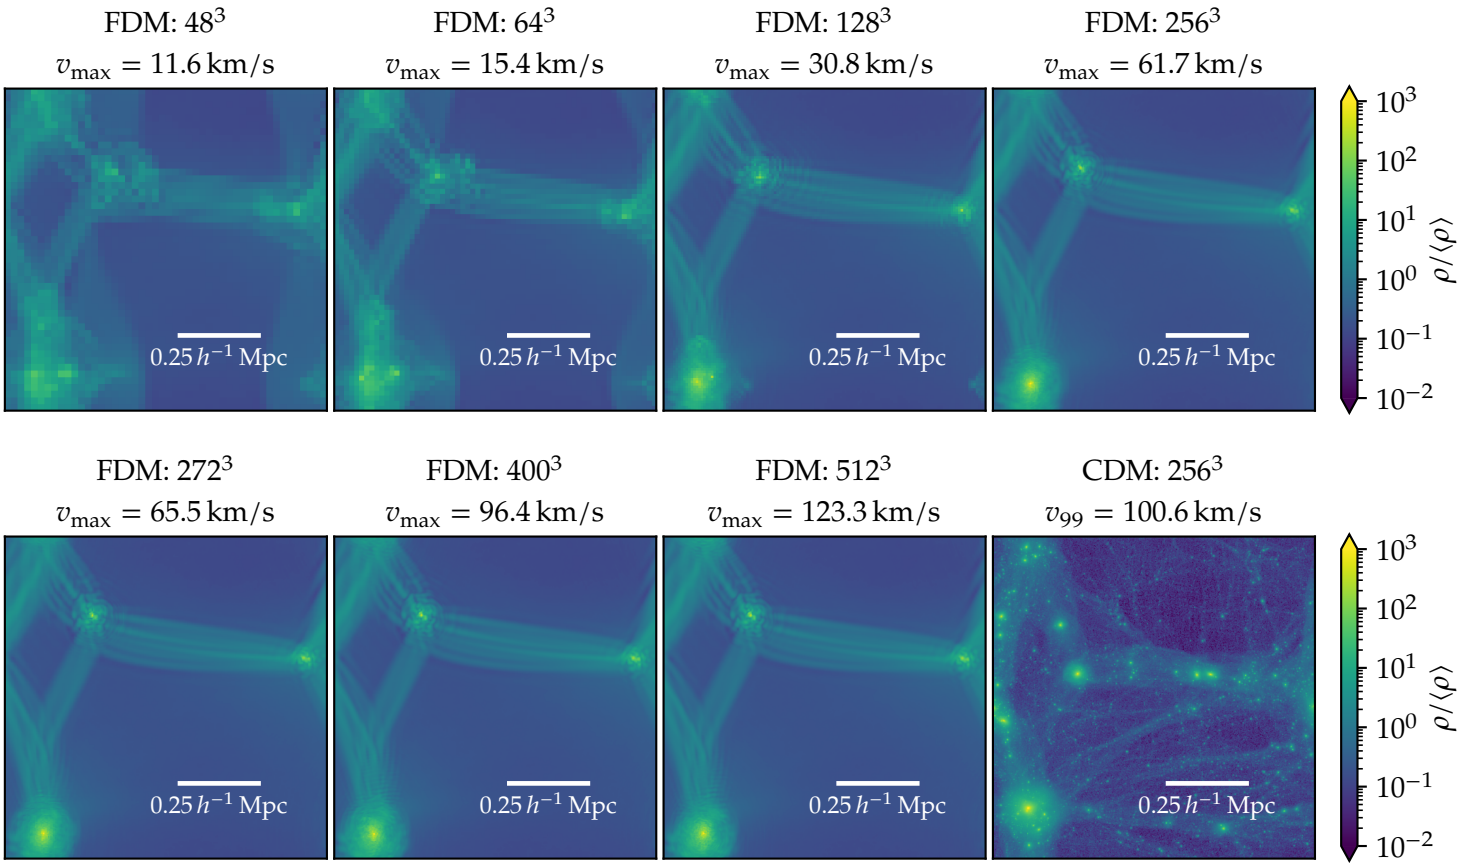

Figure 4.1: Projected dark matter density at $z = 0$ in $L = 1\,h^{-1}\,\mathrm{Mpc}$ cosmological box simulations of FDM with $mc^2 = 1.75 \times 10^{-23}\,\mathrm{eV}$ and CDM ICs for different resolutions. A high-resolution CDM simulation is shown for comparison. The 99th-percentile velocity of the particles in the CDM simulation is given as v_{99}.

$L = 1\,h^{-1}\,\mathrm{Mpc}$.[2] This not only allowed for verification of the code's correctness by comparing to CDM simulations with the same ICs and to other implementations of FDM, but also to study the behavior of the pseudo-spectral method. Of particular interest are the resolution requirements, e. g. to what extent the numerical results remain valid and the convergence of the matter power spectrum is compromised when the velocity constraint (eq. (4.7)) is violated. Moreover, the differences to the behavior of CDM simulations are of interest even in this small test volume. Although the simulation volume is too small to be cosmologically representative at $z = 0$ since even the largest scales become non-linear by then, there are opportunities to observe what differences (or similarities) become apparent in both the linear (early times)

[2]The cosmological parameters are the same as in section 4.3.4.

and non-linear (late times) regimes of both dark matter models.

Figure 4.1 shows the projected density at $z = 0$ of such test simulations across a wide range of resolutions, with a CDM simulation for comparison.[3] Notably, even at the lowest resolution, the qualitative features are preserved and structures on the largest scales remain the same. However, for the lower resolutions, structures become increasingly "smeared out" (filling a larger volume), and the resulting range of values in the density contrast is decreased, indicating that the lack of resolution interferes with the formation of more compact structures. On the other hand, for the higher resolutions, there are no appreciable discrepancies in the low-density regions and filaments, with the only differences being slight changes in the position and internal structure of the largest halos. This is in line with the idea of the velocity criterion eq. (4.7) and indicates that simulations can yield valid answers in regions with low velocities even in the presence of high-velocity regions which violate the criterion. Common to all the FDM simulations, and in contrast to the CDM simulation, almost all small-scale structure is erased due to the Heisenberg uncertainty principle with a large de Broglie wavelength ($\lambda_{\mathrm{dB}} = 1.21\,\mathrm{kpc}$ for $mc^2 = 10^{-22}\,\mathrm{eV}$, $v = 100\,\mathrm{km\,s^{-1}}$). The FDM structures consist of a few massive halos and smooth overdense filaments, while in CDM, the filaments fragment into sub-halos down to the smallest scales.

A more quantitative indication of numerical convergence is given in fig. 4.2 by the power spectra of the density field. This demonstrates that, for the given setup, grid sizes of 128^3 and smaller lead to significant deviations from higher-resolution results even on large scales. This discrepancy is visible even in the density projections of fig. 4.1. On the other hand, while it becomes difficult to see any visual differences for grid sizes of 256^3 and larger, the power spectra show that, to some degree, even the 256^3 grid suffers from lack of power on smaller scales. In this case, at least roughly, the velocity criterion seems to give a rather good indication of the resolution required to achieve convergence of the power spectrum on most scales (cf. the values of v_{max} in fig. 4.1). Generally, the (relative) lack of power is most pronounced and persists for higher resolutions on smaller scales, i. e. the power spectrum converges progressively from larger to smaller scales with increasing resolution, although even convergence on just the largest scales places significant demands on resolution (better than a 128^3 grid in this case).

[3] Due to the fact that at $z = 0$, all scales have become non-linear with such a small box size, the largest-scale structures parallel to the coordinate axes ("fundamental modes", i. e. the modes with the minimal value of $k = 2\pi/L$) are clearly visible.

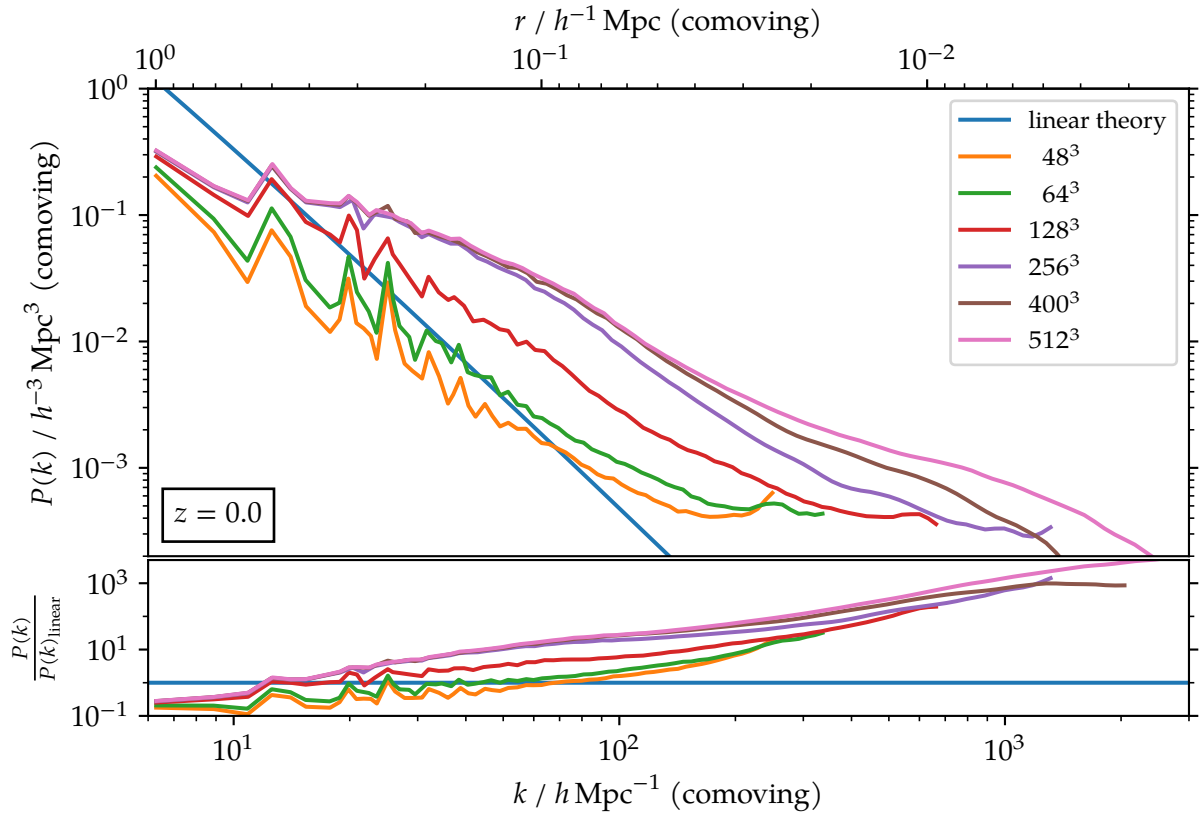

Figure 4.2: Dark matter power spectra for cosmological FDM simulations with different resolutions, box size $L = 1\,h^{-1}\,\text{Mpc}$, and FDM mass $mc^2 = 1.75 \times 10^{-23}\,\text{eV}$, at $z = 0$ evolved from CDM ICs. The legend indicates the grid size for each included simulation. The power spectrum evolved using linear perturbation theory is shown for comparison. The bottom panel shows the ratio of the power spectra to the result from linear theory.

4.3.3 Computational cost analysis

As mentioned before, the required computational resources – both run-time cost and memory – are a significant obstacle when attempting to perform large-scale simulations of FDM. The run-time cost depends on three simulation parameters: the FDM mass m, the (comoving) box size L and the number of grid points N^3 (where N is the number of points per dimension).[4]

For each individual time step, the operations in algorithms (4.5a) to (4.5e) are performed. The cost of a single time step only depends on the grid size N^3, while m and L only enter the computation through multiplication as constant factors. It is dominated by the cost of the FFT operations, which scales as $O(N^3 \ln(N^3))$. The

[4]Instead of N, the resolution $\Delta x = L/N$ can be equivalently considered.

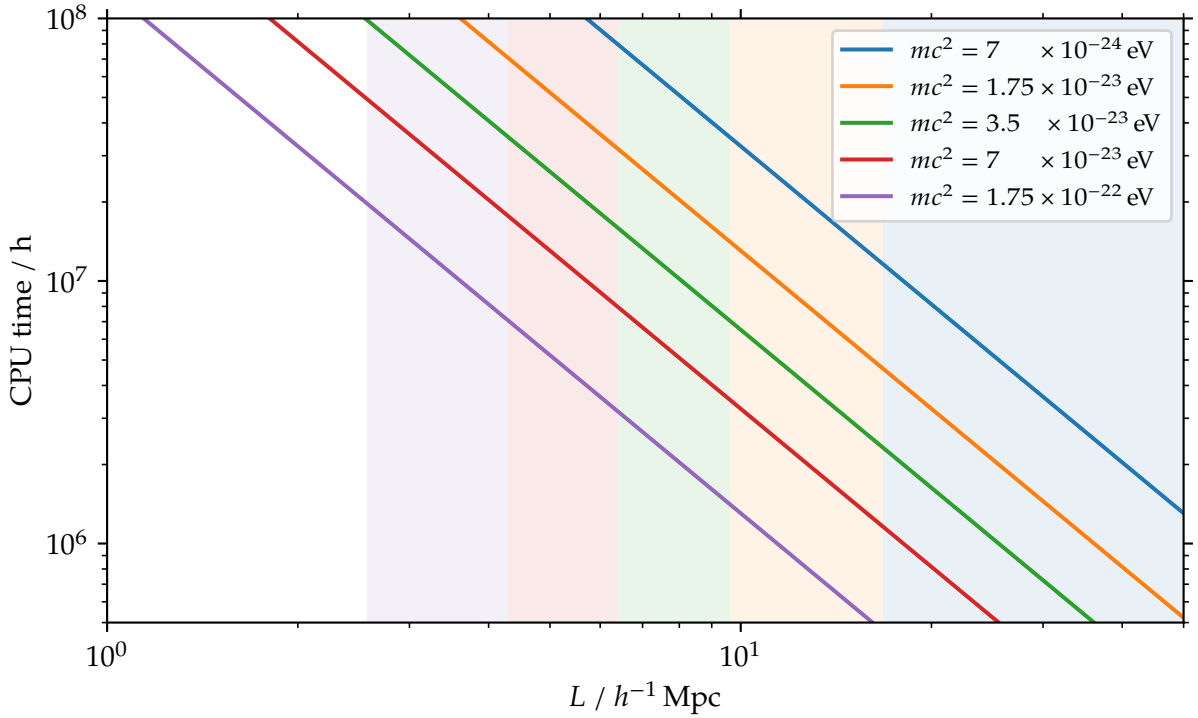

Figure 4.3: The expected computational cost of cosmological simulations when running to $z = 0$ as a function of box size L for a grid size of $N^3 = 8096^3$ and different FDM masses. Shaded areas indicate regions where the velocity constraint eq. (4.7) is violated for the corresponding mass at $z = 0$ due to lack of resolution. The concrete values of CPU time correspond to the Cobra cluster at MPCDF for reference.

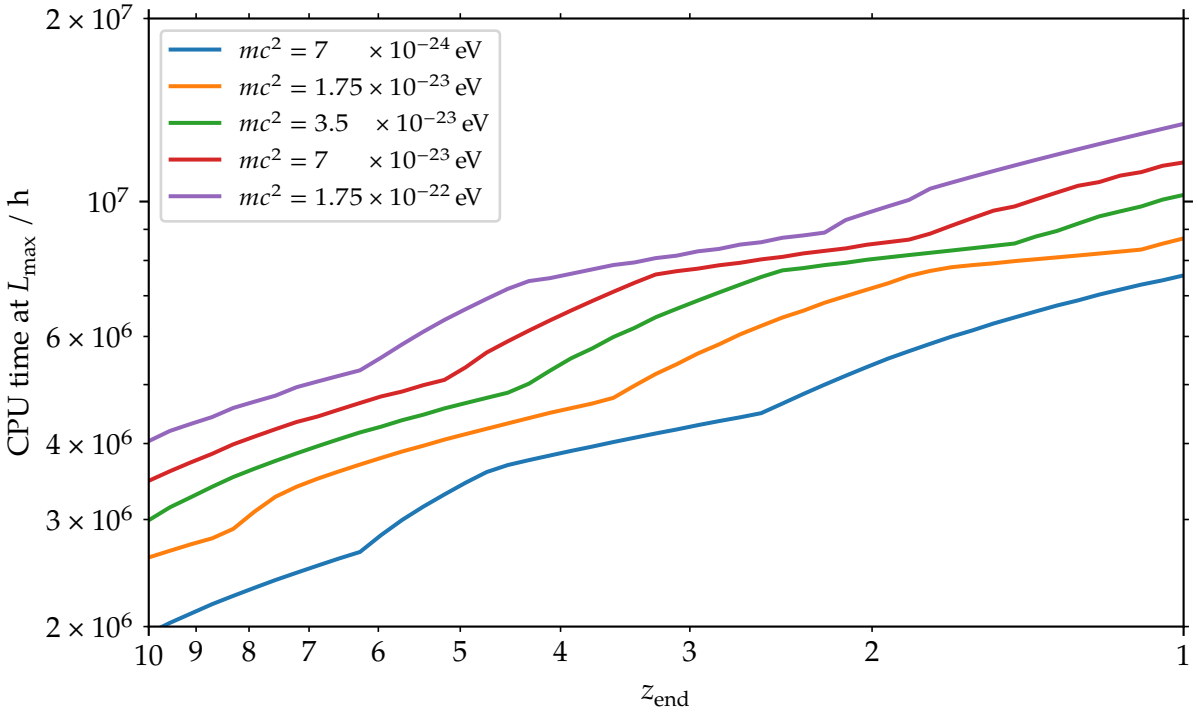

Figure 4.4: The expected computational cost of cosmological simulations with different FDM masses as a function of target redshift z_{end} for a grid size of $N^3 = 8096^3$ at the maximum allowed box size $L = L_{\max}$, as given by the velocity constraint eq. (4.7). The points where the curves intersect their associated shaded regions in fig. 4.3 would correspond to $z_{\text{end}} = 0$ in this figure.

total cost then depends on the number of time steps which have to be performed, which is determined by the time step size given by eq. (4.6). The drift criterion $\propto \Delta x^2$ is typically more stringent, especially in light of the resolution requirements (e. g. eq. (4.7)) and for higher redshifts. Thus, the total number of time steps is roughly inversely proportional to the limit given by the drift criterion, which results in the following approximate behavior of the computation time C:

$$C \propto \frac{N^5 \ln(N^3)}{mL^2}. \tag{4.10}$$

This relation was also verified in practice using measurements of small test computations.

Particularly striking is the strong dependence of eq. (4.10) on N^5, which rapidly leads to exploding computational costs for larger N. At fixed N, the cost can be decreased by choosing a larger box size L because this reduces the resolution Δx. However, as mentioned before, the resolution requirements are quite strict due to the inherent scale in the Schrödinger–Poisson system, which must be resolved in order to yield reliable results.

To give a tangible reference for the absolute costs involved on the one hand, but also the code's performance on the other hand, fig. 4.3 shows the expected cost of simulations with grid size $N^3 = 8096^3$ to $z = 0$ in CPU hours on the state-of-the-art Cobra supercomputer at the Max Planck Computing and Data Facility (MPCDF).[5] Figure 4.3 displays the CPU time for a range of masses m with variable box size L. Even when allowing for the velocity criterion eq. (4.7) to be violated, it becomes clear that such a simulation would take at least 10^7 CPU h. Even worse, the resolution requirements severely limit the possible simulation volume to $L < 10\,h^{-1}$ Mpc for the mass range of interest. Since the de Broglie wavelength increases with smaller masses, the resolution requirements can be loosened by choosing much smaller values for m, but this does not nearly go far enough as the physically viable range is bounded around $\gtrsim 10^{-22}$ eV. The computational cost for the largest possible box size $L_{\max}$ at a given mass m (intersections of the graphs with the vertical lines) decreases only weakly with decreasing m.

Figure 4.4 shows the same computation time at the largest possible box size due to eq. (4.7) for each mass, but now for the case where the simulation is stopped earlier

[5] `https://www.mpcdf.mpg.de/services/supercomputing/cobra`

Table 4.1: List of performed simulations with important characteristics. The lengths given for the box sizes and resolutions are comoving.

Type	Res. el.	L / h^{-1} Mpc	mc^2 / eV	Resolution
FDM	8640^3	10	7×10^{-23}	$1.16\,h^{-1}$ kpc
FDM	4320^3	10	$(3.5, 7) \times 10^{-23}$	$2.31\,h^{-1}$ kpc
FDM	3072^3	10	$(3.5, 7) \times 10^{-23}$	$3.26\,h^{-1}$ kpc
FDM	2048^3	10	$(3.5, 7) \times 10^{-23}$	$4.88\,h^{-1}$ kpc
FDM	4320^3	5	7×10^{-23}	$1.16\,h^{-1}$ kpc
FDM	3072^3	5	$(3.5, 7) \times 10^{-23}$	$1.63\,h^{-1}$ kpc
FDM	2048^3	5	$(3.5, 7) \times 10^{-23}$	$2.44\,h^{-1}$ kpc
FDM	1024^3	5	$(3.5, 7) \times 10^{-23}$	$4.88\,h^{-1}$ kpc
CDM	2048^3	10	—	$9.69 \times 10^3\,h^{-1}\,M_\odot$
CDM	1024^3	10	—	$7.75 \times 10^4\,h^{-1}\,M_\odot$
CDM	512^3	10	—	$6.20 \times 10^5\,h^{-1}\,M_\odot$
CDM	1024^3	5	—	$9.69 \times 10^3\,h^{-1}\,M_\odot$
CDM	512^3	5	—	$7.75 \times 10^4\,h^{-1}\,M_\odot$

than $z = 0$.[6] Unfortunately, the computational cost does not depend very strongly on the final redshift (approximately $\propto z^{-1/2}$ in the region shown); it is merely lowered by a factor of $\lesssim 5$ even when stopping at $z = 6$, and still remains at several million CPU hours for a $N^3 = 8096^3$ simulation.

4.3.4 Simulations

All of my simulations were performed using the cosmological parameters $\Omega_\mathrm{m} = 0.3$, $\Omega_\mathrm{b} = 0$, $\Omega_\Lambda = 0.7$, $H_0 = 70\,\mathrm{km\,s^{-1}\,Mpc^{-1}}$ ($h = 0.7$), and $\sigma_8 = 0.9$,[7] with ICs as described in section 4.3.1. The same random seed was shared for the generation of all ICs at $z = 127$ in order to allow for a direct comparison between different dark matter models and resolutions. In order to avoid the onset of resolution effects that would affect even my highest-resolution simulations, the simulations were run until $z = 3$. For comoving box sizes of $5\,h^{-1}$ Mpc and $10\,h^{-1}$ Mpc, and masses mc^2 of 3.5×10^{-23} eV and 7×10^{-23} eV, simulations with different resolutions – up to grid sizes of $N^3 = 8640^3$ – were performed. A detailed list of the different simulations is given in table 4.1.

[6] The unevenness in the lines results due to AREPO's time binning procedure, which "discretizes" the allowed time step values.

[7] As usual, the density parameters Ω_i indicate the values at $z = 0$.

To date, these are the largest three-dimensional cosmological simulations of structure formation including the full FDM dynamics to low redshifts, with a simulation volume which is at least 64 times larger and a grid which has at least 600 times more resolution elements than any comparable existing work. The largest simulation required more than 7×10^6 CPU h on the Cobra cluster to complete. Previous efforts (using various methods) have only reached box sizes of $2.5\,h^{-1}$ Mpc (Mina, Mota, and Winther 2022, using an AMR method, until $z = 2.5$) or $1.7\,h^{-1}$ Mpc (Mocz et al. 2020, also with a pseudo-spectral method, until $z = 5.5$) for full FDM simulations, and $2.5\,h^{-1}$ Mpc with a hybrid method (Veltmaat, Niemeyer, and Schwabe 2018, i. e. not including the FDM dynamics everywhere in the simulation volume) for similar cosmological simulations.

4.3.5 Grid-based halo finding

The halo mass function (HMF) is an important measure of large-scale structure. For typical CDM simulations, it is determined using an algorithm like friends-of-friends (FoF) in AREPO, which identifies halos by connecting simulation particles (point masses) whose distance to other particles is below a certain threshold. When using the pseudo-spectral method, however, there are no particles, and the density field is instead represented by a Cartesian grid. This means that widely-used algorithms, which operate on the particle distribution, cannot be used to analyze the FDM simulations presented here.

Because of this, a modified version of the FoF algorithm was developed for this dissertation to enable the determination of the HMF for a discretized density grid. Instead of a linking length, this grid-based halo finder uses a density threshold as a parameter. Adjacent cells in the grid are linked if their density exceeds the density threshold. The grid-based halo finder is part of the pseudo-spectral FDM AREPO module AxiREPO.

Figure 4.5 demonstrates that this new halo finder performs well in comparison to the standard particle-based FoF algorithm. Using a CDM simulation with 512^3 particles, the density was represented both using the original particle data and a Cartesian grid with 1024^3 grid points whose density values were determined using cloud-in-cell (CIC) mass assignment. The HMF was determined with the particle and grid data as input for FoF and the new grid-based halo finder, respectively. As evident in fig. 4.5, both procedures show excellent agreement. Differences arise in the lightest halos, which are limited by resolution since density is always spread out

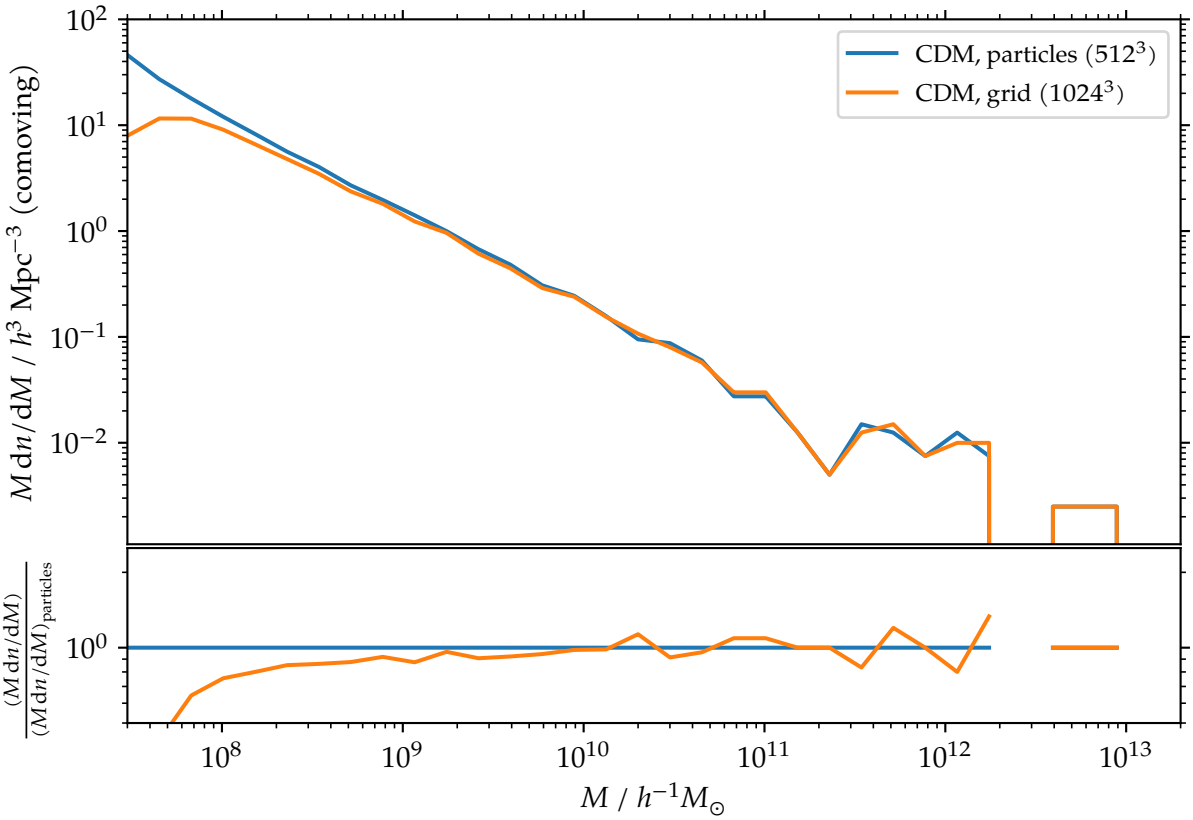

Figure 4.5: The HMF of a cosmological ΛCDM simulation at $z = 0$ with $L = 10\,h^{-1}$ Mpc and 512^3 particles, as determined using AREPO's standard FoF algorithm ("particles") and the newly implemented grid-based halo finder ("grid"). The latter used the density grid constructed from a CIC mass assignment of the simulation particles onto a grid with 1024^3 points as input, with an overdensity threshold of 60 times the matter background density. The bottom panel shows the ratio of the mass functions. Note that the masses given here are the *masses of the FoF groups* (i. e. the sum of all particle or grid cell masses in the group) to give a direct comparison of the two approaches.

over at least one cell volume, and slight variations are present for the most massive halos, where statistical effects play a role since there are only a few halos per mass bin.

4.4 Dark matter power spectrum

The matter power spectrum gives important insight into how matter clusters at different length scales, and its evolution in time is by now quite accurately known for ΛCDM (e. g. Jenkins et al. 1998), even when baryonic effects are included (Springel

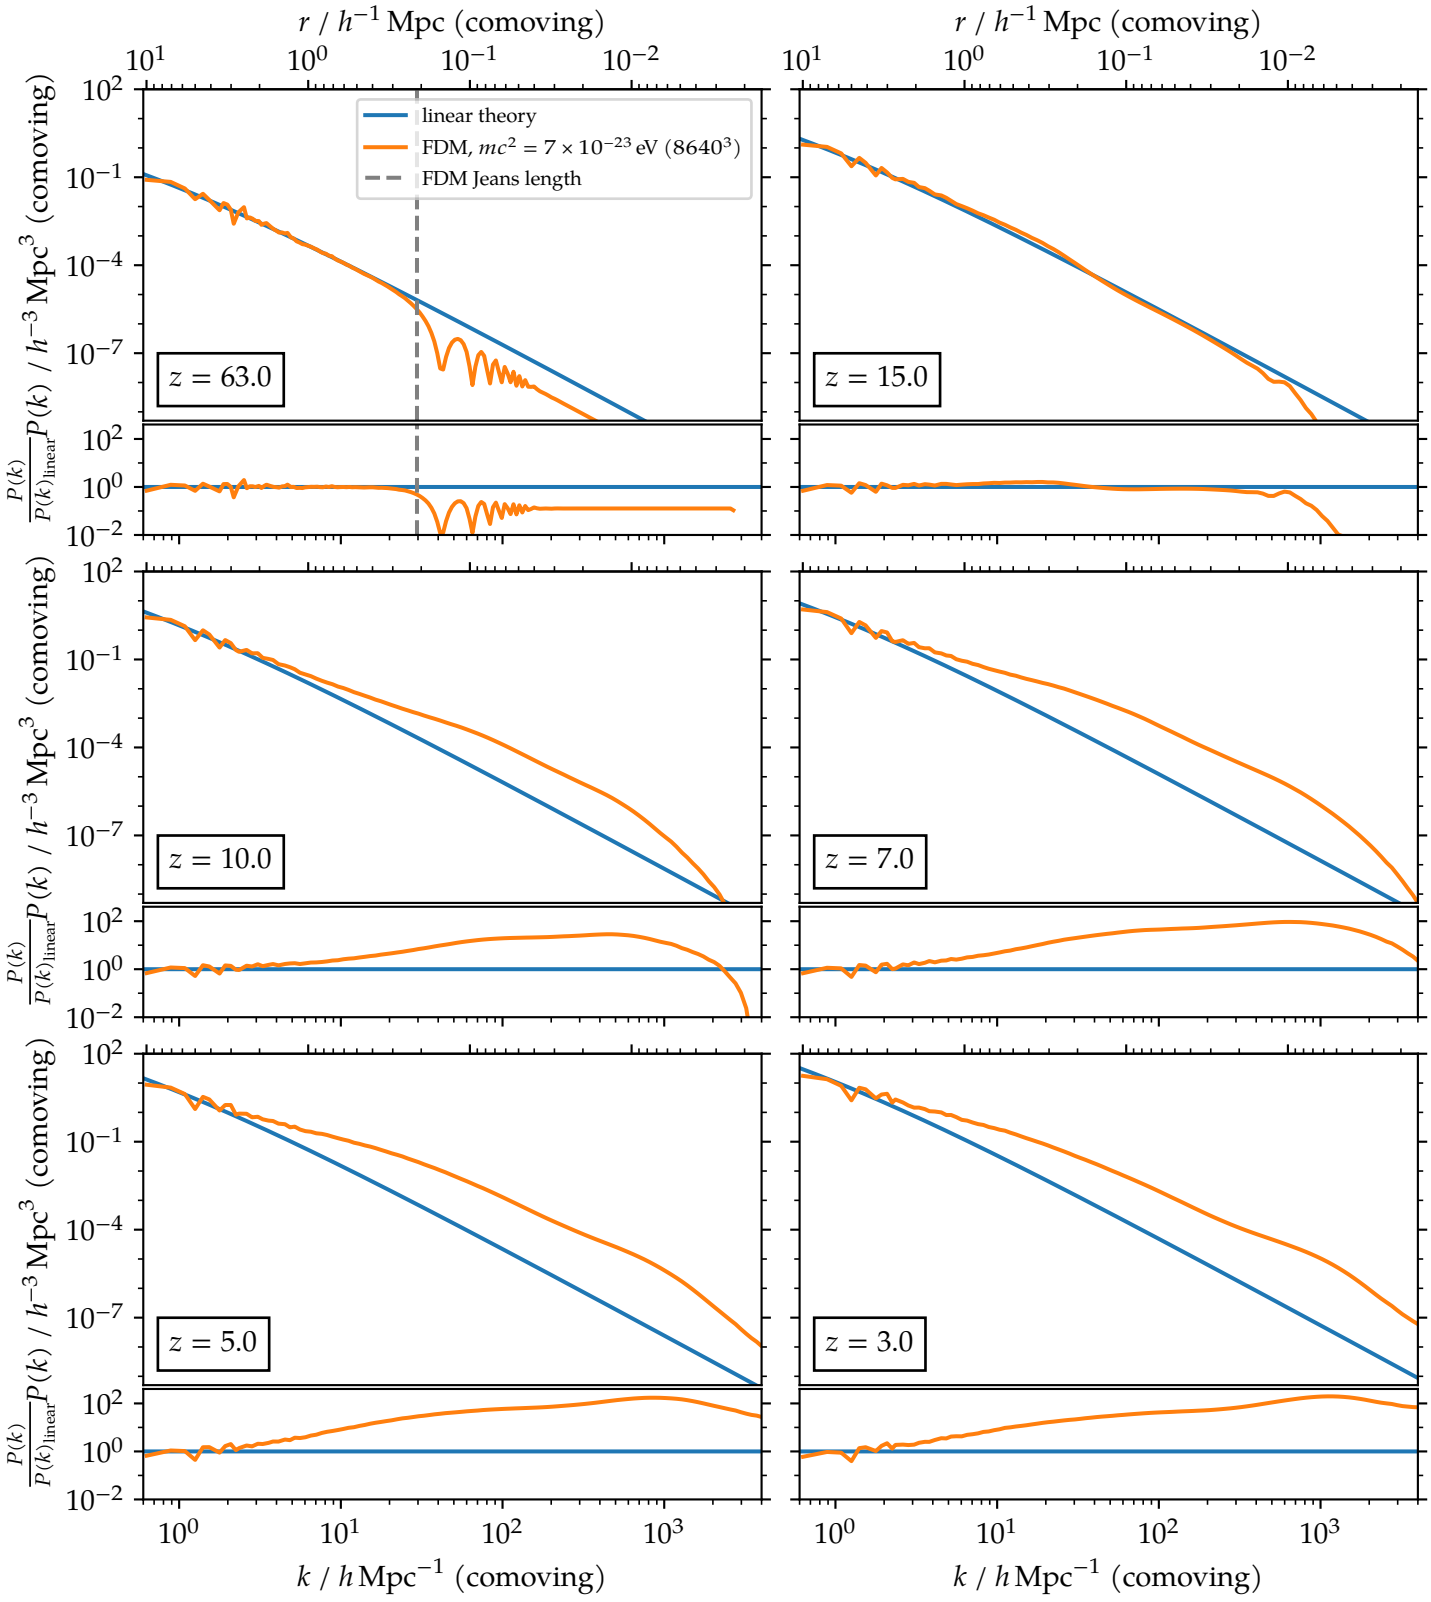

Figure 4.6: Dark matter power spectra at different redshifts for a high-resolution cosmological FDM simulation with box size $L = 10\,h^{-1}$ Mpc, FDM mass $mc^2 = 7 \times 10^{-23}$ eV, grid size $N^3 = 8640^3$, and CDM ICs. The power spectrum evolved using linear perturbation theory is shown for comparison. The lower panels show the ratio of the power spectra to the result from linear theory. For $z = 63$, the dashed line additionally indicates the FDM Jeans scale (eq. (2.14)).

et al. 2018). The power spectrum is also a particularly important diagnostic for a comparison of FDM with CDM, because the former is expected to suppress structure formation at scales smaller than λ_{dB} due to the Heisenberg uncertainty principle. Due to the use of CDM ICs, my results demonstrate to what extent the dynamics of the Schrödinger–Poisson eqs. (2.3) and (2.4) affect the clustering of matter compared to CDM, even without the initial suppression present in a "realistic" FDM cosmology (cf. section 4.7).

Figure 4.6 shows the time evolution of the matter power spectra of FDM at the highest resolution of 8640^3 cells in a $10\,h^{-1}\,\mathrm{Mpc}$ box. From the top left to the bottom right panel, I show redshifts from $z = 63$ to $z = 3$, and compare to expectations from linear theory. In particular, the bottom of each panel gives the ratio of the measured FDM power to the linear theory power spectrum used for initializing the simulation. At high redshifts and at small k, FDM accurately follows linear theory, with modes growing independently with the same linear growth factor, just like CDM does, so that the pattern of random fluctuations of the particular realization around the smooth initial input spectrum is preserved in time. Beginning at scales around $k \approx 10\,h\,\mathrm{Mpc}^{-1}$ and a redshift of $z \approx 15$, signs of mildly non-linear evolution are apparent, which manifest themselves in a stronger than linear growth of power. This non-linear evolution becomes quickly more pronounced in a way reminiscent of CDM, except that on the smallest resolved scales, for $k \geq 1000\,h\,\mathrm{Mpc}^{-1}$, the non-linear growth appears sluggish and lagging behind that seen on larger scales.

It is now important to examine on which scales and at which times these results are quantitatively reliable. As discussed earlier, the stringent numerical requirements of FDM make this more involved than for CDM, because here even the large-scale linear growth requires a fairly high resolution to get right, and it is not readily clear how numerical limitations will manifest themselves in the results. In fig. 4.7 I first compare results for FDM simulations in a $10\,h^{-1}\,\mathrm{Mpc}$ box using different resolutions, with grid sizes from 2048^3 to 8640^3. Focusing on the three redshifts of $z = 31$, $z = 10$ and $z = 3$, one can see that the results still appear well converged at the high redshifts of $z = 31$ and $z = 10$ (although the lowest-resolution simulation with $N^3 = 2048^3$ slightly starts to lag behind at $z = 10$ on smaller scales). However, this ceases to be true at the lower redshift of $z = 3$, where the growth of the lower-resolution simulations now clearly trails behind, and this effect occurs on an extended range of scales, including fairly large ones, quite unlike in CDM, where the impact of spatial resolution limits is typically constrained to fairly small scales. In particular, it becomes clear that once the onset of non-linear evolution is not properly resolved anymore in a simulation,

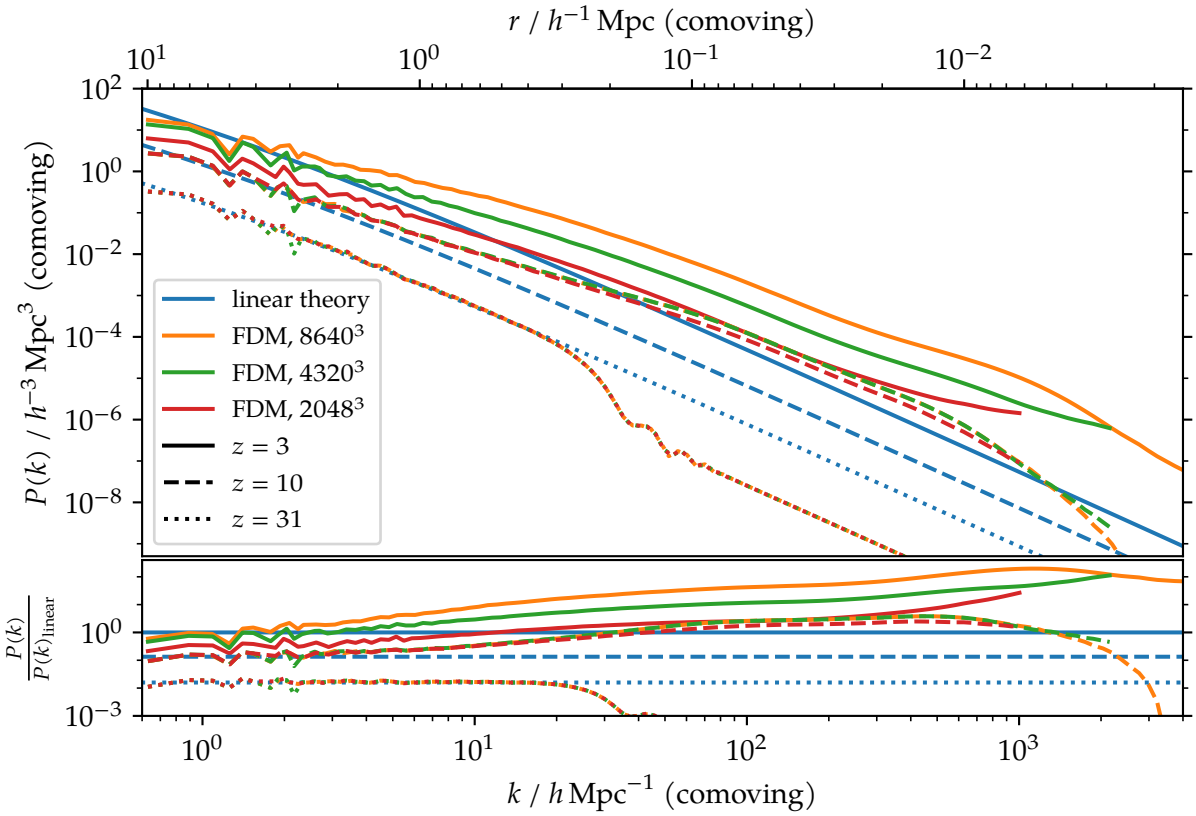

Figure 4.7: Dark matter power spectra for cosmological FDM simulations with box size $L = 10\,h^{-1}$ Mpc, FDM mass $mc^2 = 7 \times 10^{-23}$ eV and CDM ICs at different resolutions (indicated by different colors) and different redshifts (indicated by solid, dashed, and dotted lines). The power spectrum evolved using linear perturbation theory is shown for comparison. The bottom panel shows the ratio of the power spectra to the result from linear theory at $z = 3$.

the FDM spectra even stop to follow linear growth on the largest scales, making the simulations appear to be frozen in in their current state.

Further insights into these numerical effects can also be obtained by examining how they vary with box size and axion particle mass at a fixed epoch of $z = 3$. This is shown in figs. 4.8 and 4.9. Here, the impact of eq. (4.8) is clearly visible: The resolution requirements become more stringent with larger axion masses and at fixed grid size, a larger box size directly implies worse resolution. Thus, at fixed grid size, if axion mass and box size are scaled by the same factor, the resulting power spectra are identical except at small scales, which are impacted by the change in Jeans length with different axion masses. Accordingly, due to resolution effects and contrary to the physical expectation, for fixed box size and grid resolution, it is possible to see

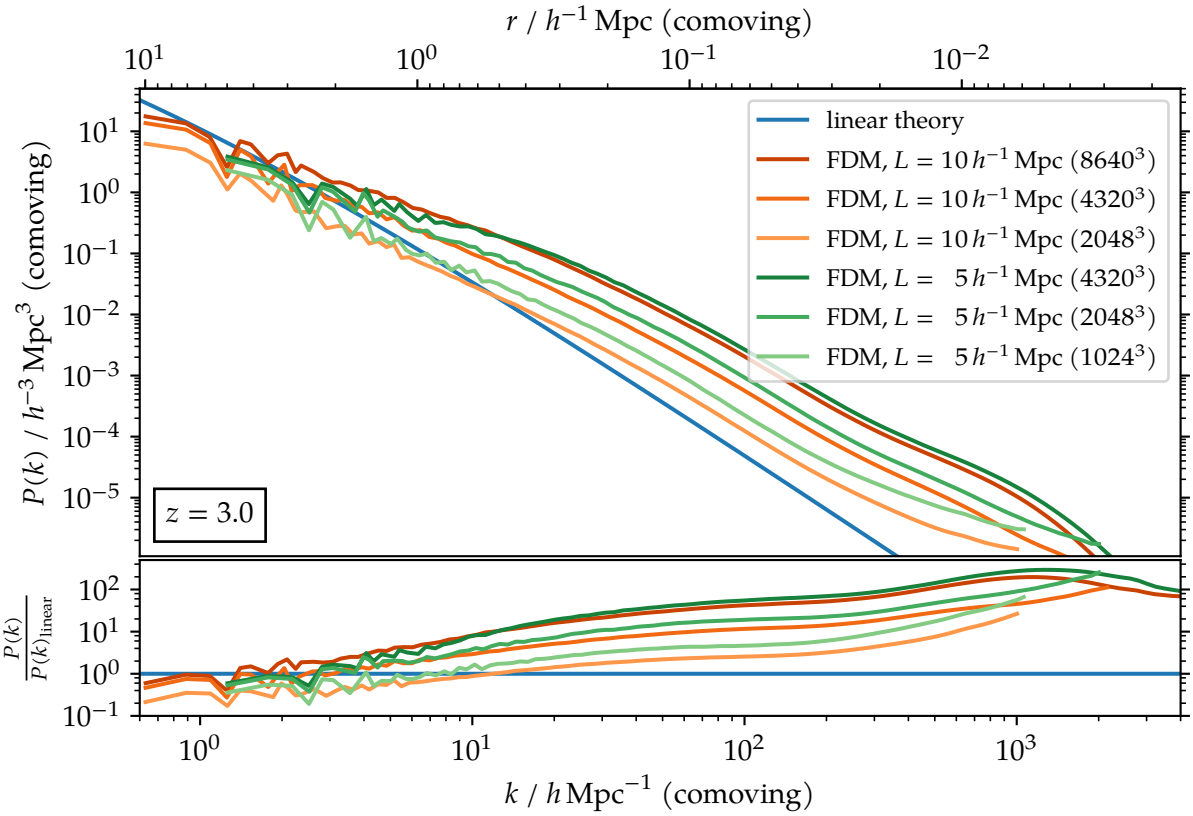

Figure 4.8: Dark matter power spectra for cosmological FDM simulations at fixed redshift ($z = 3$) with FDM mass $mc^2 = 7 \times 10^{-23}$ eV using CDM ICs, with varying box sizes ($L = 5\,h^{-1}$ Mpc and $10\,h^{-1}$ Mpc) and resolutions. The power spectrum evolved using linear perturbation theory is shown for comparison. The bottom panel shows the ratio of the power spectra to the result from linear theory.

more small-scale power when the particle mass is lowered, since lower masses have less stringent resolution requirements.

I finally come to a direct comparison of FDM with the (non-linear) CDM power spectrum in fig. 4.10. While both accurately agree with each other and with linear perturbation theory at large scales and early times, some differences start to show up as early as $z \approx 15$. The non-linear small-scale power enhancement sets in somewhat more vigorously for CDM than for FDM. For example, while mild non-linear amplification is similar at $k \approx 10\,h\,\mathrm{Mpc}^{-1}$ in CDM and FDM for $z = 15$, this effects extends to smaller scales $k \approx 100\,h\,\mathrm{Mpc}^{-1}$ in CDM, whereas FDM still pretty much tracks linear growth there. From the viewpoint of the overall time evolution, it can thus be said that the onset of non-linear structure formation is delayed for FDM, and does not proceed strictly in the same bottom-up fashion as in CDM.

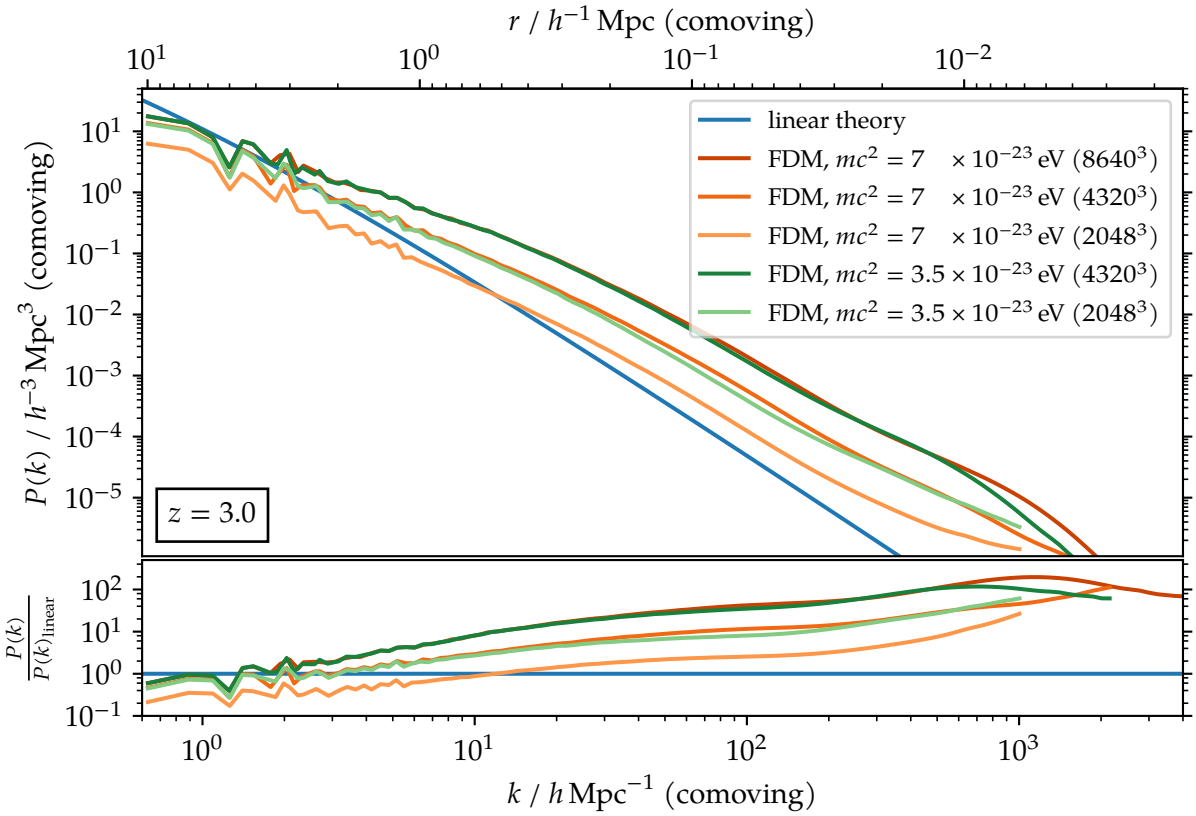

Figure 4.9: Dark matter power spectra for cosmological FDM simulations at fixed redshift ($z = 3$) with box size $L = 10\,h^{-1}\,\mathrm{Mpc}$ using CDM ICs, with varying FDM masses ($mc^2 = 3.5 \times 10^{-23}\,\mathrm{eV}$ and $7 \times 10^{-23}\,\mathrm{eV}$) and resolutions. The power spectrum evolved using linear perturbation theory is shown for comparison. The bottom panel shows the ratio of the power spectra to the result from linear theory.

Interestingly, FDM however eventually catches up in its non-linear growth and then shows a quite similar overall shape of the non-linear power spectrum. While FDM power is still somewhat suppressed at $z \leq 7$ for non-linear scales, the difference to CDM is much smaller than at earlier redshifts. This is true until a characteristic scale of around $k \approx 1000\,h\,\mathrm{Mpc}^{-1}$, where the FDM non-linear power drops significantly below the CDM power. This scale appears to be related to the "quantum pressure" effects in FDM, and is thus ultimately related to the particle mass.

Curiously, just before the non-linear FDM power decays away from CDM towards small scales, it manages to slightly exceed it. This effect could be a reflection of the interference patterns in the dark matter density field resulting from the wave-like nature of FDM, which generates transient structures of size λ_{dB} that are clearly visible

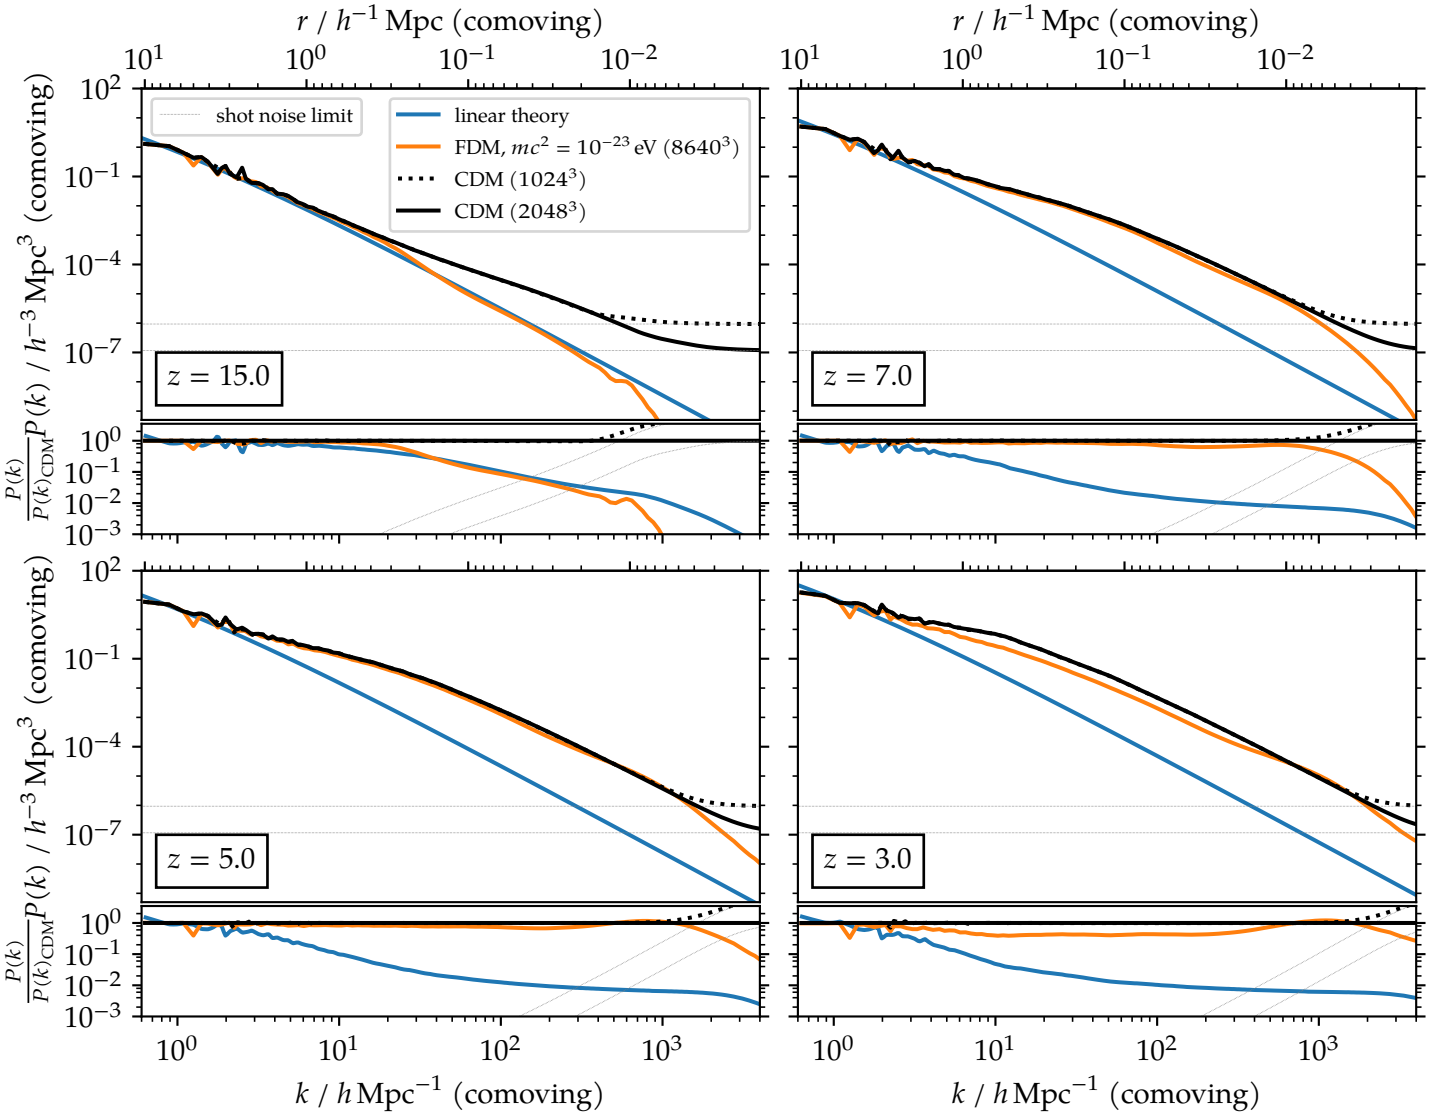

Figure 4.10: Dark matter power spectra at different redshifts for a high-resolution cosmological FDM simulation with box size $L = 10\,h^{-1}\,\mathrm{Mpc}$, FDM mass $mc^2 = 7 \times 10^{-23}\,\mathrm{eV}$, grid size $N^3 = 8640^3$, and CDM ICs, compared to cosmological N-body CDM simulations with different resolutions. The power spectrum evolved using linear perturbation theory is shown for comparison. The lower panels show the ratio of the power spectra to the result of the highest-resolution CDM simulation. For $z = 63$, the dashed line additionally indicates the FDM Jeans scale (eq. (2.14)). Faint dotted lines show the shot noise limits of the N-body simulations; the power spectrum cannot be measured accurately once it reaches this limit.

in high-resolution images that zoom in on large halos.[8] These types of order-unity density fluctuations are absent in CDM. If the associated enhancement of power localized around a characteristic scale k could be measured by some tracer sufficiently accurately, it could be a tell-tale sign of FDM.

4.5 Halo mass function

One of the most important and fundamental outcomes of structure formation are gravitationally collapsed structures of dark matter, so-called halos, which are in turn the sites of baryonic galaxy formation processes. The abundance of halos as a function of mass and epoch is thus fundamental to any cosmological model.

For the CDM scenario, the (extended) Press–Schechter formalism (Press and Schechter 1974; Sheth and Tormen 1999) has proven to be a simple and quite reliable approach to estimate the HMF and its evolution based on the linear theory power spectrum and the linear growth factor alone. Comparisons to full N-body simulations have both verified the basic approach and led to the calibration of more accurate empirical fitting formulae describing the HMF (Jenkins et al. 2001; Tinker et al. 2008; Despali et al. 2016). As a result, the HMF in pure CDM models is now accurately known and understood.

In contrast, the situation for FDM is much more murky. While a few analytic estimates have been published (Marsh and Silk 2014), based loosely on the idea of a Jeans-filtered power spectrum to account for the "quantum pressure", it is still unclear whether these approaches are quantitatively reliable. Recently, Kulkarni and Ostriker (2022) have advocated that a sharp k-space filtering in the Press–Schechter formalism with a variable cut-off may be more adequate for the FDM regime. But their predictions could only be compared to FDM mass function estimates by Schive et al. (2016), which in turn were based on an approximate technique of removing "spurious" low-mass halos in a collisionless N-body simulation with truncated initial fluctuation spectrum, similar to how warm dark matter models are often treated (Wang and White 2007). Ultimately, full non-linear simulations of the Schrödinger–Poisson system are required to obtain quantitatively reliable results. However, due to the numerical challenges involved in large-volume simulations of FDM, such determinations have not been obtained thus far.

[8] A similar effect (although in a different cosmological context) has been found in Mocz et al. (2020).

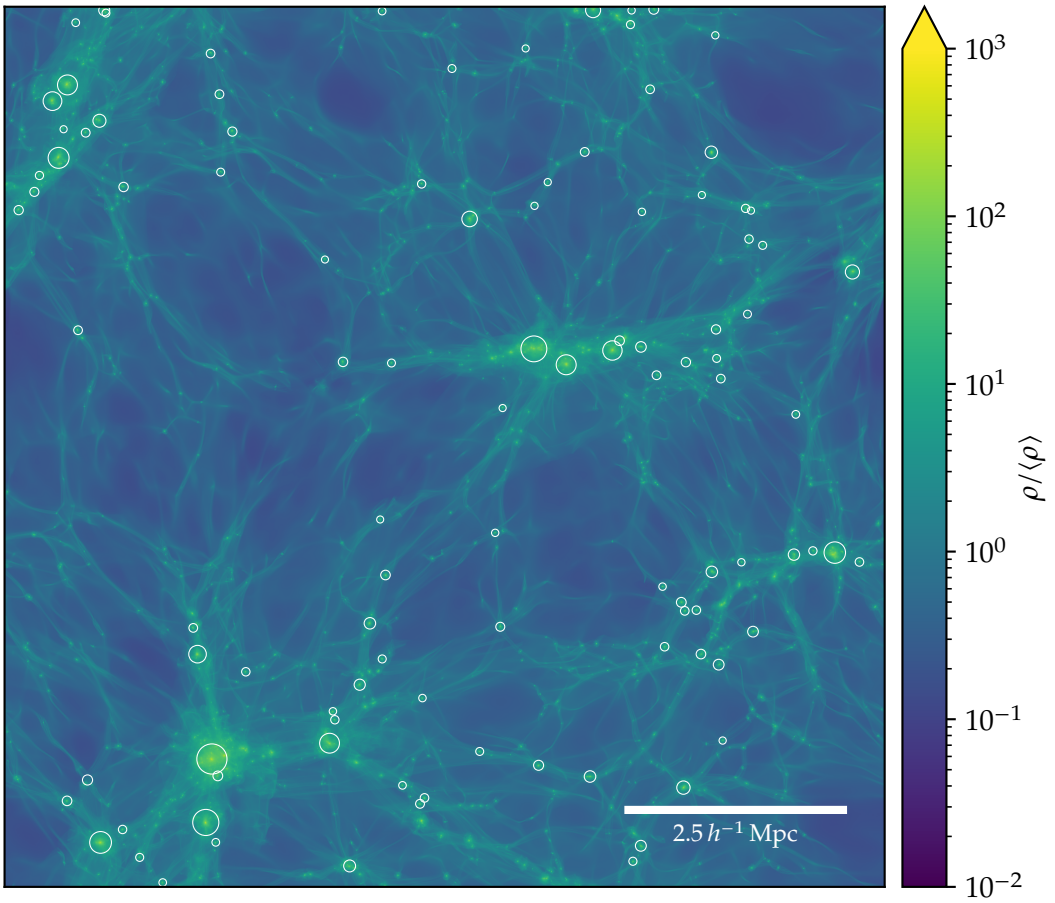

Figure 4.11: Projected dark matter density at $z = 3$ of a $3\,h^{-1}$ Mpc (comoving) slab in a high-resolution cosmological box simulation of FDM with box size $L = 10\,h^{-1}$ Mpc, FDM mass $mc^2 = 7 \times 10^{-23}$ eV, grid size $N^3 = 8640^3$, and CDM ICs. The largest halos identified with the halo finder are marked with circles whose radii indicate the halos' virial radii R_{200}.

My new simulations for the first time allow corresponding measurements, although the results are not representative of a fully consistent FDM cosmology due to the use of CDM ICs. However, my simulations allow me to demonstrate how strongly the FDM dynamics alone impact the HMF, and accordingly, to what extent their omission might affect the result. To this end, I study the HMF using the grid-based halo finder introduced in section 4.3.5, whose operation on my largest simulation is demonstrated in fig. 4.11. In fig. 4.12, I show the obtained mass functions for the FDM and CDM cases for a set of example simulations.[9] As expected from the power

[9]The halo masses are determined using the spherical overdensity definition, i. e. identifying the

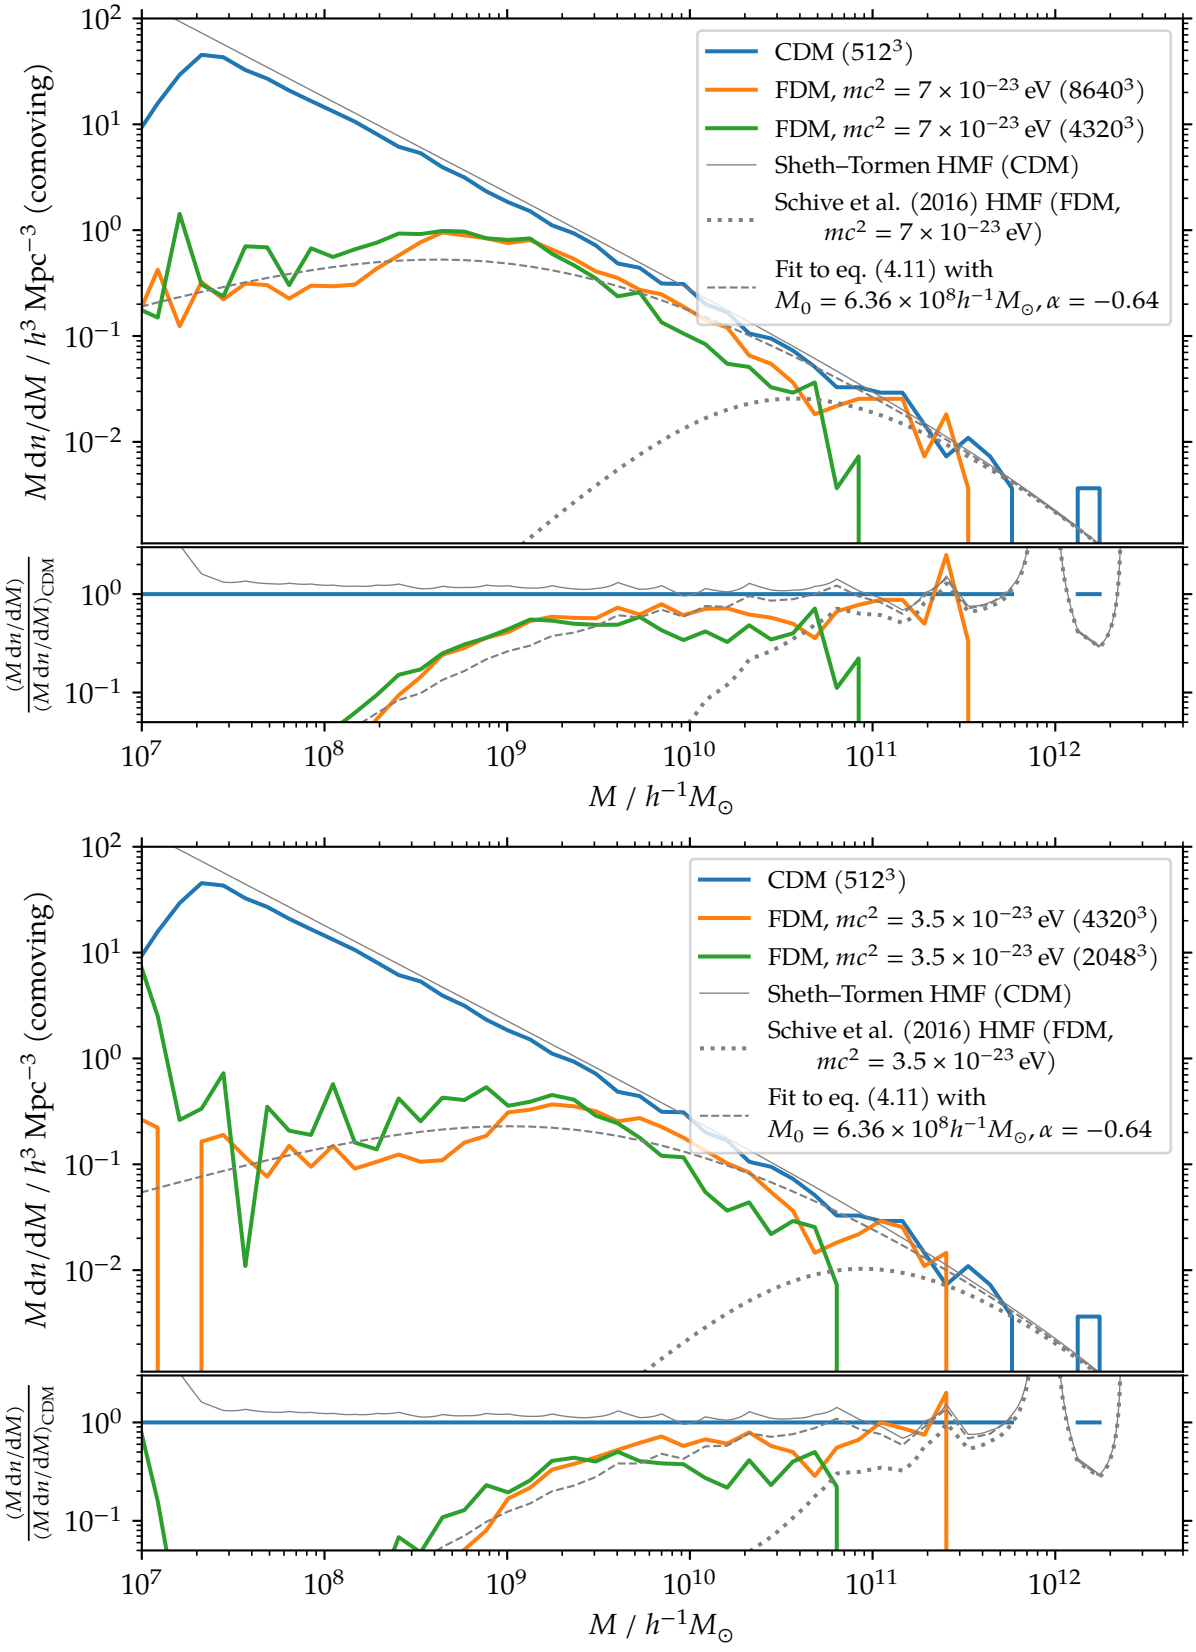

Figure 4.12: HMF of cosmological FDM and CDM simulations at $z = 3$ with a box size of $L = 10\,h^{-1}\,\mathrm{Mpc}$, CDM ICs, and FDM masses of $mc^2 = 7 \times 10^{-23}\,\mathrm{eV}$ and $3.5 \times 10^{-23}\,\mathrm{eV}$ (top and bottom panels, respectively). The HMF derived for CDM by Sheth and Tormen (1999) is shown for comparison. The dotted lines show the fitting function determined by Schive et al. (2016) for the given FDM mass, while the dashed lines show a fit of the data to a similar function (eq. (4.11)) with two free parameters. The lower sub-panels show the ratios of the mass functions to the CDM result.

spectra, the HMF of FDM exhibits a lack of low-mass halos, but agrees for massive halos. To obtain essentially perfect agreement with CDM for the most massive halos, however, requires sufficiently high resolution in the FDM calculation for the target redshift. As seen earlier, the numerical resolution requirements become ever more stringent towards lower redshift, and once they start to be compromised, the growth on large (even still linear) scales becomes damped, which then manifests itself in halo masses which are biased low. This is clearly seen in fig. 4.12 in the comparison of the low- and high-resolution FDM results for the HMF.

More interesting, however, is the abundance reduction of low-mass halos in FDM. Indeed, the delayed onset of structure formation becomes apparent even more clearly in the HMF than in the power spectrum, with almost no objects present before $z = 7$. At late times, there is a dearth of low-mass halos in the results, but the deficit appears not as strong as the cut-off predicted by Kulkarni and Ostriker (2022). Rather, it appears that power transfer to smaller scales is sufficiently strong in FDM that the break in the HMF is weaker than predicted by simplified models based on "quantum pressure"-filtered versions of the Press–Schechter formalism. This reinforces the notion that simply applying a Jeans filtering at the linear level and assuming that all smaller scales never grow and thus can be ignored does not necessarily yield sufficient quantitative accuracy to predict the non-linear outcome of FDM. On the other hand, the simulations shown in fig. 4.12 do not yet include the effect of ICs appropriate for the FDM model, which feature a strong cut-off even in the initial power spectrum (see section 4.7) and are thus expected to result in additional suppression of low-mass halos.

Also shown in fig. 4.12 are two fits, given by dotted and dashed lines, in a fashion similar to Schive et al. (2016). While the dotted line uses exactly the same function as given in Schive et al. (2016) for the corresponding particle mass m, the dashed line is a modified fit of the form

$$\left.\frac{\mathrm{d}n}{\mathrm{d}M}\right|_{\mathrm{FDM}}(M,z) = \left(1 + \left(\frac{M}{M_0}\left(\frac{mc^2}{10^{-22}\,\mathrm{eV}}\right)^{\frac{4}{3}}\right)^{\alpha}\right)^{-2.2} \left.\frac{\mathrm{d}n}{\mathrm{d}M}\right|_{\mathrm{CDM}}(M,z), \tag{4.11}$$

where the Sheth and Tormen (1999) mass function has been used for the CDM mass function $(\mathrm{d}n/\mathrm{d}M)|_{\mathrm{CDM}}(M,z)$. In other words, the mass parameter M_0 and

densest cell (FDM) or the particle with the minimum gravitational potential (CDM) as the center of a halo, the halo's virial mass M_{200} is defined as the enclosed mass of a sphere around the halo center with a radius such that the enclosed region has a mean density of $\rho_{200} = 200\langle\rho\rangle = 200\Omega_{\mathrm{m}}\rho_{\mathrm{crit}}$.

the inner exponent α are allowed to vary compared to the values found in Schive et al. (2016) ($M_0 = 1.6 \times 10^{10}\, M_\odot$ and $\alpha = -1.1$ in Schive et al. (2016)). The result shows that while an appropriate fit to the mass function can be obtained using a very similar functional form, the shape changes slightly, and the peak mass M_0 shifts to much lower halo masses. This indicates that either the method used by Schive et al. (2016) underestimates the number of low-mass halos, or that the inclusion of FDM-appropriate initial conditions (as discussed in section 4.7) causes a strong impact on the formation of low-mass halos. This question can only be resolved by performing another set of full FDM simulations with self-consistent initial conditions.

4.6 Halo profiles

The internal density structure of cosmological dark matter halos is a further key outcome of non-linear structure formation whose importance can hardly be overstated. This is because of its crucial influence on the size, kinematics and morphology of galaxies forming in the dark matter halos, as well as on the gravitational lensing strength of these objects. These properties are of course decisive for the viability of a cosmological model in the first place. For cold dark matter models, the spherically averaged density profiles have $\rho(r) \propto r^{-1}$ density cusps at their centers (Navarro, Frenk, and White 1996), giving rise to a particular rotation curve shape of galaxies.

One of the often cited motivations for considering FDM is the so-called cusp–core problem in ΛCDM, combined with the expectation that FDM naturally produces cored density profiles that may be potentially easier to reconcile with rotation curve data of certain galaxies. By considering a spherically-symmetric ground state of FDM, one arrives at solutions called solitons, whose density profiles are flat towards their centers. The production of such solitonic cores has been predicted on theoretical grounds and has been verified with simulations of small cosmological volumes (Schive, Chiueh, and Broadhurst 2014). Here, I am interested in testing whether I also see them in my comparatively large-volume simulations, which have however quite limited ability to resolve the internal structure of individual halos.

In fig. 4.13 I therefore consider spherically averaged density profiles, stacked for several different mass ranges,[10] in my highest-resolution simulation at $z = 3$. The stacking greatly reduces halo-to-halo scatter and allows me to clearly identify the mean density profile. For comparison, I also show in the right panel equivalent

[10] The "stacking" is performed by determining the mean density across all halos in a given mass bin for each radial bin.

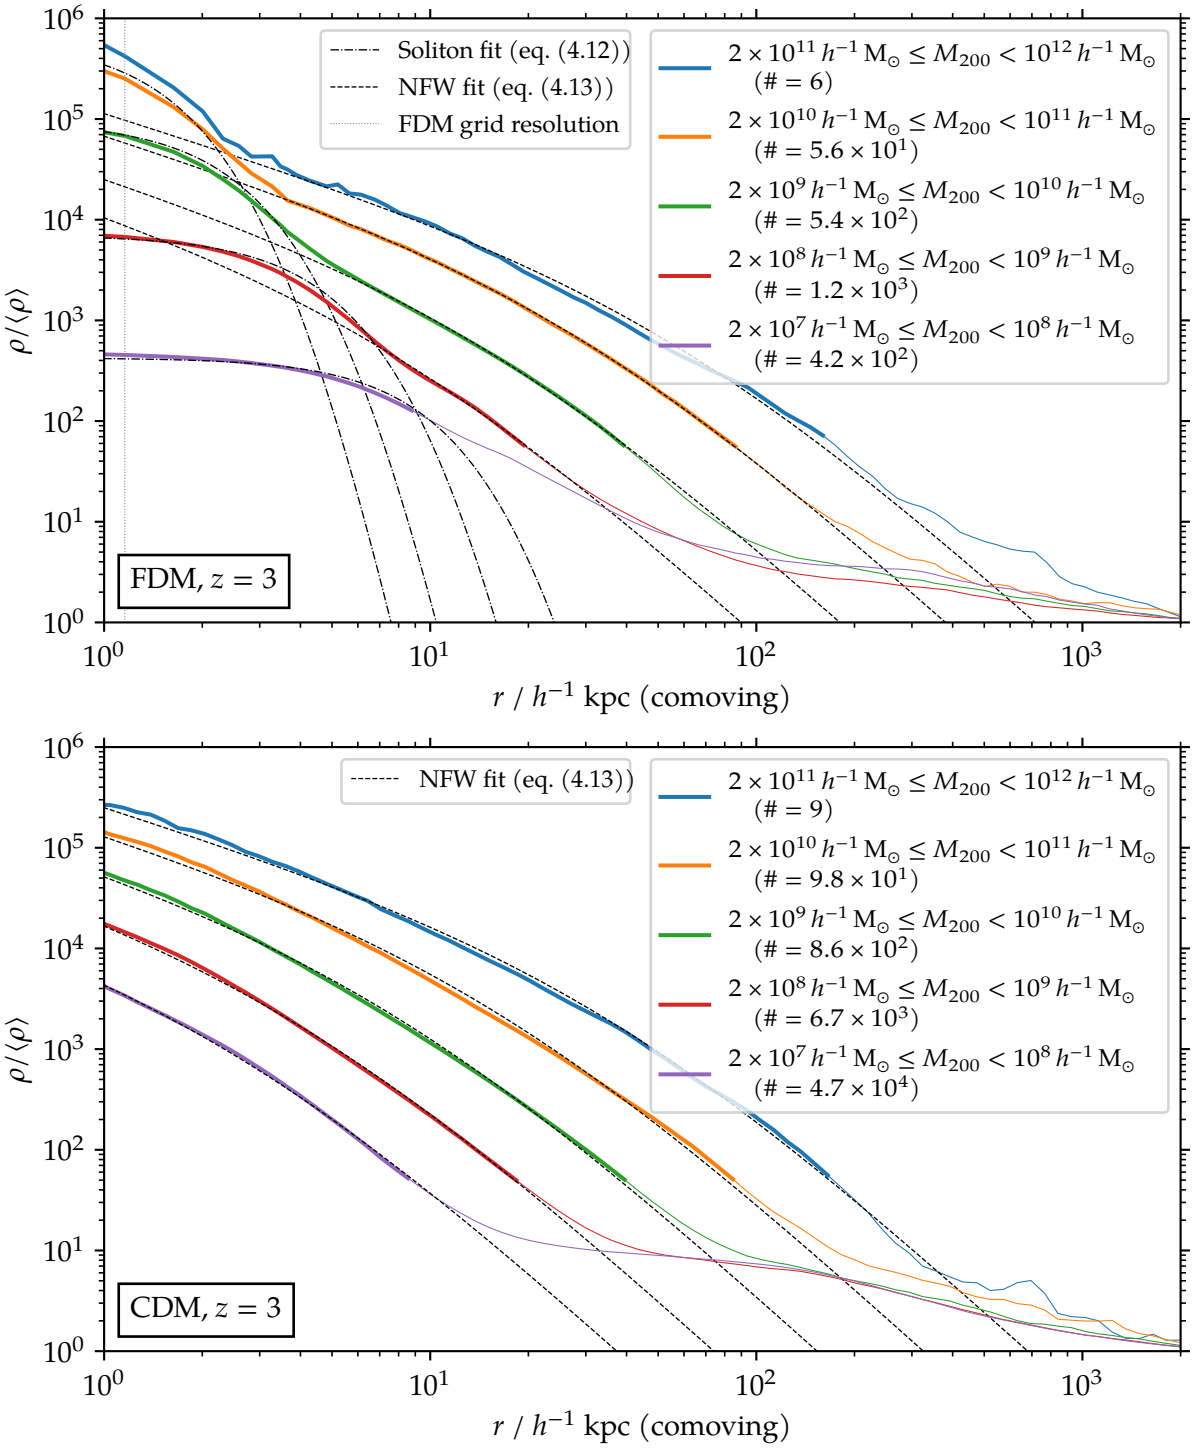

Figure 4.13: "Stacked" halo profiles at $z = 3$ for several bins of virial mass M_{200} in $L = 10\,h^{-1}$ Mpc cosmological box simulations of FDM ($mc^2 = 7 \times 10^{-23}$ eV, $N^3 = 8640^3$) and CDM (512^3 particles), using CDM ICs. The number of halos contained in each mass bin is stated in parentheses. The top panel shows the results for FDM, while the bottom panel shows those for CDM. Thin dashed lines show NFW fits (eq. (4.13)) to the region within the virial radius R_{200}, which is indicated with thick lines. For FDM, the inner region of the halo has been excluded from the NFW fit and instead been fit to the soliton density profile (eq. (4.12)), as indicated by dot-dashed lines. The FDM grid resolution is shown as a vertical dotted line.

density profiles for halos in a corresponding CDM N-body simulation. As before, the centers of halos have been identified as the densest cell in the FDM case, or the particle with the minimum gravitational potential in CDM, while the total halo mass has been determined with the spherical overdensity definition in both cases. The measurements of the density profiles include all mass around the halo centers, i. e. beyond R_{200}, the mean profiles eventually slow their decline and asymptote to the cosmic mean density.

In the FDM case, the formation of solitonic cores in low-mass halos can clearly be seen, as illustrated by the dot-dashed lines, which are fits to the analytic approximation of the soliton density profile by Schive et al. (2014):[11]

$$\rho_{\mathrm{c,soliton}}(r_\mathrm{c}) = \frac{1.9\times10^9\, a^{-1}\left(\frac{mc^2}{10^{-23}\,\mathrm{eV}}\right)^{-2}\left(\frac{r_\mathrm{s}}{\mathrm{kpc}}\right)^{-4}}{\left(1+0.091\left(\frac{r_\mathrm{c}}{r_\mathrm{s}}\right)^2\right)^8} M_\odot\,\mathrm{kpc}^{-3}. \tag{4.12}$$

Note that this profile depends on the core radius r_s only, without further free parameters. My simulations thus accurately reproduce earlier work on the shape of the innermost mean profile, where the "quantum pressure" dominates.[12] In the outer parts of the halos, the run of the mean density follows the NFW form (Navarro, Frenk, and White 1996)

$$\rho_{\mathrm{c,NFW}}(r_\mathrm{c}) = \frac{\rho_0}{\frac{r_\mathrm{c}}{R_\mathrm{s}}\left(1+\frac{r_\mathrm{c}}{R_\mathrm{s}}\right)^2} \tag{4.13}$$

rather well. It is thus tempting to use a combination of a solitonic core and a NFW profile as comprehensive description of the non-linear density structure of FDM halos, and to use this to forecast the rotation curve of galaxies and the core size of galaxies as a function of their virial mass (Burkert 2020). I caution, however, that this neglects the potentially significant impact of order-unity density fluctuations

[11] As before, ρ_c refers to the comoving density; r_c is the comoving distance from the soliton center.

[12] Since I use the mean of many halo density profiles in mass bins of a given width, the physical interpretation of fitting the soliton profile eq. (4.12), which describes a *single* soliton (with a given mass/radius), to such mean profiles is not necessarily immediately clear. The halos in a given mass bin have a range of (soliton) masses, which are distributed according to the HMF. However, since the soliton profile is simply flat in the center and for most of its domain of applicability, such a fit will simply produce a central density corresponding to the HMF-weighted mean soliton mass of the binned halos. Although it is not necessarily to be expected that the drop in density beyond the flat central plateau of the stacked profile will be accurately described by the soliton profile due to the varying soliton radii contained in each mass bin, the general agreement with the soliton profile demonstrates the presence of the cores described in previous work in my simulations.

of the FDM wave function outside the solitonic core, which are completely lost in the spherical averaging and the halo stacking. Likewise, it is unclear how severely baryons can impact the innermost density profile. These baryonic effects are quite uncertain in CDM, but even less is known about them in FDM, where they could potentially be stronger due to the lower dark matter densities in the centers of halos. It will thus likely require cosmological hydrodynamic simulations of galaxy formation in FDM to arrive at firm predictions about rotation curve shapes. My results here confirm that such calculations are in principle feasible down to intermediate redshifts in sufficiently large volumes to allow population studies of galaxies, especially if the resolution can be made spatially adaptive (also see e. g. Mocz et al. 2020).

4.7 Fuzzy dark matter initial conditions

Finally, I would like to briefly consider the additional source of differences in FDM relative to CDM introduced due to modifications of the initial conditions. Thus far, I had focused for this chapter on studying differences between CDM and FDM resulting just from the different dynamics induced by the Schrödinger–Poisson system of equations as opposed to the Vlasov–Poisson system governing collisionless dynamics of classical particles. To this end, I had used *identical* initial perturbations, retaining the CDM power spectrum for definiteness. However, a physically consistent model of a FDM cosmology will witness modifications of its initial conditions relative to CDM, particularly in terms of a significant reduction in small-scale power below a scale set by the "quantum pressure", yielding something resembling the truncated power spectra of warm dark matter models. In linear theory, this is determined by the Jeans scale k_J (eq. (2.14)), above which perturbations evolve in an oscillatory rather than growing fashion. It is however not clear how scales $k > k_\mathrm{J}$ react to non-linear power transfer from larger scales once scales around k_J become mildly or fully non-linear. This can only be accurately treated with explicit (e. g. spectral) FDM simulations like the ones which I have carried out here.

Figure 4.14 shows a comparison of two FDM simulations carried out either with CDM-like initial conditions or instead with ICs predicted by the Boltzmann code axionCAMB (Hložek et al. 2017) for an axion cosmology with an adopted particle mass of $2.5 \times 10^{-22}\,\mathrm{eV}\,c^{-2}$ in this comparison. For definiteness, I chose a $2\,h^{-1}$ Mpc box size with 2400^3 grid points to represent the wave function; the remaining cosmological parameters are the same as before (section 4.3.4). I consider the time evolution down to redshift $z = 3$, and show in each redshift panel of the plot also the linear theory

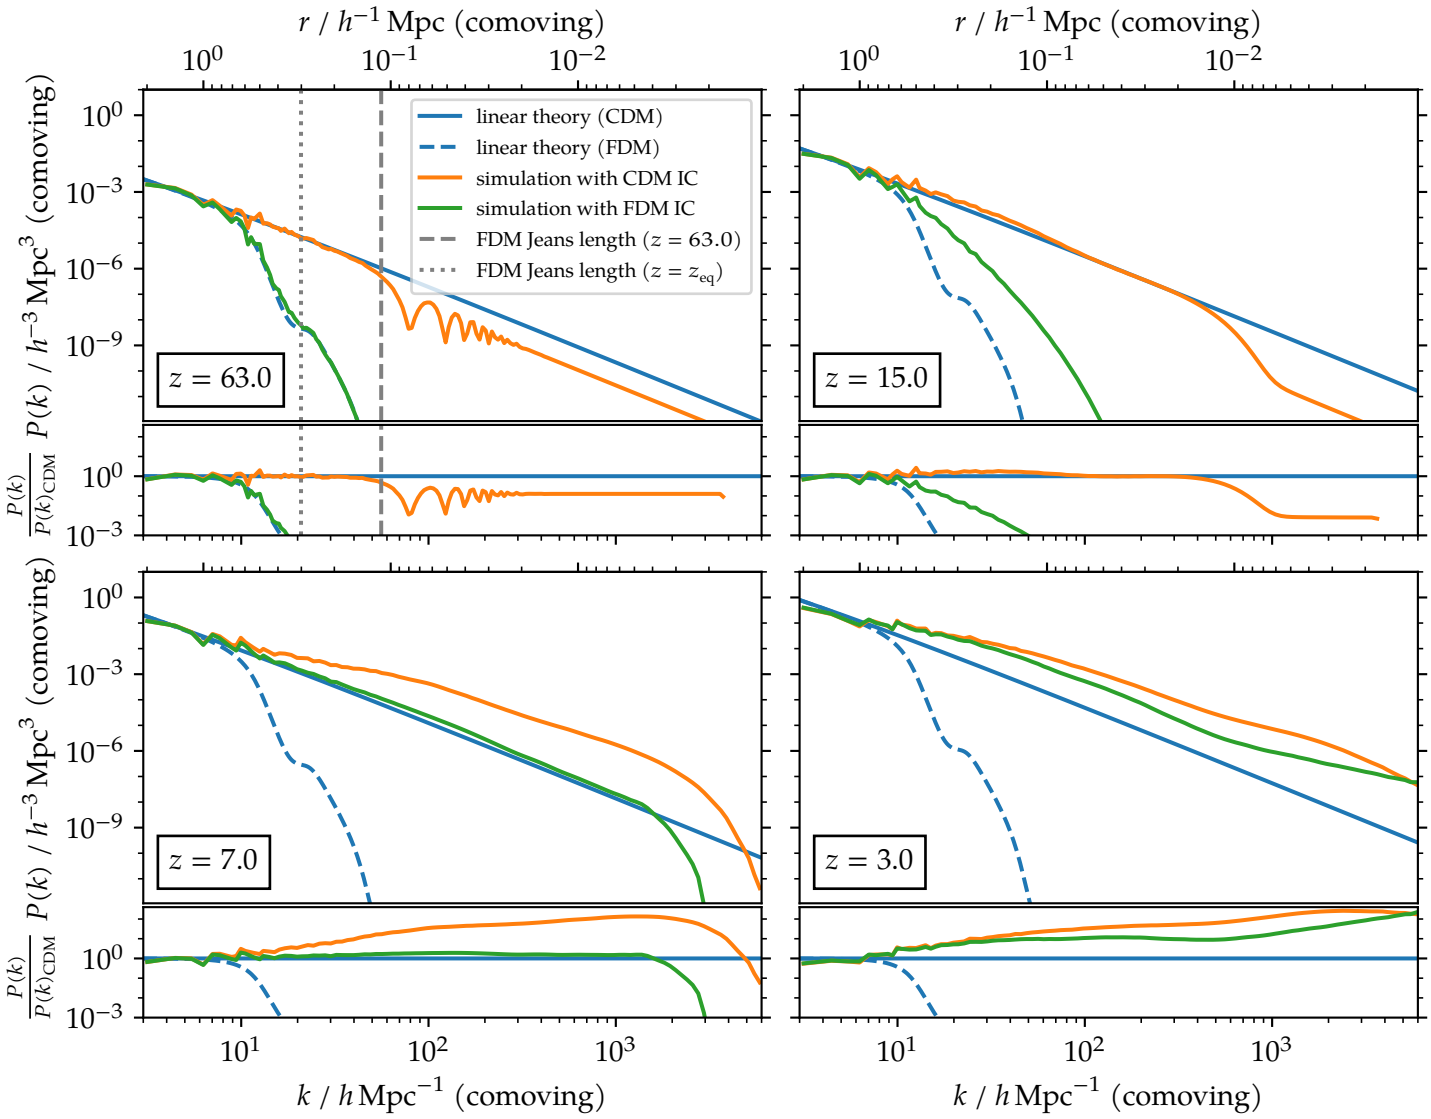

Figure 4.14: Dark matter power spectra at different redshifts for cosmological FDM simulations with both FDM and CDM ICs (FDM mass $mc^2 = 2.5 \times 10^{-22}\,\mathrm{eV}$, box size $L = 2\,h^{-1}\,\mathrm{Mpc}$, grid size $N^3 = 2400^3$). The CDM and FDM initial power spectrum evolved using linear perturbation theory (scale-independent growth factor) are shown for comparison. The lower panels show the ratio of the power spectra to the result from CDM linear theory. For $z = 63$, the dashed line additionally indicates the current FDM Jeans scale (eq. (2.14)), while the dotted line indicates the Jeans scale at matter–radiation equality, $z = z_{\mathrm{eq}}$.

power spectrum of the corresponding CDM cosmology, for comparison. At $z = 63$, the FDM model with CDM-like ICs follows linear growth accurately only to a scale of order the Jeans scale set by the "quantum pressure". On much smaller scales, the growth has completely stalled, while in the transition region oscillatory behavior is seen. This is the expected behavior based on linear theory, where scales with $k \gg k_\mathrm{J}$ should not grow, while those with $k \ll k_\mathrm{J}$ follow the same linear growth as CDM. In comparison, the FDM model with self-consistent ICs shows negligible small-scale power and only agrees with the CDM linear power on the largest scales that can be followed with this box size.

At redshift $z = 15$, the FDM simulation with self-consistent ICs still shows a stalled growth at the smallest scales, but scales around k_J have already entered mild non-linear evolution, as reflected in small excess power relative to the linear theory CDM power spectrum. Apparently, this already couples significantly to smaller scales, so that scales beyond k_J begin to catch up with the growth, such that in this instance, the power spectrum happens to agree coincidentally with the linear CDM power at around $k \approx 200\,h\,\mathrm{Mpc}^{-1}$, while scales $k \gtrsim 1000\,h\,\mathrm{Mpc}^{-1}$ are not yet affected and still prevented from growth. Similarly, the self-consistent FDM model begins to witness power transfer to smaller scales, as evidenced by the change in shape of its small-scale power spectrum.

By redshift $z = 7$, these trends have greatly accelerated. Now the non-linear evolution of the FDM model started from CDM-like initial conditions has produced a large power excess relative to linear theory for nearly all k resolved in the simulation, except for the largest scales (which are still linear), and the smallest scales (which are still suppressed by "quantum pressure"). Interestingly, however, the model with self-consistent initial conditions is clearly in the process of catching up to this evolution, with its small-scale power being progressively filled in by the non-linear evolution on larger scales. Finally, at redshift $z = 3$, both simulations have become even closer in their total matter power spectrum. These trends are expected to continue towards lower redshifts (which are unfortunately inaccessible with this resolution), so that the discrepancy introduced by starting from CDM-like initial conditions instead of self-consistent FDM ones becomes largely forgotten on small scales, simply because it is overwhelmed by non-linear evolution. However, there may still be an important impact on other quantities, such as the HMF (which, however, requires a larger box size, as in section 4.5, to measure). How large this is remains to be seen (cf. chapter 5). In principle, excess oscillatory perturbations on small scales that cannot grow should not give rise to additional collapsed halos, but due to the non-linear

coupling to larger scales this may only be approximately true.

4.8 Summary and Conclusions

In this chapter, I have carried out simulations of fuzzy dark matter models with the pseudo-spectral method. This can be considered the most accurate approach to solve the Schrödinger–Poisson system numerically, without neglecting certain aspects of the temporal evolution of the wave function. In particular, this can account for oscillatory, order-unity fluctuations of the local density due to the quantum-mechanical effects in the axion-like dynamics.

Unfortunately, the numerical resolution requirements to faithfully follow the FDM dynamics are much harder to fulfill than for the familiar N-body techniques applicable in the CDM case. Even large-scale modes require a very fine grid, otherwise the velocities appearing in the dynamics cannot properly be represented because then the spatial variation of the quantum-mechanical phase factor is not resolved. Yet, making the grid fine enough to resolve the de Broglie wavelength λ_{dB} of the largest velocities drives down the time step, which depends quadratically and not linearly on the spatial resolution. In practice, this means that one cannot trivially simulate large cosmological volumes at low resolution in a way similar to the standard practice in CDM. Rather, one is relegated to very small volumes and to the high-redshift regime, whereas a push to large box sizes and low redshift quickly becomes extremely expensive, as was highlighted in this chapter.

In this chapter, I could nevertheless compute the largest-volume simulations with full FDM dynamics thus far by using rather large grid sizes of up to 8640^3 cells. This allowed me to gain new insights into the evolution of the non-linear power spectrum in these cosmologies, and to make the first direct measurements of the halo mass function in such models, especially with respect to how the FDM dynamics affect these observables without starting from an already-suppressed power spectrum in the ICs. My main findings can be summarized as follows:

- Once sufficient resolution is available, the FDM power spectrum in my spectral simulations follows the CDM evolution very closely on large scales, even once highly non-linear evolution has set in on smaller scales.
- Lacking resolution in pseudo-spectral simulations of FDM manifests itself in a spurious “freezing” of further structure formation once the resolution requirements become violated. The evolution of the power spectrum then

stops, and it remains roughly constant with decreasing redshift. This can be detected by comparing the power on the largest scales to linear perturbation theory or a CDM simulation, and observing that the FDM model falls behind.

- The maximum simulation volume which can be achieved in the pseudo-spectral method without violating resolution requirements increases for lower particle masses m. The computational time required for the corresponding maximum box size varies only slowly for different values of m. Stopping a simulation before $z_{\mathrm{end}} = 0$ relaxes the resolution requirements, thus allowing for larger box sizes, and reduces computational time, but not very strongly (the cost approximately scales as $\propto z_{\mathrm{end}}^{-1/2}$ at fixed box size, which parallels the scaling of the maximum particle velocity in linear perturbation theory).

- On small scales, FDM structure formation is suppressed compared to CDM due to the effects of the "quantum pressure", which also induces a small-scale cut-off in the initial power spectrum of self-consistent cosmological FDM models. The onset of non-linear structure formation is thus generally delayed in FDM compared to CDM. However, on scales comparable to λ_{dB}, the difference between FDM and CDM is reduced, and interference patterns on these scales in FDM can even lead to a small temporary excess of power compared to CDM. Mildly trans-linear and fully non-linear evolution in FDM in any case leads to significant power transfer to scales smaller than the initial FDM Jeans scale, even to the extent that oscillatory perturbations on these scales, if present, can become completely overwhelmed and buried by the power transfer from larger scales.

- I could, for the first time, measure the halo mass function directly from spectral FDM simulations. For massive halos, I find the same abundance of halos as in CDM, consistent with expectations based on the agreement of their large-scale linear power spectra. The formation of the first halos is however delayed for FDM with respect to CDM, and there are considerably fewer low-mass halos forming compared to CDM. The low-mass cut-off of the HMF I find is less pronounced than found by methods which take into account the suppressed FDM ICs, even for those which neglect the non-linear FDM dynamics. This demonstrates that the FDM ICs have a larger impact on the HMF than the dynamics at later times, although the suppression caused by the latter is still sizable. However, since the cut-off appears at much lower halo masses compared to the FDM ICs, the impact of this effect is much more difficult to observe.

- Despite their large grid sizes, my simulations have only limited resolving power for the internal structure of dark matter halos. However, I could still clearly detect characteristic solitonic cores in spherically averaged density profiles of halos, with a profile shape that matches results of previous work well. These cores are particularly large for low-mass halos, where they also substantially reduce the central dark matter densities relative to CDM.

My results thus highlight that it is, in principle, possible to compute accurate cosmological simulation results for FDM in representative volumes, albeit at considerable computational cost. I will apply my implementation to FDM initial conditions (similar to section 4.7) in larger volumes in chapter 5, providing physical observables in a fully consistent FDM scenario. However, to conclusively understand the combined dynamics of FDM and baryons in the centers of galaxies, and thus to arrive at reliable predictions about whether these scenarios can resolve, e. g., the cusp–core tension, will ultimately require to add baryonic physics to these simulations, and to make them spatially more adaptive inside individual halos such that the resolution can be increased there in a targeted fashion. First attempts in this direction have been carried out (e. g. Mocz et al. 2019; Veltmaat, Schwabe, and Niemeyer 2020), but the associated numerical challenges are still formidable.

4.9 Fuzzy dark matter and the core–halo mass relation

This section includes portions of work previously published in Chan, Ferreira, May, et al. (2022). The large-scale cosmological simulations and resulting data were performed and provided by the author of this dissertation, while additional small-scale cosmological and soliton merger simulations were carried out by Hei Yin Jowett Chan, and the statistical analysis to estimate the uncertainty in the core–halo mass relation was performed by Elisa Ferreira.

Abstract

In the fuzzy dark matter (FDM) model, gravitationally collapsed objects always consist of a solitonic core located within a virialized halo. Although various numerical simulations have confirmed that the collapsed structure can be described by a cored NFW-like density profile, there is still disagreement about the relation between the core mass and the halo mass. To fully understand this relation, a large sample of

cored halos is assembled, based on both idealized soliton mergers and cosmological simulations with various box sizes. A sizable dispersion in the core–halo mass relation is found that increases with halo mass, indicating that the FDM model allows cores and halos to coexist in diverse configurations. A new empirical equation is provided for a core–halo mass relation with uncertainties that can encompass all previously-found relations in the dispersion, and emphasize that any observational constraints on the particle mass m using a tight one-to-one core–halo mass relation should suffer from an additional uncertainty on the order of 50 % for halo masses $\gtrsim 10^9\,(8\times10^{-23}\,\mathrm{eV}/(mc^2))^{3/2}\,\mathrm{M}_\odot$. Tidal stripping is suggested as one of the effects contributing to the scatter in the relation.

4.9.1 Introduction

The CDM model is one of the essential components of the standard cosmological paradigm. In this model, dark matter is described as a cold, pressureless, non-interacting fluid that dominates the matter content of the universe. The CDM model is extremely successful in explaining the observed large-scale structure of our universe (Planck Collaboration: Aghanim, Akrami, Ashdown, et al. 2020; BOSS Collaboration: Alam et al. 2017; Pillepich et al. 2018). However, on small scales, the behavior of dark matter is still weakly constrained and its properties are less understood. A prominent manifestation of this is a series of possible incompatibilities found between predictions from CDM-only simulations and observations (Bullock and Boylan-Kolchin 2017).

The FDM model is proposed to be a promising alternative to CDM (for reviews see e. g. Hui et al. 2017; Niemeyer 2020; Ferreira 2021; Hui 2021). In this model, dark matter is composed of ultra-light particles. With a particle mass as light as $10^{-22}\,\mathrm{eV}\,c^{-2}$, this candidate has a de Broglie wavelength of $\sim 1\,\mathrm{kpc}$, behaving as a wave on astrophysical scales, while on large scales it behaves like CDM, as required by observations. This wave behavior on small scales leads to a series of phenomenological consequences, like the suppression of structure formation on those scales, and the formation of a core in the interior of each galaxy halo, where the field is in its ground state (soliton). With these features, the FDM model not only presents many predictions that can be tested using observations, but depending on its mass, it might reconcile some of the small-scale incompatibilities, like the cusp–core problem.

The dynamics of structure formation in the FDM model are governed by the non-relativistic Schrödinger–Poisson system of equations. Although the computational

cost of solving the coupled system in a cosmological box is known to be much more expensive than for CDM simulations (May and Springel 2021), Schive, Chiueh, and Broadhurst (2014) were able to perform cosmological FDM simulation on an adaptive refined mesh to gain detailed insights into the non-linear structure formation. Their self-gravitating virialized FDM halos are well-resolved to confirm the existence of a solitonic core at the center of each halo, for which the density structure is approximated by the so-called soliton profile with an outer NFW-like profile. In addition, simulations have confirmed that FDM indeed mimics the non-linear power spectrum of CDM on large scales, but suppresses structure on small scales depending on the particle mass (Widrow and Kaiser 1993; Schive, Chiueh, and Broadhurst 2014; Mocz et al. 2018).

Regardless of the different numerical approaches and initial setup, several independent simulations have been performed to confirm the core–halo structure of a FDM halo, but there is still disagreement on the relation between the core mass and the halo mass, expressed as $M_\mathrm{c} \propto M_\mathrm{h}^\alpha$ (Schive et al. 2014; Schwabe, Niemeyer, and Engels 2016; Mocz et al. 2017; Nori and Baldi 2021). The relation depends on the mechanism of interaction between the core and the NFW region, which is not well understood yet. It might also depend on the formation and merger history of the halos, as shown in Du et al. (2017). Recent literature pointed out that the soliton is in a perturbed ground state interacting with the NFW region, i. e. the excited states, by means of wave interference (Li, Hui, and Yavetz 2021). The resulting oscillation of the soliton further complicates the analytical understanding of the relation.

The disagreement on the core–halo mass relation is of particular observational importance because many previous constraints on the particle mass of FDM are made by dynamic modeling of dark matter-dominated galaxies, which relies on the soliton profile and core–halo mass relation predicted by simulations. For instance, analyses of dwarf spheroidal galaxies that have often found a particle mass of $mc^2 \sim 10^{-22}$ eV or smaller (Chen, Schive, and Chiueh 2017; González-Morales et al. 2017; Safarzadeh and Spergel 2020) are in tension with measurements like the Lyman-α forest measurement $mc^2 \geq 10^{-20}$ eV (Rogers and Peiris 2021), which constrains the FDM mass by probing a different prediction, the suppression of structures. For UFD galaxies that have even smaller stellar-to-total mass ratios, some studies predicted similar particle masses as found for dwarf spheroidals (Calabrese and Spergel 2016), while others (Safarzadeh and Spergel 2020) have found that the particle mass should be heavier, with the strongest bound coming from Hayashi, Ferreira, and Chan (2021) with a particle mass as heavy as $mc^2 = 1.1^{+8.3}_{-0.7} \times 10^{-19}$ eV from Segue I. Constraints

from ultra-diffuse galaxies also suggest a FDM mass of $mc^2 \sim 10^{-22}$ eV (Broadhurst et al. 2020). Except for the Lyman-α bounds, the constraints cited above depend on the assumed core–halo mass relation. Although the origin of such incompatibilities might also be due the influence of baryons in these systems, the core–halo relation is another important aspect, and any change or uncertainty in this relation will influence the bounds on the FDM mass cited above.

Here, new FDM halo simulations are performed, and combined with the largest cosmological FDM simulations with full wave dynamics to date (May and Springel 2021, presented earlier in this chapter), to obtain a large sample of collapsed objects. The core–halo mass relation is revisited, and find a scatter that can encompass all previously-found relations (i. e. Schive et al. 2014; Mocz et al. 2017; Mina, Mota, and Winther 2022; Nori and Baldi 2021).

4.9.2 Theory

A consequence of the finite Jeans length and corresponding suppression of small-scale structure formation can be seen in the suppression of small-scale power in the power spectrum of these models, and consequently the suppression of the formation of smaller halos. The effect of this suppression can also be seen inside halos, where there is a highly non-linear evolution. The interior of each halo forms a core, where there is no further structure formation and the FDM field is in its ground state. A gravitationally bound object thus consists of two components in the FDM model: The inner part – where quantum pressure dominates – is called the core, while in the outer part, gravity dominates and structure formation can happen.

In Schive et al. (2014), a fitting function for the core–halo mass relation was obtained:

$$M_\mathrm{c} = \frac{1}{4\sqrt{a}} \left[\left(\frac{\zeta(z)}{\zeta(0)} \right)^{1/2} \frac{M_\mathrm{h}}{M_{\mathrm{min},0}} \right]^{1/3} M_{\mathrm{min},0}\,, \tag{4.14}$$

where M_c and M_h are again the core and halo masses, and

$$M_{\mathrm{min},0} \sim 4.4 \times 10^7 (mc^2/(10^{-22}\,\mathrm{eV}))^{-3/2} M_\odot\,, \tag{4.15}$$

and the outer exponent $\alpha = 1/3$ represents the (logarithmic) slope of the relation $M_\mathrm{c} \propto M_\mathrm{h}^\alpha$. In order to compare with Schive et al. (2014), their definition of halo mass $M_\mathrm{h} = (4\pi r_\mathrm{h}^3/3)\zeta(z)\rho_{\mathrm{m}0}$ is followed here, where r_h is the halo's virial radius, $\rho_{\mathrm{m}0}$ is the background matter density and $\zeta \sim 180$ (350) for $z = 0$ (≥ 1).

Note that the SP equations follow a scaling symmetry

$$\{x, t, \rho, m, \psi\} = \{\alpha x, \beta t, \beta^{-2}\rho, \alpha\beta^{-2}m, \beta^{-1}\psi\}\,. \tag{4.16}$$

Therefore, this symmetry can be used to re-scale the resulting structure of a simulation to another particle mass.[13]

Previous studies were able to confirm the empirical density profile eqs. (4.12) and (4.13) using different simulations. However, they disagree about the form of the core–halo mass relation, calling the validity of eq. (4.14) obtained by Schive et al. (2014) into question. Schwabe, Niemeyer, and Engels (2016) performed idealized soliton merger simulations and were unable to reproduce eq. (4.14). Mocz et al. (2017) used a larger halo sample with simulations of a similar setup and obtained a slope of $\alpha = 5/9$, disagreeing with eq. (4.14). Mina, Mota, and Winther (2022) found the same slope of $5/9$ using cosmological simulations with a box size of $2.5\,\mathrm{Mpc}\,h^{-1}$. Finally, Nori and Baldi (2021) performed zoom-in simulations by including an external quantum pressure term in an N-body code, and obtained a relation with yet another value of the slope, $\alpha = 0.6$. Such disagreement between different studies indicates that there is still a fundamental lack of understanding of the core–halo structure in the FDM model, and also generates uncertainty in any constraints on the FDM mass which were obtained using eq. (4.14) or similar relations. Therefore, the main motivation of this section is to revisit and clarify the core–halo mass relation.

4.9.3 Results for the core–halo mass relation

Figure 4.15 shows the core–halo mass relation obtained from the soliton merger and cosmological simulations. All data are scaled to $mc^2 = 8 \times 10^{-23}\,\mathrm{eV}$ using eq. (4.16) in order to enable a direct comparison with the data and fitting relation from Schive et al. (2014). For reference, the "core–halo" mass relation of a soliton-only profile is also shown, represented by the solid black line. This curve indicates the minimum halo mass for a certain core mass, and any halos located to the right of the soliton-only core–halo relation must have the usual cored NFW structure. For halos in the soliton merger simulations with mass $\gtrsim 10^8\,M_\odot$, the relation has a steeper slope than $\alpha = 1/3$, confirming the results from Mocz et al. (2017). However, halos from the large-scale cosmological simulation predict a core–halo relation with a large enough dispersion that can cover a range of data produced by both the soliton merger simulations and

[13] Since only the goal is to re-scale only the mass, $\beta = 1$ is fixed and only α changed to perform the scaling. This α is unrelated to the slope of the core–halo mass relation.

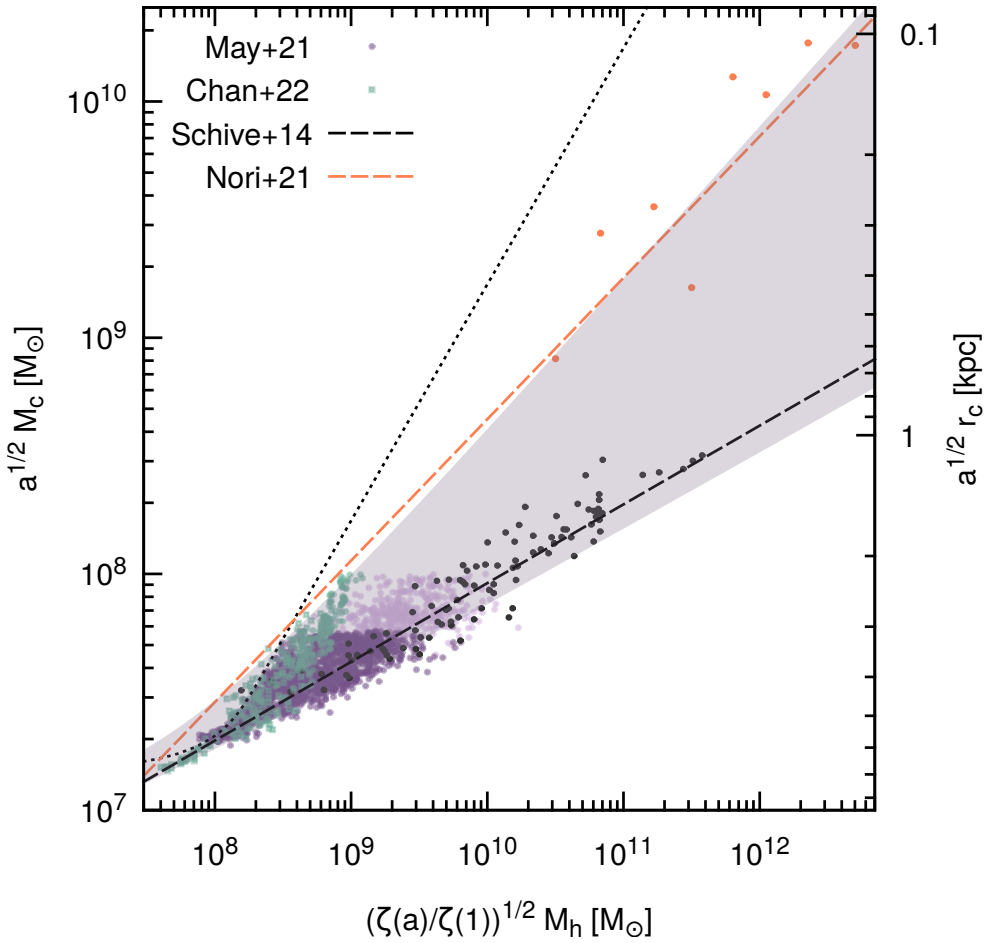

Figure 4.15: Core–halo relation scaled to $mc^2 = 8 \times 10^{-23}$ eV via eq. (4.16). Green dots are halos simulated in Chan, Ferreira, May, et al. (2022) with cores resolved by at least $3\Delta x$. Purple and and faint purple dots are halos from the large-box cosmological simulation (May and Springel 2021) with cores resolved by at least $2\Delta x$ and Δx respectively. The pink shaded region is enclosed by the empirical fits to the purple and green dots, with the maximum and minimum values of the parameters in eq. (4.14). The solid dotted line corresponds to the soliton-only relation obtained from a pure core profile. The black and orange dashed lines are fitting relations corresponding to the black and orange dots obtained from Schive et al. (2014) and Nori and Baldi (2021) respectively.

Schive et al. (2014). The range of the dispersion can span as large as one order of magnitude in halo mass for $M_c \sim 5 \times 10^7 M_\odot$. This dispersion, which fills in the space in between the soliton-only line and the relation from Schive et al. (2014), indicates the diversity of the cored NFW structure in the FDM simulations. The tight "one-to-one" core–halo relations found by different groups, with different slopes,

therefore only describe a part, but not all populations of halos in the FDM model.

An empirical equation is suggested that has the following form: $M_\mathrm{c} = \beta + (M_h/\gamma)^\alpha$. The parameter β takes the limit of the relation for small halo masses into account, although low-mass halos are rare in a FDM universe due to the suppression in the initial power spectrum. α is the slope that can be compared to previous works. Including the scaling symmetry in eq. (4.16) and the redshift dependence according to Schive et al. (2014) results in

$$\begin{aligned} a^{1/2} M_\mathrm{c} = \beta \left(\frac{mc^2}{8 \times 10^{-23}\,\mathrm{eV}} \right)^{-3/2} \\ + \left(\sqrt{\frac{\zeta(z)}{\zeta(0)}} \frac{M_\mathrm{h}}{\gamma} \right)^\alpha \left(\frac{mc^2}{8 \times 10^{-23}\,\mathrm{eV}} \right)^{3(\alpha-1)/2} M_\odot \,. \end{aligned} \tag{4.17}$$

The best-fit parameters for the halos from the large-box cosmological simulation give $\beta = 8.00^{+0.52}_{-6.00} \times 10^6\, M_\odot$, $\log_{10}(\gamma/M_\odot) = -5.73^{+2.38}_{-8.38}$ and $\alpha = 0.515^{+0.130}_{-0.189}$, which is shown as a pink shaded region in fig. 4.15.

The effect of the large dispersion is encompassed in the uncertainty of the model parameters. This uncertainty is not the statistical uncertainty of the fit, but an "overestimation" of the uncertainty in the parameters that can reflect the large dispersion of the data. Indeed, the statistical uncertainty would be the incorrect quantity to consider in this case, since there is no assumption that there is an underlying "true" set of values for the parameters with statistical fluctuations, but rather propose that different halo populations could *systematically* follow different relations depending on their histories and properties (see section 4.9.4). To obtain a more appropriate description of the core–halo diversity, kernel density estimation (KDE) was employed, estimating the probability distribution function of the core masses with respect to the central value of the corresponding binned halo mass. Each of these distributions reveals the dispersion of core masses for each halo mass.The minimum and maximum curves $M_\mathrm{c}(M_\mathrm{h})$ that fit all of these distributions are then obtained, and the minimum and maximum vales for the parameters b, γ and α extracted from these curves. The difference to the global fit is the uncertainty in the parameters.

Nori and Baldi (2021), Mocz et al. (2017), and Schive et al. (2014) determined slopes α of 0.6, 0.556 and 0.333, respectively. Given the large dispersion seen in the data, all of these slopes are compatible when taking into account the uncertainty assigned to the fitting parameters. So when considering the proposed fitting function, all of the other cases in the literature are covered as well. Is should be emphasized that

the results show that a general halo population is not well-described by any single one-to-one core–halo mass relation. Further investigation is required to determine which halo populations follow which relations (if any), and under what conditions – cf. section 4.9.4.

This large spread and uncertainty in the fitting function can affect the constraints on the FDM mass obtained from these relations. A a rough estimate of the error is provided in the following. For the same halo mass $M_\mathrm{h} = 10^9\,M_\odot$ in fig. 4.15, the least massive core mass is $M_\mathrm{c} = 3 \times 10^7\,M_\odot$ and the most massive is $M_\mathrm{c} = 10^8\,M_\odot$. Applying these values to the core density in eq. (4.12) gives a 50 % difference in particle mass m. Therefore, any observational constraints made using the relation eq. (4.14) should include an additional uncertainty on the order of 50 % in the results, unless the halo mass is smaller than $10^9 (8 \times 10^{-23}\,\mathrm{eV}/(mc^2))^{3/2}\,M_\odot$. Therefore, when obtaining the FDM mass using the core–halo relation, one needs to take into account the dispersion of these values, shown in the uncertainty in the fitting parameters, which will translate to a higher uncertainty in the FDM mass.

Due to limited spatial resolution, it was only possible to observe the dispersion to increase with halo mass until $M_\mathrm{c} \sim 6 \times 10^7\,M_\odot$. It would be important for potential future higher-resolution simulations to examine if the dispersion keeps increasing along the soliton-only relation or not. Again, the increasing dispersion is of importance to observational studies since it will also lead to an increasing uncertainty in the core–halo relation.

4.9.4 The origin of the dispersion

Different core–halo structures have been found in different simulations:

- As mentioned before, Schive et al. (2014) and Mocz et al. (2017) find different results for the slope α (1/3 vs. 5/9), even for similar simulation setups (soliton mergers).

- Mina, Mota, and Winther (2022) claim to confirm a slope of $\alpha = 5/9$, as found in the soliton merger simulations of Mocz et al. (2017), but using a *cosmological* simulation, contradicting the result of $\alpha = 1/3$ from Schive et al. (2014). However, the number of halos in their sample is very small.

- Schwabe, Niemeyer, and Engels (2016) performed soliton merger simulations similar to Schive et al. (2014) (and later Mocz et al. 2017)[14] and could not reproduce the previously-found value of the slope α, or indeed any universal relation.
- Nori and Baldi (2021) studied the dynamics of eight simulated halos and concluded with a similar comment: Schive et al. (2014) and Mocz et al. (2017) only captured a partial representation of the core–halo relation in a realistic cosmological sample.
- Yavetz, Li, and Hui (2021) used the Schwarzschild method to construct self-consistent FDM halos and found that a stable core–halo structure can exist even when the adopted core–halo mass relation deviates from Schive et al. (2014).

These examples illustrate that the diversity of the possible core–halo slopes found in different works seems to originate from the type of simulations performed, which results in halos and cores that have different properties. The diversity of core–halo structure found in these simulations is exhibited in this chapter, where the difference between the core–halo mass relation from halos formed in soliton merger simulations (green points in fig. 4.15) and in cosmological simulations (pink points in fig. 4.15) can clearly be seen.

There are a few possible explanations for this diversity of halos: merger history (Du et al. 2017; Yavetz, Li, and Hui 2021), tidal stripping effects, and the relaxation state of the halo (Nori and Baldi 2021). Formation and merger history is an explanation that seems very plausible to be a relevant factor. Larger cosmological simulations, like the one from May and Springel (2021) (presented earlier in this chapter), present halos that could have very different merger histories, and a large dispersion is expected. This is different from the soliton merger simulations, where a complicated merger history would not be expected. It is left for future work to try to identify the different merger histories and try to clarify how this relates to the different incarnations of the core–halo mass relation.

Another possible factor that can also contribute to the dispersion found is stripping. The following will be an attempt to provide an argument to support tidal stripping as one element responsible for the dispersion, based on the setups of various simulations. By comparing the box sizes and the resulting slopes α between the small-volume cosmological simulations of Chan, Ferreira, May, et al. (2022) with Mocz et al. (2017)

[14] Although Schwabe, Niemeyer, and Engels (2016) made use of "sponge" boundary conditions instead of periodic boundary conditions.

and Schive et al. (2014), which are 335 kpc, 1765 kpc and ≥2000 kpc (box sizes) after re-scaling via eq. (4.16), and ∼0.9, 0.556 and 0.333 (slopes) respectively, one finds that smaller simulation box sizes are correlated with a steeper slope in the core–halo relation. This can be explained by the stripping effect on the halo by its own gravity due to the periodic boundary conditions: the self-stripping effect becomes more effective at removing mass from the NFW region as the box size decreases. This skews the core–halo structure towards smaller halo masses, steepening the core–halo relation. A more rigorous test to prove the above argument requires simulations with increased spatial resolution and box sizes up to at least 2 Mpc, which current numerical schemes are unable to feasibly achieve.

The self-stripping effect is a numerical artifact, but there is no doubt that a stable core–halo structure can exist within such environments. In more realistic cosmological simulations, dwarf satellites also experience a similar effect from their host halos in the form of tidal stripping. Therefore, stripping effects by tidal forces can be suggested as one of the contributing factors causing the dispersion obtained from the large-box simulation in May and Springel (2021). One subtlety is that the tidal effect is an interaction between host halos and sub-halos with at least two orders of magnitude difference in mass, but the halo finder used in May and Springel (2021) does not identify sub-halos. However, it is known that sub-halos in CDM simulations can temporarily move outside of the virial radius of the host halo after the first pericentric passage (Bosch 2017). Ejected sub-halos should also exist in an FDM cosmology, and are therefore identified by the halo finder. An in-depth analysis of the tidal effect on the core–halo relation, or FDM sub-halos in general, would require building merger trees, which is still not yet studied in any FDM cosmological simulations. This investigation is left to future work.

4.9.5 Core radius–halo mass relation

As suggested by Burkert (2020), the FDM model may fail to explain the observed trend of the core radius–halo mass relation measured from dwarf galaxies. Following Mina, Mota, and Winther (2022), the core radius–halo mass relation is presented as measured from the FDM halo samples. As shown in fig. 4.16, the scatter is still observed, but the decreasing trend, which is a fundamental property of quantum pressure-induced cores, is in disagreement with the positive scaling predicted by low surface brightness (LSB) galaxies (Salucci et al. 2007; Di Paolo, Salucci, and Erkurt 2019).

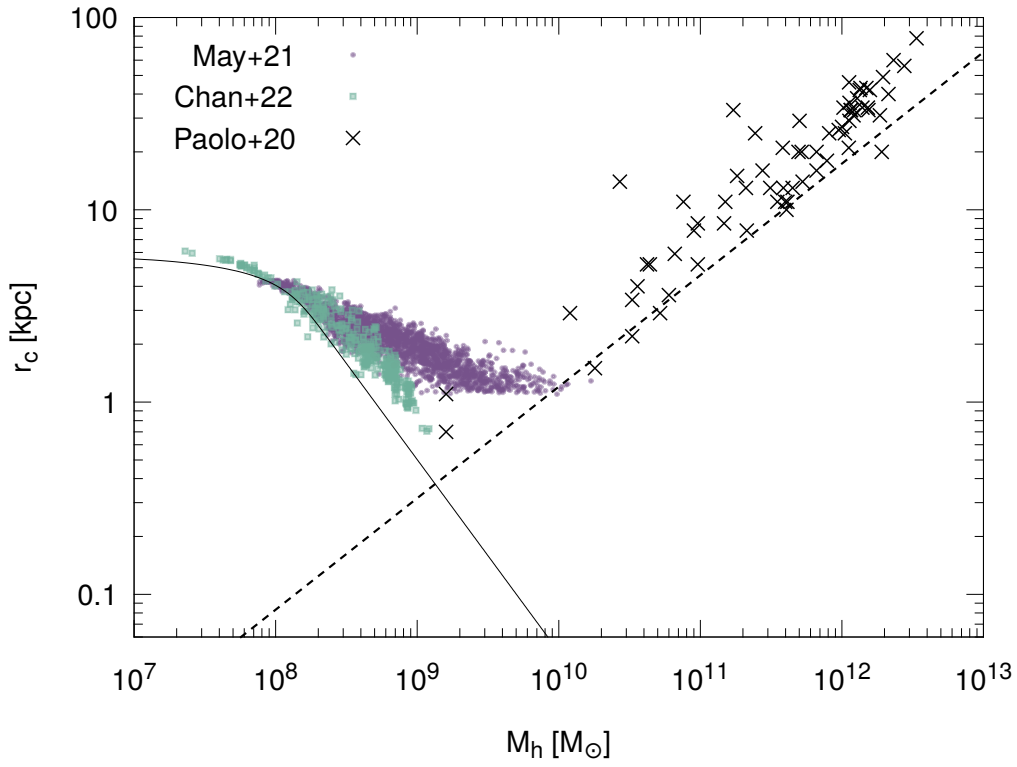

Figure 4.16: Core radius vs. halo mass. Green and purple points are properties of halos from simulations of Chan, Ferreira, May, et al. (2022) and May and Springel (2021). The black line shows the relation predicted by a soliton-only density profile. The dashed line is an empirical function predicted by low surface brightness (LSB) galaxies (Salucci et al. 2007). Black crosses are from Di Paolo, Salucci, and Erkurt (2019).

The disagreement is expected because the negative scaling, where less massive galaxies are cored, allows the FDM model to solve the core–cusp problem, but the relation from LSB galaxies has the opposite behavior, where massive galaxies have larger cores. In addition, LSB galaxies are predicted in CDM simulations to have experienced tidal heating and supernova feedback (Martin et al. 2019). Therefore, the relation between core radius and halo mass poses a challenge to the FDM model, but more importantly, it motivates future FDM simulations to include baryonic physics to verify if LSB-like galaxies can be formed or not.

The fuzzy dark matter halo mass function and filaments

5

The contents of this chapter have been published as May and Springel (2023) (with slight modifications).

Abstract

Fuzzy dark matter (FDM) is a dark matter candidate consisting of ultra-light scalar particles with masses around $10^{-22}\,\mathrm{eV}/c^2$, a regime where cold bosonic matter behaves as a collective wave rather than individual particles. Although constraints on FDM are accumulating in many different contexts, very few have been verified by self-consistent numerical simulations. I present new large numerical simulations of cosmic structure formation with FDM, solving the full Schrödinger–Poisson (SP) equations using the AxiREPO code, which implements a pseudo-spectral numerical method. Combined with my previous simulations, they allow me to draw a four-way comparison of matter clustering, contrasting results (such as power spectra) for each combination of initial conditions (FDM vs. CDM) and dynamics (SP vs. N-body). By disentangling the impact of initial conditions and non-linear dynamics in this manner, I can gauge the validity of approximate methods used in previous works, such as ordinary N-body simulations with an FDM initial power spectrum. Due to the comparatively large volume achieved in my FDM simulations, I am able to measure the FDM halo mass function from full wave simulations for the first time, and compare to previous results obtained using analytic or approximate approaches. I also investigate the density profiles of these filaments and compare to their ΛCDM counterparts.

5.1 Introduction

A number of tensions between theory and observations on small cosmological scales, such as the "cusp–core problem", the "missing satellite problem", or the "too-big-to-fail problem" (Weinberg et al. 2015; Del Popolo and Le Delliou 2017; Boylan-Kolchin, Bullock, and Kaplinghat 2011) have sometimes raised questions about the validity of the "standard cosmological model" based on ΛCDM, which has otherwise been extremely successful in describing a wide variety of cosmological observations across a broad range of physical scales (e. g. Frenk and White 2012; Bull et al. 2016). Distinguishing effects due to baryonic physics from genuine failures of ΛCDM has proven to be a significant challenge, however (e. g. Santos-Santos et al. 2020; Grand et al. 2021; Sales, Wetzel, and Fattahi 2022). Further, CDM as a model only makes predictions in the context of cosmology, with countless possible implementations of galaxy formation physics (see Vogelsberger et al. 2020, for a recent review).

Although large-scale structure studies have solidified the correspondence of observations with the behavior of CDM, dark matter models can still deviate on less well constrained, smaller (i. e. galactic) scales. In this regard, FDM offers a wide range of new phenomena which have an impact on some of the "small-scale problems" (see e. g. Marsh 2016; Hui et al. 2017; Niemeyer 2020; Hui 2021; Ferreira 2021, for reviews). Due to the small particle masses, interesting wave effects occur which are unique to this class of dark matter models. Early numerical simulations have already shown that ultra-light scalars form cores in the centers of dark matter halos (Schive, Chiueh, and Broadhurst 2014), possibly explaining observed dwarf galaxy rotation curves (but see Burkert 2020). In addition, the cut-off in the FDM transfer function, suppressing small-scale power in a similar fashion as WDM, has the potential to solve issues like the "missing satellite problem". In a sense, FDM combines properties similar to the features of WDM (suppression on small scales) and SIDM (cored density profiles) with respect to the small-scale challenges mentioned above (Boyarsky et al. 2019; Tulin and Yu 2018).

However, FDM also exhibits a number of additional unique phenomena. Objects other than cores, such as quantized vortices, are another differentiating feature with interesting prospects for detection (Hui et al. 2021), and relative fluctuations of order one in the density field at the scale of the FDM wavelength can have a strong impact on visible matter e. g. through dynamical heating. Furthermore, light (pseudo-)scalar particles are a common feature of theories in particle physics, from the original axion in QCD to a plethora of axion-like particles predicted by unified and early-universe

theories such as string theories (Marsh 2016).

After FDM garnered great interest due to motivation from astrophysics and particle physics, as summarized above, constraints on the mass m of the scalar particles at values around $mc^2 \gtrsim 10^{-21}$ eV accumulated using many different contexts and observables for some time (Ferreira 2021). It has turned out that the scaling relations of FDM halo cores are difficult to reconcile with observations when assuming a single value for the particle mass m (Burkert 2020), and although this does not rule out FDM as a dark matter candidate by itself, it does weaken its motivation as a possible solution to the "cusp–core problem". Relatedly, recent results from dynamical modeling of UFD galaxies (where FDM cores become increasingly large) have been used to claim rather strong constraints on m, with values up to $\approx 10^{-19}$ eV (Hayashi, Ferreira, and Chan 2021; Zoutendijk et al. 2021). However, most of these analyses do not take the significant scatter in the FDM core–halo mass relation, discovered in cosmological simulations, into account (Chan, Ferreira, May, et al. 2022), which would weaken these bounds. A very recent result from Dalal and Kravtsov (2022), claiming $mc^2 > 3 \times 10^{-19}$ eV due to the heating effect of stellar orbits caused by potential fluctuations in FDM halos, has more serious implications, since it does not depend on the uncertain core–halo relation or delicate measurements of the UFD central density.

However, most of these constraints (including the most recent) have in common that they have not yet been verified by self-consistent numerical simulations, instead relying on approximate, idealized, or simplified numerical and analytic approaches. While cosmological FDM simulations have been carried out using a variety of numerical methods, many attempts were quite limited in scope (Zhang, Liu, and Chu 2018; Laguë et al. 2021, table 1), so that the effects of FDM in (mildly) non-linear regimes of structure formation are still poorly understood compared to CDM.

Apart from an overall still lower level of research attention, an important reason impeding insight into FDM lies in the very large computational costs incurred when numerically solving the corresponding equations of motion – these costs are much higher than the ones associated with corresponding ΛCDM calculations. Due to the computational requirements, the cosmological volumes studied in simulations with the full FDM SP dynamics have been especially limited (Woo and Chiueh 2009; Schive, Chiueh, and Broadhurst 2014; Veltmaat, Niemeyer, and Schwabe 2018; Mocz et al. 2020). Although recent advances in hybrid numerical techniques have made it feasible to embed simulated FDM halos within much larger simulated boxes (Schwabe and Niemeyer 2022), the fundamental issues driving the cost in the regions

where the full FDM equations of motion are treated remain. In this chapter, I carry out simulations that smoothly connect the non-linear state reached in isolated FDM halos to the still linear large-scale structure, thereby bridging, in particular, the regime of mildly non-linear evolution where differences in the temporal evolution compared to CDM can be expected. I complement the previous simulations (chapter 4) by including the self-consistent transfer function expected for an FDM cosmology. As before, I carry out very large FDM simulations with a spectral method on a uniform grid, which fully retains the quantum-mechanical effects. I compare and contrast with my previous results for central measures of matter clustering, namely the power spectrum and the HMF, and highlight the unique challenges and phenomena related to filaments which arise when using an FDM initial power spectrum.

While methods which forego a treatment of the full wave dynamics have been able to conduct simulations with volumes much closer to those attainable using traditional *N*-body and SPH approaches for CDM (Schive et al. 2016; Veltmaat and Niemeyer 2016; Zhang et al. 2018; Nori and Baldi 2018; Nori et al. 2019; Nori and Baldi 2021), these do not capture inherent wave phenomena such as interference effects, which can have a significant impact on the overall evolution at least on small scales (Li, Hui, and Bryan 2019), leaving the validity of results obtained in this way unclear in the absence of similar computations solving the fundamental wave equations. While the hybrid method from Schwabe and Niemeyer (2022) improves upon the computational limitations, by its nature it also does not incorporate the full FDM evolution, and can only reproduce it in a statistical sense. In particular, while all simulations can easily incorporate the impact of the suppressed small-scale power spectrum present with FDM in the ICs, such methods either lack the wave nature of FDM entirely or only approximate it. Using my simulations, I am now equipped to fully clarify the reliability of such approximative results, disentangling the two essential physical differences distinguishing FDM from CDM in cosmological numerical simulations: the dynamics (equations of motion) and the ICs.

This chapter is structured as follows. In section 5.3 I detail the numerical methodology. In section 5.4 I compare matter clustering in FDM and CDM cosmologies at the level of the power spectrum, for different sets of initial conditions. I then turn to a discussion of the challenges involved in measuring the HMF in self-consistent cosmological FDM simulations in section 5.5. In section 5.6 I consider the density structure of filaments in FDM as compared to CDM, while I return to halo profiles in section 5.7. Finally, I present a summary of the findings and my conclusions in section 5.8.

5.2 Theoretical background

Since the theoretical background and fundamental equations relevant for FDM have already been covered in section 2.1, the associated commentary in the published manuscript corresponding to this chapter (May and Springel 2023) will not be reiterated here and the reader is referred to section 2.1.

5.3 Numerical methodology

The numerical simulations are performed using the same setup as in chapter 4. Two types of simulations are presented: (1) FDM simulations using the AxiREPO code[1] first introduced in May and Springel (2021) (cf. chapter 4, which is implemented as a module in AREPO (Springel 2010) and numerically solves the SP equations, eqs. (2.3) and (2.4), using a pseudo-spectral method, and (2) "standard" N-body CDM simulations using the unmodified AREPO code, which implements the TreePM numerical method to solve the Vlasov–Poisson equations. For each type of simulation, I additionally compare and contrast two different kinds of ICs, corresponding to an FDM or CDM universe (see section 5.3.1).

As in the previous chapter 4, the simulation volume consists of a cubic box of side length L with periodic boundary conditions, sampling the matter distribution in the universe. The box is filled with dark matter (FDM or CDM) whose average comoving density is the cosmic mean background matter density

$$\langle \rho \rangle = \rho_{\mathrm{m}} = \Omega_{\mathrm{m}} \rho_{\mathrm{crit}} = \Omega_{\mathrm{m}} \frac{3H_0^2}{8\pi G}. \tag{5.1}$$

In the FDM case (pseudo-spectral SP solver), the fields ψ ("wave function") and Φ ("potential") are discretized on a uniform Cartesian grid with N^3 points, enabling the use of the FFT. In the CDM case, as usual, the phase space of dark matter particles is sampled using (much) more massive "simulation particles", whose trajectories are evolved using Newtonian gravitational dynamics.

Concerning FDM SP simulations, it is important to make note of the tremendous computational requirements involved. Firstly, for the pseudo-spectral method, ensuring the absence of "aliasing" in the complex exponentials imposes a restriction

[1] The AxiREPO code will be made public in the near future (including the halo finder).

on the size of the time step Δt,

$$\Delta t < \min\left(\frac{4}{3\pi} \frac{m}{\hbar} a^2 \Delta x^2,\ 2\pi \frac{\hbar}{m} a \frac{1}{|\Phi_{\max}|} \right), \tag{5.2}$$

where $\Delta x = L/N$ is the spatial grid resolution and $\Phi_{\max}$ is the maximum value of the potential. Although the exact form of eq. (5.2) is specific to the pseudo-spectral method, the dependence $\Delta t \propto \Delta x^2$ seems to hold for all numerical approaches to the SP system of equations, eqs. (2.3) and (2.4), and serves as an illustration of the Schrödinger equation's nature as a diffusion equation.

Secondly, a constraint on the validity of the spatial discretization emerges from the relationship between the velocity field and the gradient of the wave function's phase (eq. (2.10)), which implies that velocities v larger than a certain value $v_{\max}$ cannot be represented in a simulation with spatial resolution Δx:

$$v_{\max} = \frac{\hbar}{m} \frac{\pi}{\Delta x}. \tag{5.3}$$

In other words, the "worst" spatial resolution $\Delta x_{\max}$ that still yields acceptable results is (roughly) given by the de Broglie wave length corresponding to $v_{\max}$:

$$\begin{aligned} \Delta x_{\max} &= \frac{\pi \hbar}{m v_{\max}} = \frac{1}{2} \lambda_{\mathrm{dB}}(v_{\max}) \\ \Leftrightarrow \Delta x &< \frac{\pi \hbar}{m v_{\max}}. \end{aligned} \tag{5.4}$$

Or, rephrasing yet again, the de Broglie wavelength must be resolved for all velocities appearing in the simulation. As with the time step criterion, similar considerations mandate resolving the de Broglie wavelength also for other numerical methods for the SP equations, yielding a similar constraint.

Combined, the time step criterion (5.2), $\Delta t \propto \Delta x^2$, and the velocity criterion (5.4), $\Delta x \propto 1/v$, render simulations of large cosmological objects or representative cosmological volumes prohibitively expensive. Contrary to N-body simulations, the additional requirement on the spatial resolution makes it impossible to resort to coarser, low-resolution simulations while still obtaining valid results on all (large) scales that are still resolved. Instead, a lack of sufficient resolution will affect results even on the largest scales and within the linear regime of structure formation, as shown in chapter 4.

Although numerical methods have been developed for FDM which are in principle

more versatile or sophisticated, such as AMR or hybrid techniques (Schwabe and Niemeyer 2022), the unique computational demands of FDM have the effect that, except for specific cases where very high resolution is desired in a small region ("zoom-in"), these methods still suffer from the tremendous cost involved in FDM simulations. While it may be argued that a uniform grid is wasteful in regions where a lower resolution would be sufficient, the velocity criterion (5.4) actually imposes significant demands on resolution even in low-density regions. The pseudo-spectral method used here, although not very flexible due to its limitation to a single uniform resolution, nevertheless has the advantage of unmatched accuracy in the spatial integration (at the given level of resolution), in addition to its use of a very simple and optimized algorithm in the form of the FFT, so it is at present one of the best ways to address the unique challenges of the SP system of equations.

5.3.1 Initial conditions

As in chapter 4, the ICs were generated at the starting redshift $z = 1/a - 1 = 127$ using the N-GenIC code (Springel 2015), which employs the Zel'dovich approximation (or second-order Lagrangian perturbation theory) to generate a random realization of density fluctuations consistent with a prescribed power spectrum in terms of a perturbed, but otherwise regular particle distribution. I analyze and compare four different kinds of simulations (cf. section 5.3.2), differentiated by the solved dynamics (FDM/Schrödinger–Poisson and CDM/N-body) and the initial power spectrum ("standard" ΛCDM and "self-consistent" FDM ICs with a cut-off at small scales). For each case, a different approach for the ICs was necessary:

- CDM (N-body) dynamics: In this case, the particle distribution generated by N-GenIC can simply be used directly.

- FDM (Schrödinger–Poisson) dynamics: The initial wave function is constructed using the prescription

$$|\psi(\vec{x})| = \sqrt{\frac{\rho(\vec{x})}{m}}, \tag{5.5}$$

$$\nabla \arg(\psi(\vec{x})) = \nabla\theta(\vec{x}) = \frac{m}{\hbar}\vec{v}(\vec{x}), \tag{5.6}$$

  where ρ is the matter density and $\vec{v}$ the velocity (cf. eqs. (2.6) and (2.10); Widrow and Kaiser 1993; Mocz et al. 2018; May and Springel 2021). In this case,

the random realization of the density fluctuation $\delta(\vec{x})$ in N-GenIC is actually used without creating a particle distribution at all, in order to avoid unnecessary transformation steps. For the absolute value of the complex wave function, the correspondence to the density via eq. (5.5) is straightforward. The wave function's phase requires a few more steps, with the following result using eq. (5.6) (in Fourier space)

$$\mathcal{F}[\theta](\vec{k}) = -i\frac{m}{\hbar}\frac{\vec{k}}{k^2} \cdot \mathrm{FFT}[\vec{v}](\vec{k}) \tag{5.7}$$

where $\vec{v}$ is the velocity field in the Zel'dovich approximation.

- CDM initial power spectrum: The simulations for this case are the ones previously presented in chapter 4. The ICs were generated with an input power spectrum following Efstathiou, Sutherland, and Maddox (1990) and Efstathiou, Bond, and White (1992), i. e. of the form

$$P_{\mathrm{CDM}}(k) \propto k\left[1 + \left(ak + (bk)^{\frac{3}{2}} + c^2k^2\right)^{\nu}\right]^{-\frac{2}{\nu}}, \tag{5.8}$$

with $a = 6.4/\Gamma\, h^{-1}\,\mathrm{Mpc}$, $b = 3.0/\Gamma\, h^{-1}\,\mathrm{Mpc}$, $c = 1.7/\Gamma\, h^{-1}\,\mathrm{Mpc}$, $\Gamma = \Omega_{\mathrm{m}}h = 0.21$,[2] and $\nu = 1.13$. Here $h = H_0/(100\,\mathrm{km\,s^{-1}\,Mpc^{-1}})$ encodes the Hubble constant.

- FDM initial power spectrum: In this case, I make use of the power spectrum predicted in linear theory by the code axionCAMB (Hlozek et al. 2015; Grin, Marsh, and Hlozek 2022) for an axion cosmology. However, in order to be able to compare the simulations as directly as possible, the aim is to minimize differences in the CDM and FDM ICs arising from the two different methods (a fitting function eq. (5.8) and a numerical Boltzmann code). To this end, instead of using the power spectrum $P_{\mathrm{axionCAMB,FDM}}(k)$ from axionCAMB directly, I also calculated the corresponding ΛCDM power spectrum $P_{\mathrm{axionCAMB,CDM}}(k)$ using the same code and obtained the final input spectrum $P_{\mathrm{FDM}}(k)$ as follows:

$$P_{\mathrm{FDM}}(k) = \frac{P_{\mathrm{axionCAMB,FDM}}(k)}{P_{\mathrm{axionCAMB,CDM}}(k)} P_{\mathrm{CDM}}(k). \tag{5.9}$$

[2] See section 5.3.2.

Table 5.1: List of performed simulations with important characteristics: simulation type (SP/*N*-body with FDM or CDM ICs), number of resolution elements (grid cells or *N*-body particles), box size, FDM particle mass, and resolution (grid cell size or *N*-body particle mass). The lengths given for the box sizes and resolutions are comoving. The simulations with CDM ICs have previously been presented in chapter 4.

Type	IC	Res. el.	L / h^{-1} Mpc	mc^2 / eV	Resolution
SP	FDM	8640^3	10	7×10^{-23}	$1.16\,h^{-1}$ kpc
SP	FDM	6144^3	10	5×10^{-23}	$1.63\,h^{-1}$ kpc
SP	FDM	4320^3	10	3.5×10^{-23}	$2.31\,h^{-1}$ kpc
SP	FDM	4320^3	10	2×10^{-23}	$2.31\,h^{-1}$ kpc
SP	FDM	4320^3	10	1×10^{-23}	$2.31\,h^{-1}$ kpc
SP	CDM	8640^3	10	7×10^{-23}	$1.16\,h^{-1}$ kpc
SP	CDM	4320^3	10	3.5×10^{-23}	$2.31\,h^{-1}$ kpc
N-body	FDM	2048^3	10	7×10^{-23}	$9.69 \times 10^3\,h^{-1}\,M_\odot$
N-body	FDM	2048^3	10	5×10^{-23}	$9.69 \times 10^3\,h^{-1}\,M_\odot$
N-body	FDM	2048^3	10	3.5×10^{-23}	$9.69 \times 10^3\,h^{-1}\,M_\odot$
N-body	CDM	2048^3	10	—	$9.69 \times 10^3\,h^{-1}\,M_\odot$

This ensures that any variations between eq. (5.8) and axionCAMB are canceled out, and only the differences arising from the choice of the CDM or FDM model remain.

5.3.2 Simulations

The simulations were performed using the cosmological parameters $\Omega_\mathrm{m} = 0.3$, $\Omega_\mathrm{b} = 0$, $\Omega_\Lambda = 0.7$, $H_0 = 70\,\mathrm{km\,s^{-1}\,Mpc^{-1}}$ ($h = 0.7$), and $\sigma_8 = 0.9$,[3] with ICs as described in section 5.3.1. The same cosmological parameters and random seeds were shared for the generation of all ICs at $z = 127$ in order to allow for a direct comparison between different dark matter models and resolutions, including those previously presented in chapter 4. For a comoving box size of $10\,h^{-1}$ Mpc, and different masses m, simulations with different resolutions – up to grid sizes of $N^3 = 8640^3$ – were performed. As before, the simulations were run until $z = 3$, where the FDM simulations should still be largely unaffected by resolution effects. Furthermore, even the largest modes in simulations of this box size would become non-linear before $z = 0$, making the reliability and benefit of evolving simulations to this point dubious in any case. A

[3] As usual, the density parameters Ω_i indicate the values at $z = 0$.

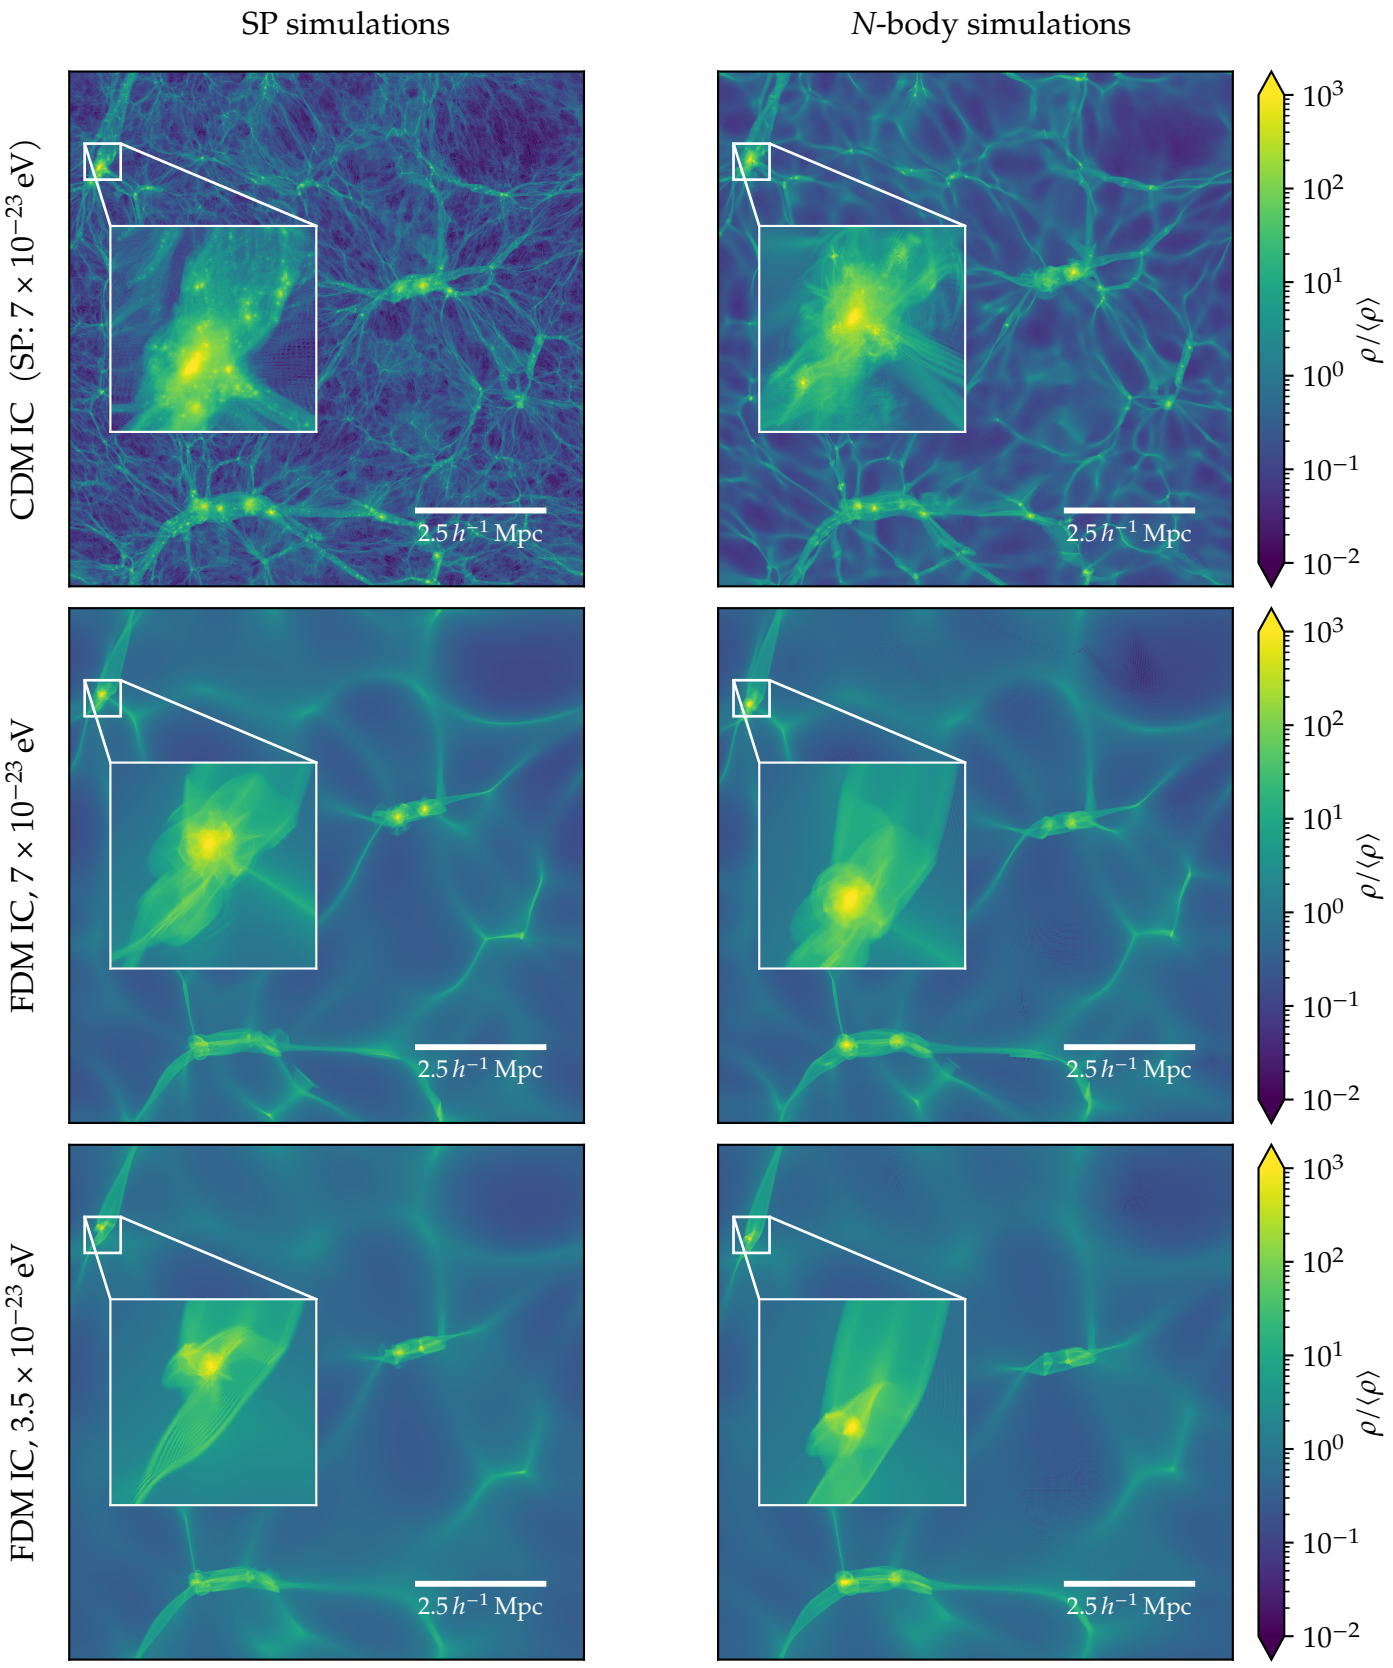

Figure 5.1: Projected dark matter densities along a thin ($100\,h^{-1}$ kpc) slice in cosmological box simulations for different dynamics (SP, left column; or N-body, right column) and ICs (FDM or CDM, rows), for box sizes $L = 10\,h^{-1}$ Mpc at $z = 3$. Insets show magnified views of a small sub-region, highlighting scales where FDM wave effects are more prominent.

detailed list of the different simulations is given in table 5.1. A visual overview showing the projected density across a slab of the simulation volume is displayed for a subset of the simulations in fig. 5.1. The images readily make differences between the four different kinds of simulations I have carried out apparent, and they show the impact of different values of m as well.

Most of the simulations listed in table 5.1 focus on the masses $mc^2 = 7 \times 10^{-23}$ eV and 3.5×10^{-23} eV, respectively, used in chapter 4, but I also performed additional "self-consistent" FDM simulations (SP + FDM ICs) for values of $mc^2 = 5 \times 10^{-23}$ eV, 2×10^{-23} eV and 10^{-23} eV in order to investigate the dependence of FDM phenomenology on the particle mass. For the lighter masses 10^{-23} eV and 2×10^{-23} eV, the resolution (with respect to eq. (5.4)) is actually much better than that of the other simulations, allowing me to evolve them further in time beyond $z = 3$ (although once again with the caveat that all spatial modes eventually become non-linear). For N-body simulations, a very high mass resolution (large number of particles) was used when compared to most similar cosmological simulations. This was done in order to be able to measure the power spectrum down to very small scales, comparable to those accessible using a 8640^3 grid. Due to my earlier work in chapter 4 using the same parameters, it was not necessary to investigate the numerical convergence of the simulations again in detail.

Although progress has been made e. g. in hybrid methods (Schwabe and Niemeyer 2022), which do not solve the equations of motion for the full wave dynamics everywhere, my simulations remain the largest cosmological SP simulations of structure formation with FDM. The new simulations with FDM ICs allow me to perform a four-way comparison between SP and N-body dynamics on the one hand, and FDM and CDM ICs on the other hand, and to quantify the impact of both aspects. Apart from being able to study a self-consistent FDM cosmology, this also allows me to gain insight into the extent to which FDM can be approximated (e. g.) by using N-body simulations with a modified initial power spectrum, and to examine the validity of previous work which took this approach (e. g. Schive et al. 2016).

5.3.3 Halo identification with suppressed small-scale power

In order to identify dark matter halos, the FoF-like halo finder developed in May and Springel (2021) (cf. chapter 4), which is able to work on a Cartesian grid, was used. Instead of a linking length for particle distances, this grid-based halo finder

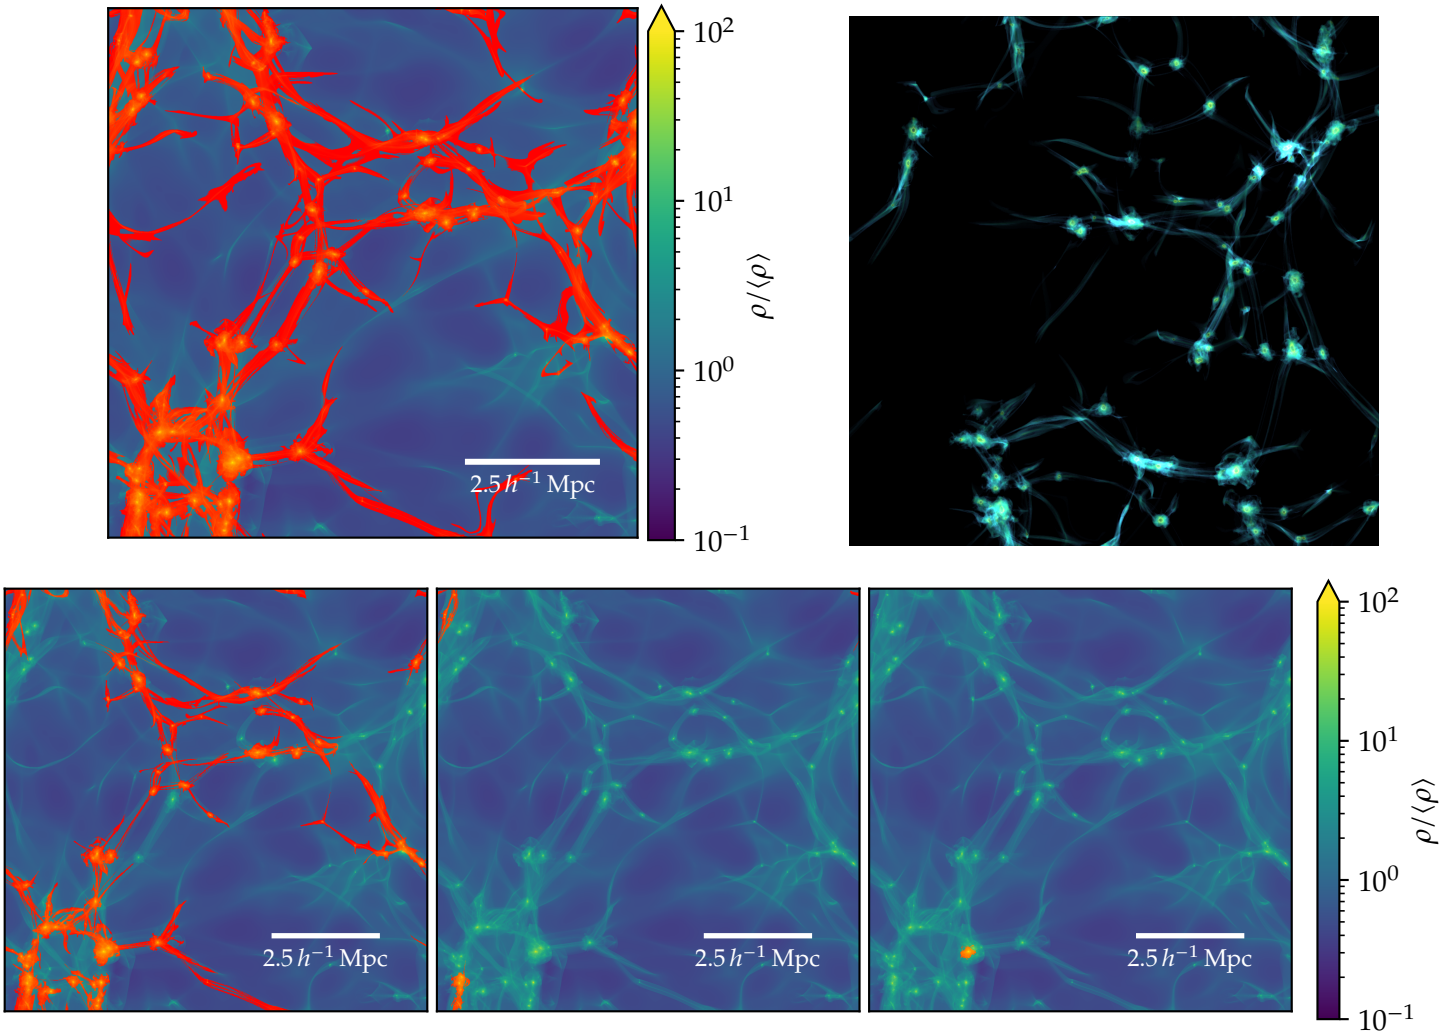

Figure 5.2: Projected dark matter density and volume rendering (different viewing angle to illustrate 3D structure) showing the largest FoF group using a halo finder overdensity threshold of $60\rho_\mathrm{m}$ (top row), and $100\rho_\mathrm{m}$, $200\rho_\mathrm{m}$ and $300\rho_\mathrm{m}$ (bottom row) in a cosmological box simulation of FDM at $z = 3$ with box size $L = 10\,h^{-1}\,\mathrm{Mpc}$, grid size $N^3 = 8640^3$, FDM mass $mc^2 = 7 \times 10^{-23}\,\mathrm{eV}$, and FDM ICs. The areas marked in red/orange in the projections and shown in the volume rendering indicate regions spanned by the largest FoF group identified by the halo finder using the given overdensity threshold.

uses a density threshold as an analogous parameter. Only grid cells above the given overdensity threshold are considered, and they are always linked if they are adjacent.

While this approach worked very well for SP simulations with CDM-like ICs and an overdensity threshold of $60\rho_\mathrm{m}$, the results with FDM ICs were very different. In this case, halos are linked via *continuous, smooth, dense filaments* throughout the entire

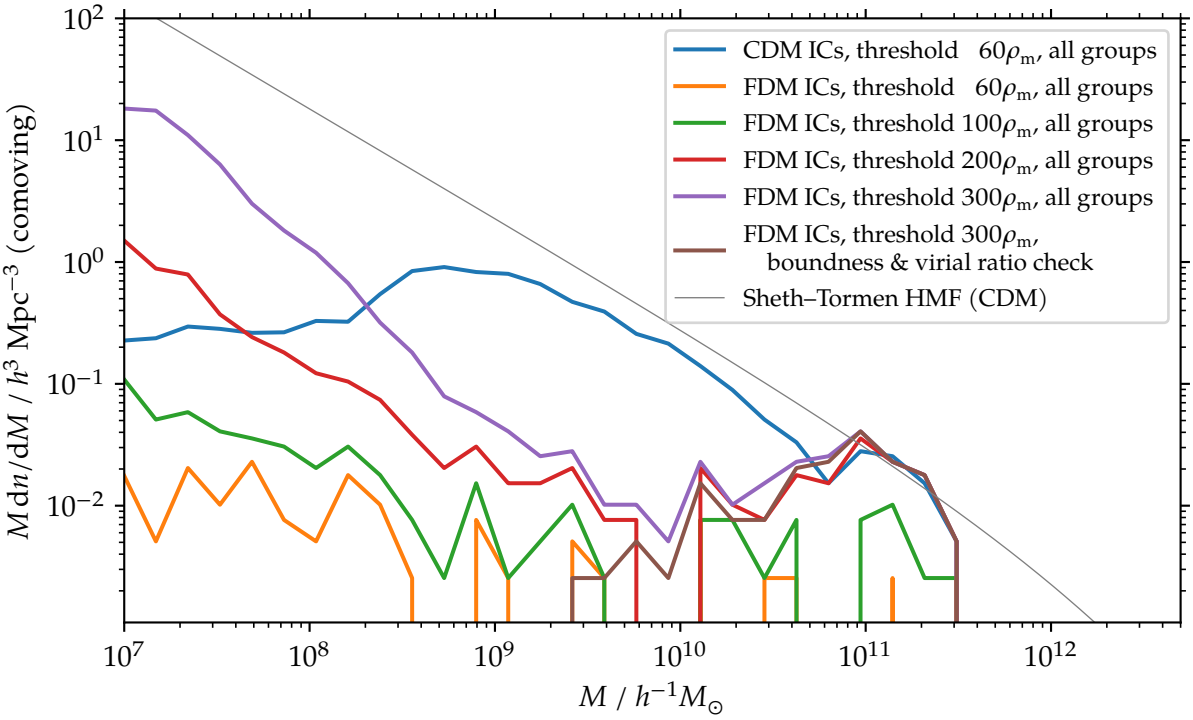

Figure 5.3: Halo mass function of a cosmological FDM simulations at $z = 3$ with a box size of $L = 10\,h^{-1}$ Mpc and FDM ICs for different values of the halo finder's overdensity threshold and after filtering using the gravitational binding criterion. Different simulations are shown as labeled.

simulation volume. The central densities of these filaments exceed the threshold value of $60\rho_\mathrm{m}$, leading to much larger regions of space being linked by the halo finder than intended. Indeed, as illustrated in fig. 5.2, the largest FoF group traces a network of filaments and spreads across the entire simulation box, incorporating numerous halos and filaments. This has dire consequences for the endeavor of actually identifying individual halos: Because each FoF group is counted at most as a single halo, the largest group subsumes many halos and thus leads to a severe under-counting. As shown in fig. 5.3, the resulting HMF shows very few or even no halos at all across wide mass ranges.

In order to break the filamentary links between halos, it was necessary to use a higher value for the overdensity threshold. Although the FoF groups are smaller as a consequence, encompassing only the denser inner parts of halos, this removes the filaments, which do not typically reach such high densities, from the groups. Empirically, I arrived at a value of $300\rho_\mathrm{m}$ to reliably prevent spurious connections between and merging of FoF groups (cf. fig. 5.3).

However, this countermeasure invokes another problematic effect: Now, FoF groups are broken up into many small parts which do not all correspond to actual and reasonably complete halos. Figure 5.3 shows how increasing the group finder's threshold leads to an ever larger proliferation of low-mass objects in the HMF. Clearly, these objects are spurious artifacts, since, as visible in fig. 5.3, their number keeps increasing with the threshold parameter, eventually even exceeding the CDM ICs case, which can be viewed as an upper limit for the HMF expected for FDM ICs.

On visual inspection of a sample of these objects, it turned out that the vast majority of them are localized density fluctuations arising from the wave nature of FDM, e. g. in the form of constructive interference fringes, which can reach very high densities. The top panel of fig. 5.4 shows how a large number of transient interference maxima is erroneously identified as halos. Although the same patterns are present in the SP simulations with CDM ICs, the lower threshold value of $60\rho_\mathrm{m}$ ensured that these objects remained connected to halos within a single FoF group. As such, it should be that while the identification of these objects by the halo finder is a consequence of the used algorithm in combination with the chosen value of the density threshold parameter, they are *physical* (if transient) objects, contrary to the "spurious halos" encountered e. g. in N-body simulations of WDM. In turn, SP simulations (regardless of ICs) do not suffer from the spurious fragmentation encountered in N-body simulations.

In the end, filtering this set of FoF groups using a gravitational binding criterion proved necessary and successful for reliably identifying halos.[4] To this end, I compute the "self-potential" of each FoF group within R_{200}, i. e. the gravitational potential V_s generated only by the matter within R_{200},[5]

$$V_\mathrm{s} = -\frac{1}{2} \sum_{\substack{i,j \\ i\neq j}}^{|\vec{r}_i - \vec{r}_j| < R_{200}} \frac{G m_i m_j}{|\vec{r}_i - \vec{r}_j|}, \tag{5.10}$$

as well as the kinetic energy T_v in the group's center-of-mass frame using eq. (2.8) and the gradient energy T_ρ, and excluded any FoF group which did not meet the

[4] Another check is to filter any spatially overlapping halos in order to avoid double-counting or potentially mis-identifying sub-halos as halos, but with the data presented here, this was not necessary after applying the gravitational binding criterion.

[5] The self-potential calculation was here done using direct summation, approximating each grid cell as a point particle for simplicity, which is accurate up to hexadecupole order in the multipole expansion (Barnes and Hut 1989).

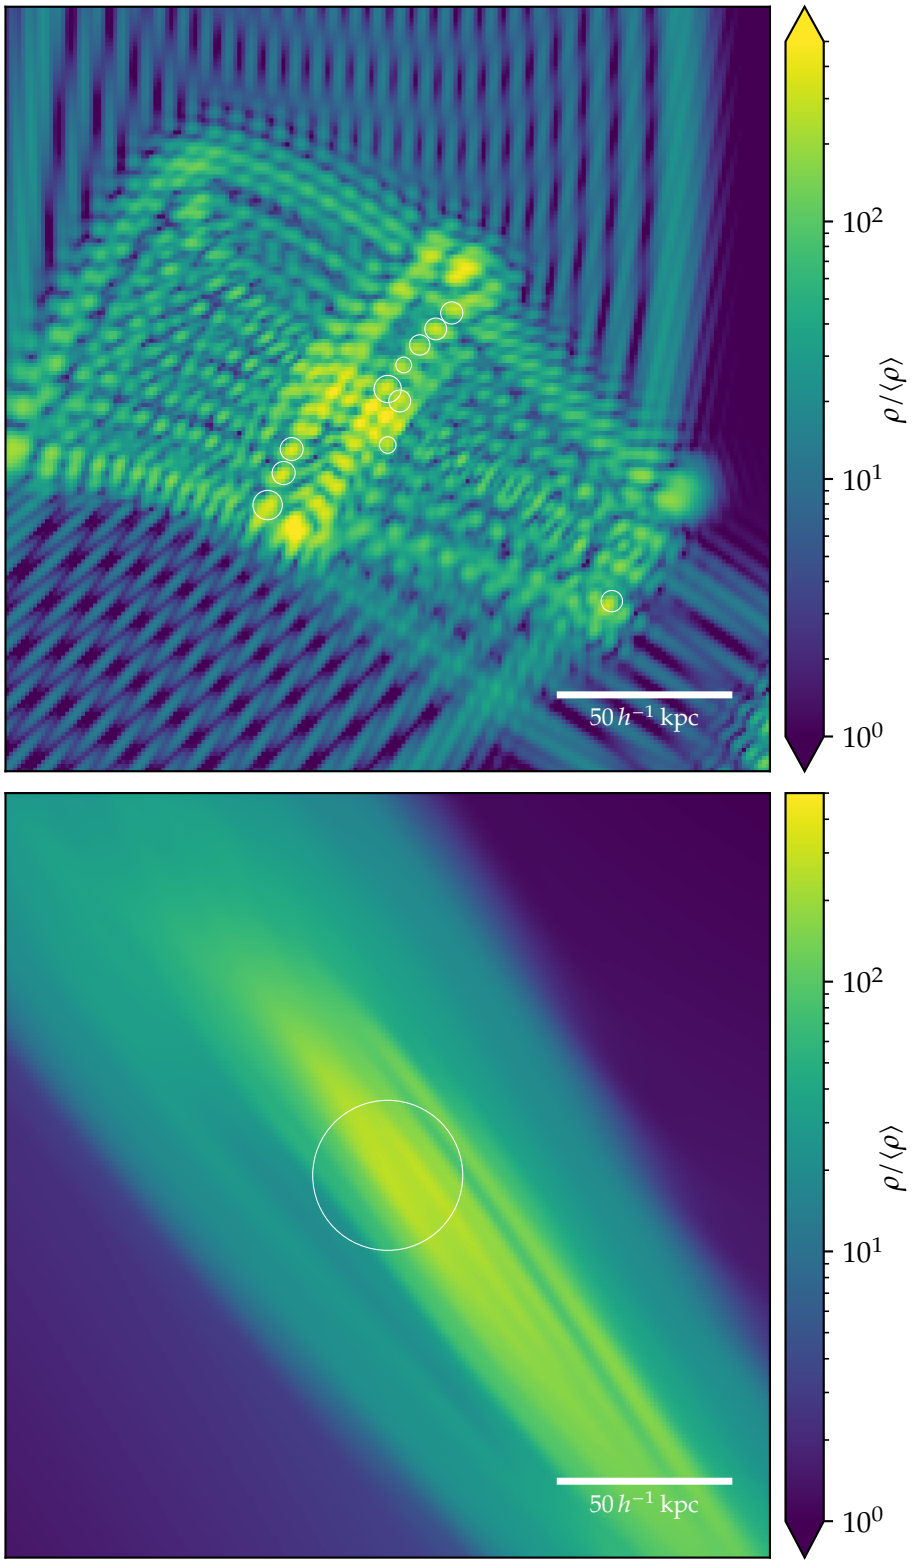

Figure 5.4: Examples of spurious FoF groups identified in FDM interference patterns (top panel) and filaments (bottom panel) in a cosmological FDM simulation (projected density). The former arise in great numbers as a consequence of raising the halo finder's overdensity threshold to $300\rho_\mathrm{m}$, while the latter appear due to a lack of small-scale power analogous to WDM. The circles mark the groups' locations, with the radius corresponding to R_{200}.

binding criterion

$$E = T_v + T_\rho + V_\mathrm{s} < 0, \tag{5.11}$$

cf. eq. (2.12). This reliably excludes the gravitationally unbound, transient wave interference patterns.

It should be noted that the computation of velocities from the complex wave amplitude ψ using the gradient of the phase as in eq. (2.10) is not straightforward. Not only is the phase a periodic variable, making it necessary to take into account wrap-around at 2π when computing gradients, but it is also undefined in regions of destructive interference, where $\rho \to 0$. In the end, the most robust method turned out to be using the (equivalent) definition via the momentum density in eq. (2.8). Not only does this avoid the treatment of the periodic phase variable, but it is also well-defined even for $\rho = 0$, preventing pathological cases (extremely large velocities) when computing the numerical derivative. Although obtaining the velocity from the momentum density requires dividing by ρ, the results of this procedure turn out to be much more well-behaved in practice.

Finally, in addition to the above, a stricter cut was enforced on the virial ratio $-2T/V$ in the form of

$$-2T/V_\mathrm{s} < 1.6 \tag{5.12}$$

in order to eliminate a number of objects which were clearly highly perturbed objects and not reasonably relaxed halos.[6] Similar cuts are regularly applied in CDM models when density profiles of halos are studied (e. g. Neto et al. 2007). These objects are distinct from the wave interference phenomena, and instead could be called very high-density filamentary structures, arising from the lack of power on small scales in the FDM ICs case. Their high virial ratios indicate that, while gravitationally bound, they have not reached a virialized state. In contrast, for CDM, such filaments do not remain smooth, but fragment into (sub-)halos down to the smallest scales. Examples of both kinds of non-halo objects are shown in fig. 5.4.

Figure 5.5 shows the virial ratios for FoF groups in an FDM simulation as a function of M_{200} (determined using a spherical overdensity algorithm). Although a virial ratio of ≈ 1, corresponding to $2T \approx |V|$, would be expected for virialized objects, there is a significant bias of many objects around $-2T/V_\mathrm{s} \approx 1.2$. Indeed, this effect has been observed for CDM halos in N-body simulations as well (Bett et al. 2007; Neto et al. 2007), and seems to be due to the definition of the self-potential, which

[6] Since eq. (5.11) corresponds to $-2T/V_\mathrm{s} < 2$, eq. (5.12) automatically enforces gravitational binding as well.

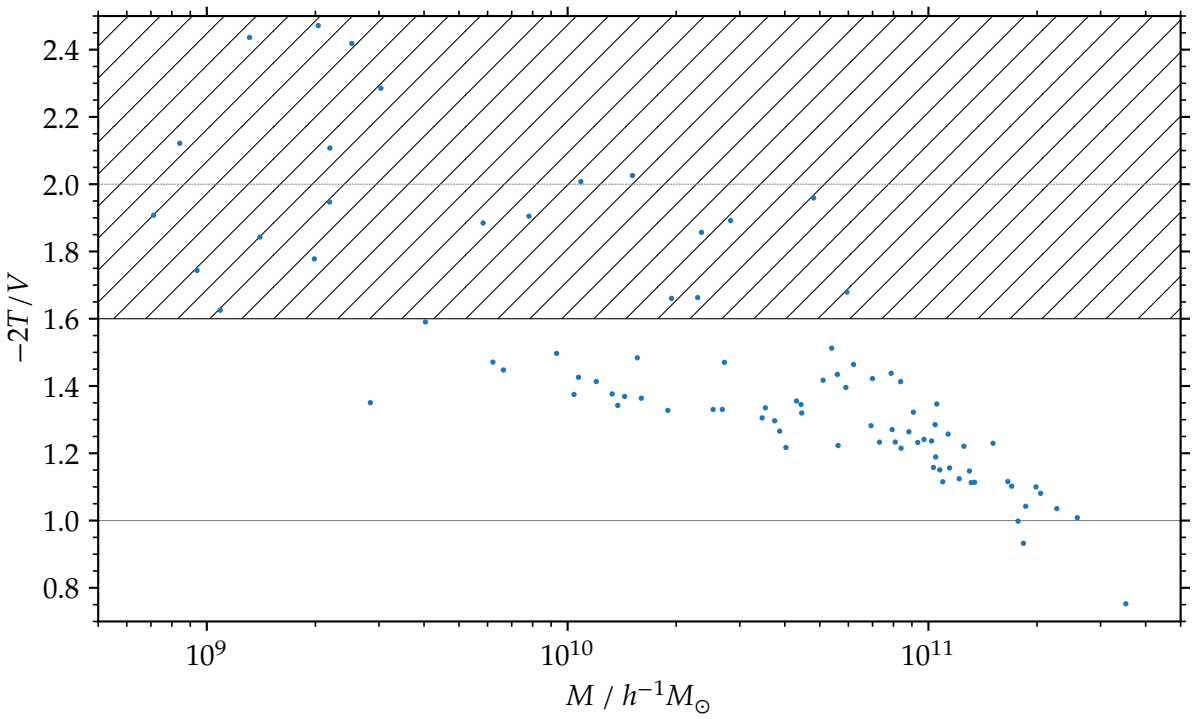

Figure 5.5: Virial ratios $-2T/V$ for FoF groups identified by AxiREPO's halo finder with an overdensity threshold of $300\rho_\mathrm{m}$ in a cosmological FDM simulation at $z = 3$ with $L = 10\,h^{-1}\,\mathrm{Mpc}$, $N^3 = 8640^3$, $mc^2 = 7 \times 10^{-23}\,\mathrm{eV}$, and FDM ICs. The potential energy V is the "self-potential", and both T and V are summed within a radius of R_{200} for each halo. Points above the dashed line correspond to FoF groups which are not gravitationally bound ($E = T + V > 0$). Points in the hatched region ($-2T/V > 1.6$) are excluded from my definition of "halos".

neglects the contribution from a possible large-scale tidal field (Stücker, Angulo, and Busch 2021). A more sophisticated approach, such as the "boosted potential binding check" from Stücker, Angulo, and Busch (2021), can potentially improve the results, but was not necessary in this case, perhaps due to the low number of remaining halos after selection, which allows for easy manual inspection of the end result. Notably, these problems of (1) enormous FoF groups, linked by smooth filaments, stretching across 10s of Mpc, and (2) additional, non-virialized structures identified by the halo finder, have been encountered in N-body simulations of WDM before (Angulo, Hahn, and Abel 2013, especially figs. 2 and 4). There have also been speculations that stars themselves may form in WDM filaments (Gao and Theuns 2007) (and similarly for FDM, Mocz et al. 2020), giving rise to a qualitatively different mode of galaxy formation compared to CDM.

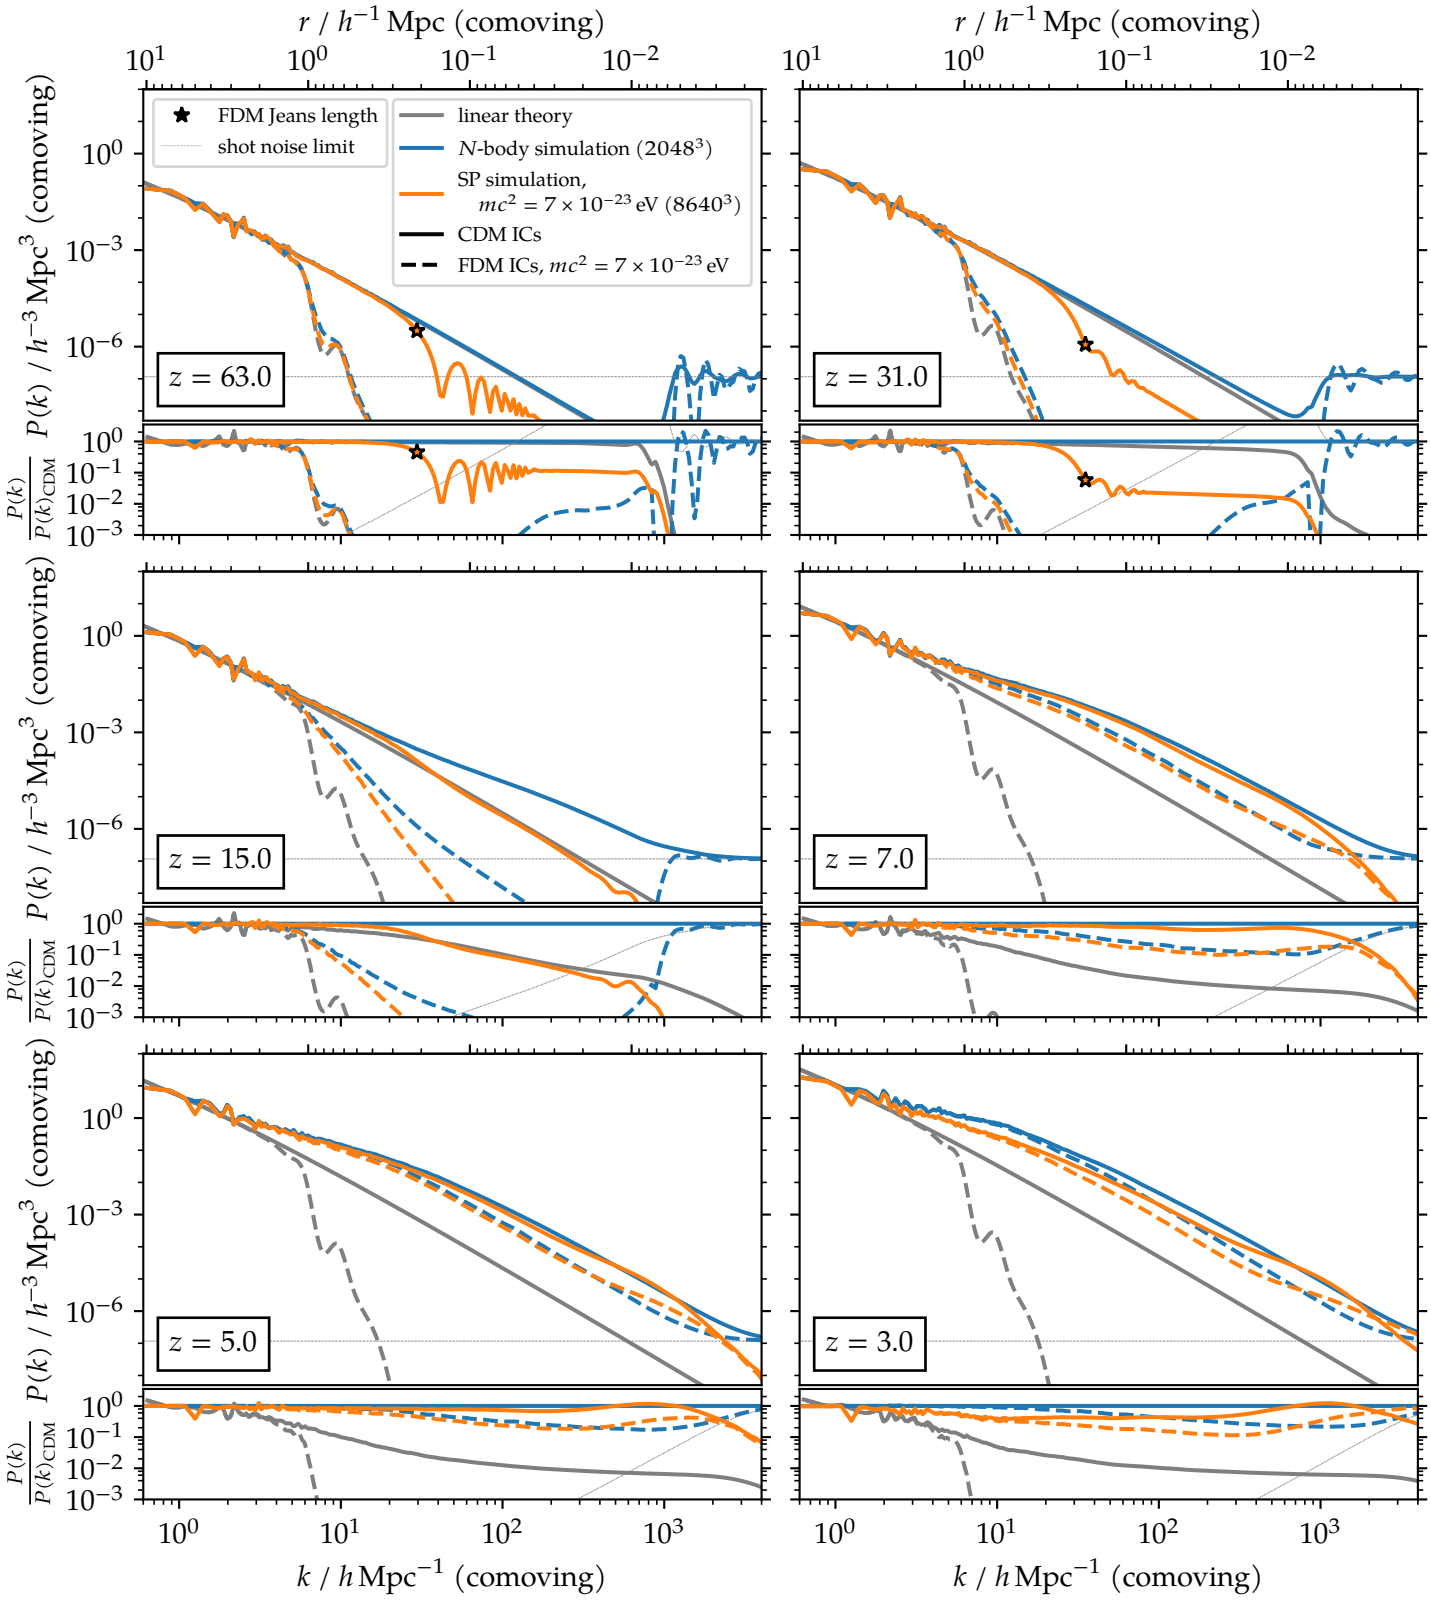

Figure 5.6: Four-way comparison of dark matter power spectra at different redshifts for cosmological FDM (wave) and CDM (N-body) simulations with FDM and CDM ICs in $L = 10\,h^{-1}$ Mpc boxes. The power spectrum evolved using linear perturbation theory (scale-independent growth factor) is shown for comparison. The lower panels show the ratio of the power spectra to the CDM result (N-body simulation with CDM ICs). For $z = 63$ and 31, a star additionally indicates the FDM Jeans scale, eq. (2.14). Faint dotted lines show the shot noise limits of the N-body simulations; the power spectrum cannot be measured accurately once this limit is reached.

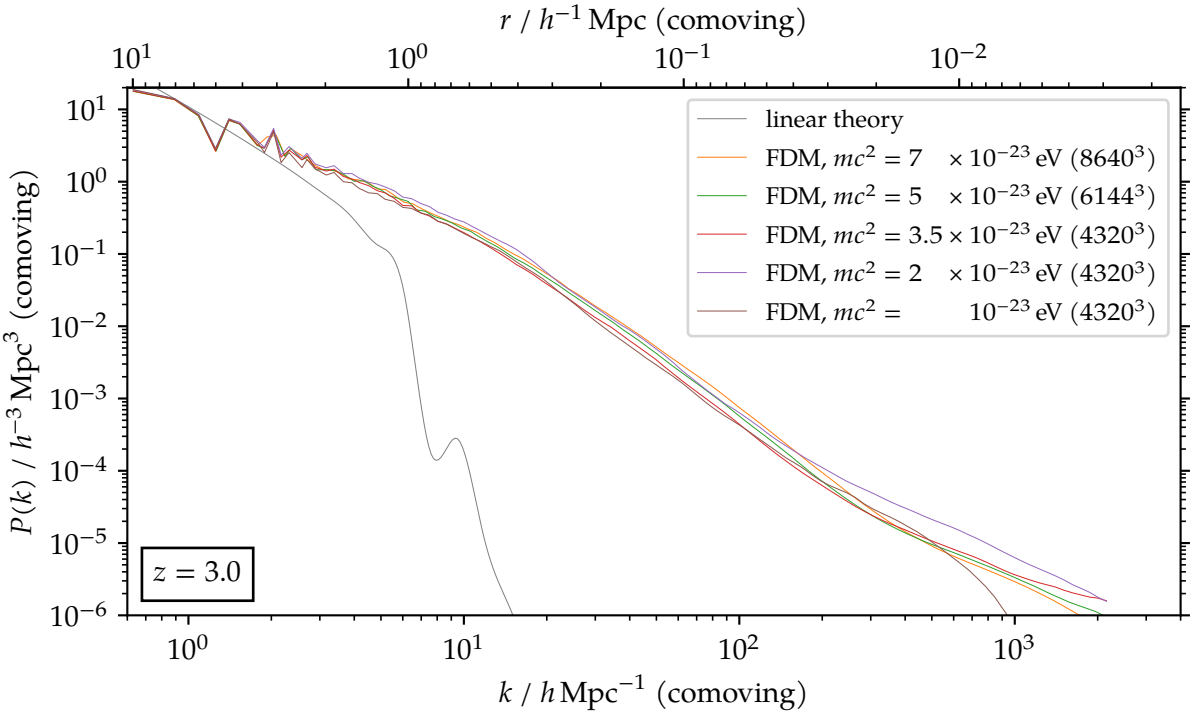

Figure 5.7: Dark matter power spectra for cosmological FDM simulations at $z = 3$ with different particle masses m in $L = 10\,h^{-1}$ Mpc boxes. The power spectrum evolved using linear perturbation theory (scale-independent growth factor) is shown for comparison.

5.4 Four-way comparison of the power spectrum

The matter power spectrum is a crucial measure of matter clustering at different length scales. For ΛCDM, even including baryonic effects, its behavior and temporal evolution have been determined rather accurately using numerical simulations (e. g. Jenkins et al. 1998; Hellwing et al. 2016; Springel et al. 2018). For FDM, linear structure formation is scale-dependent even in linear perturbation theory, with a suppression at scales smaller than the FDM Jeans length $\lambda_\mathrm{J} = 2\pi/k_\mathrm{J}$ (2.14), and even during non-linear evolution on scales where the wave nature of FDM is relevant, i. e. at or below the de Broglie wavelength λ_dB (2.13). Intuitively, this can be understood as a consequence of the analogue of the Heisenberg uncertainty principle (Schrödinger formulation), or, equivalently, the presence of the so-called "quantum pressure" (Madelung fluid formulation).

Combined with my previous results (chapter 4), my new simulations with FDM ICs, using both SP and N-body solvers, allow me to draw a four-way comparison of observables, contrasting results for each combination of ICs (FDM vs. CDM) and

dynamics (SP/"FDM" vs. *N*-body/"CDM"). Although only two of the four cases are strictly self-consistent in a physical sense (namely, FDM and CDM with their corresponding ICs), this approach allows me to disentangle the two essential physical differences distinguishing FDM from CDM in cosmological numerical simulations: the initial conditions and the dynamics on small scales. This is particularly relevant to evaluate to what extent approximate methods, such as the use of *N*-body simulations with an FDM initial power spectrum (Schive et al. 2016; Mocz et al. 2020), can yield reasonable results.

Figure 5.6 shows the measured matter power spectra for both FDM and CDM cases, where simulations with $mc^2 = 7 \times 10^{-23}\,\mathrm{eV}$ are used to represent FDM. The solid lines (CDM ICs) correspond to the results from chapter 4, whereas the dashed lines are the new simulations with FDM ICs added in this chapter. Reassuringly, all cases accurately agree with each other and with linear perturbation theory on large scales and at early times, and it can be observed again that the onset of non-linear structure formation is delayed for FDM, and does not proceed strictly in the same bottom-up fashion as for CDM.

As a reflection of this fact, the differences between the four cases exhibit variations across time. For example, only the *N*-body CDM simulation significantly exceeds linear growth at redshifts below 15, although the SP CDM case has similar, mild non-linear amplification on large enough scales of $k \approx 10\,h\,\mathrm{Mpc}^{-1}$ to $20\,h\,\mathrm{Mpc}^{-1}$ at $z = 15$. However, by $z = 7$, considerable non-linear enhancement is present on small scales $k \gtrsim 5\,h\,\mathrm{Mpc}^{-1}$ in all cases.

Indeed, for $z \geq 7$, the power spectra for simulations with the same ICs track each other comparatively closely on all scales, although the SP simulations are still suppressed by some tens of percent for $k \gtrsim 10\,h\,\mathrm{Mpc}^{-1}$. There seems to be little qualitative change in the relative evolutions of the power spectra beyond $z = 7$. Notably, for the most part during this time, the difference between SP and *N*-body is significantly smaller than that between the different sets of ICs, and is of similar size in both cases.

At $z = 3$, SP and *N*-body results drift more apart compared to earlier times ($z = 7$, $z = 5$). Interestingly, at this time, the relative difference between the dynamics and the ICs is now roughly the same, meaning that each of the two physical "ingredients" has a similar impact at this time. Since the time evolution stops at $z = 3$, it is however unclear whether this is the onset of a new, sustained phenomenological difference or whether it is a transient, numerical artifact. Since the resolution requirements for SP grow more stringent towards later times, it is also possible that these slight

relative changes in the power spectra are the first hints of resolution effects affecting the results.

An interesting difference for the SP simulations with different ICs is the location of the "bump" on small scales, which starts to appear at $z = 7$. At the location of this bump, the small-scale power in SP simulations matches or even slightly exceeds the CDM power (for FDM ICs, the SP simulation quite significantly surpasses its N-body companion on these scales). It is commonly attributed to the presence of FDM interference patterns and granules of size λ_{dB} (Mocz et al. 2020), which are unique to the wave nature of FDM and not present in CDM. However, curiously, there is a shift in the location of this bump for the two SP simulations, with the bump in the FDM ICs case appearing on smaller scales. Following the interpretation that the bump is linked to λ_{dB}, this would imply that the FDM simulation displays more power on smaller wavelengths, which (with the other parameters being equal) would correspond to higher velocities. The physical origin of this effect remains unclear and is a subject for future investigation.

Figure 5.7 shows the dependence of the power spectrum on different values of the particle mass m at $z = 3$. In part due to the limited range in values of interest which are feasible to simulate, the differences are somewhat minor. Slightly deviating from the trend are the lowest-mass simulations with $mc^2 = 10^{-23}\,\mathrm{eV}$ and $2 \times 10^{-23}\,\mathrm{eV}$, whose resolution relative to the de Broglie wavelength is higher than that of the other simulations. In particular, at very small scales, the simulation with $2 \times 10^{-23}\,\mathrm{eV}$ surpasses even the higher-mass cases. This is an indication that the comparably "lower-resolution" simulations are starting to be mildly impacted by resolution effects at these scales. Notably, only the $10^{-23}\,\mathrm{eV}$ case exhibits a distinct small-scale cut-off in the spectrum, demonstrating that extremely high ("relative") resolution is required to obtain a converged power spectrum measurement at these small scales.

5.5 The halo mass function

The formation of gravitationally collapsed structures of dark matter, so-called halos, is one of the most significant results of cosmic structure. In particular, halos serve as the loci of collapse of baryonic matter and thus provide the environment for the formation and evolution of galaxies. The HMF, which is a measure of halo abundance as a function of mass, and its evolution across time, are thus critical benchmarks of a cosmological model.

While the (extended) Press–Schechter formalism (Press and Schechter 1974; Sheth and Tormen 1999) has been demonstrated to be a simple and reliable method to estimate the HMF and its evolution (depending only on the power spectrum and growth factor in *linear* perturbation theory), with its validity being confirmed by comparisons to non-linear N-body simulations (which has also allowed for the calibration of more accurate empirical fits, e.g. Jenkins et al. 2001; Tinker et al. 2008; Despali et al. 2016), there is no similar, simple approximation for FDM. None of the few existing (semi-)analytic estimates, such as using a Jeans-filtered power spectrum to account for the effects of "quantum pressure" (Marsh and Silk 2014; Du, Behrens, and Niemeyer 2017), or a sharp k-space filtering in the Press–Schechter formalism with a variable cut-off (Kulkarni and Ostriker 2022), have been verified using the full non-linear SP evolution, rendering their quantitative reliability still unclear. The closest substitution of this goal thus far has been reached by comparing to the FDM HMF estimate in Schive et al. (2016), which employed the approximate technique of collisionless N-body simulations with truncated initial fluctuation spectrum (but *without* the full SP dynamics), similar to those performed in this chapter as well ("N-body with FDM ICs"). This type of simulation technique is also how WDM models are often computed (e. g. Lovell et al. 2014), but importantly it exhibits a number of by now well-known difficulties (Wang and White 2007), such as the formation of "spurious" low-mass halos that necessitate special removal procedures and ultimately introduce a significant source of uncertainty.

Supplementing my previous results (chapter 4), my new simulations allow for a measurement of the HMF in a fully self-consistent cosmological wave simulation of FDM for the first time. For this purpose, I made use of the FoF-based grid halo finder in AxiREPO to identify collapsed structures in my SP simulations. However, as discussed in section 5.3.3, the cut-off in the FDM initial power spectrum actually introduces considerable complications in identifying bound structures – some of which similarly plague WDM simulations (Wang and White 2007; Angulo, Hahn, and Abel 2013) –, which mandate the application of additional filtering steps on the initial group catalog determined by the halo finder. After implementing these steps, it turns out that only a limited number of halo candidates in the raw catalog survive in my FDM simulation and can be considered physically robust, bound structures. Projections of these 68 halos at $z = 3$ are displayed in fig. 5.8, making it clear that the majority of them exhibit a morphology in their outer parts that is quite distinct from what is typically seen in CDM simulations. In particular, the halos are usually embedded in thick filaments that show clear patterns of interference ridges.

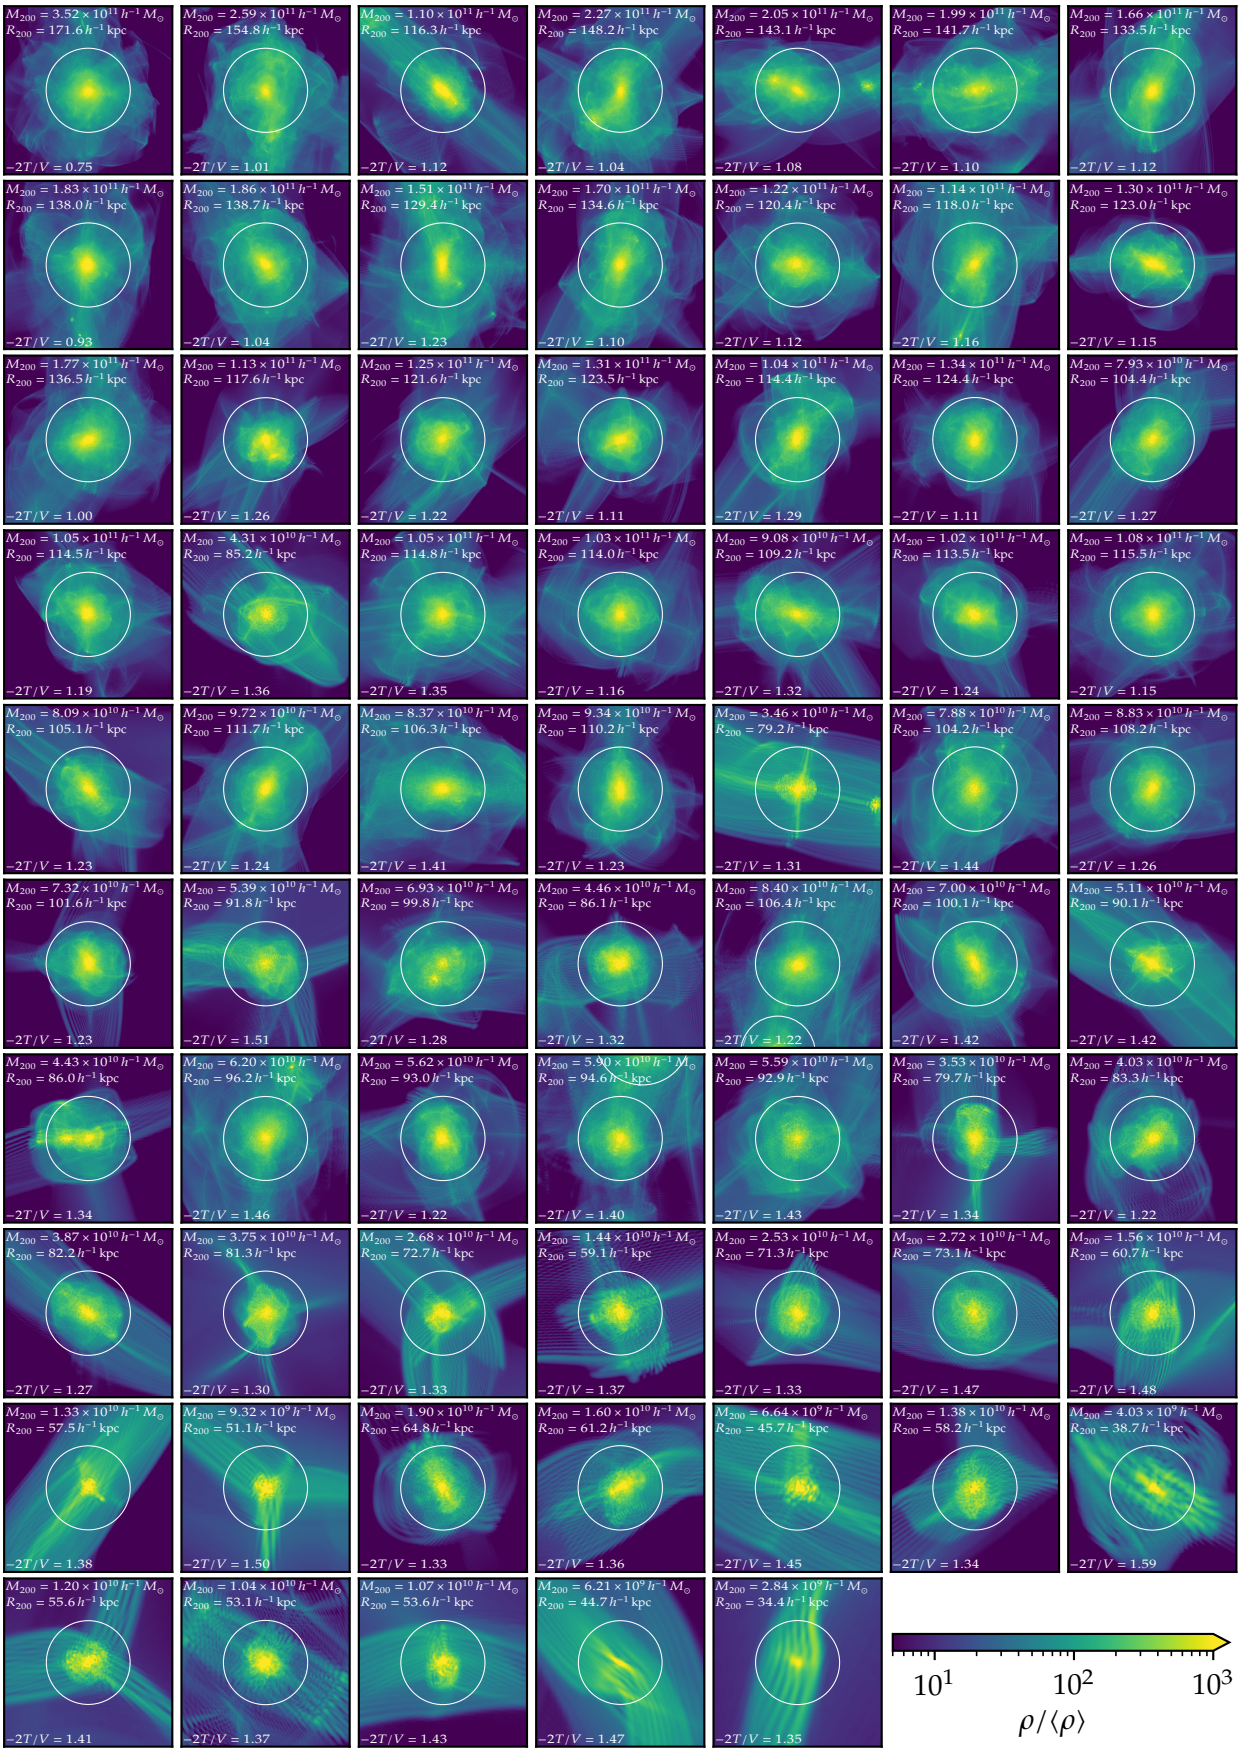

Figure 5.8: Projected dark matter densities for all 68 identified halos at $z = 3$ in a cosmological FDM simulation with box size $L = 10\,h^{-1}$ Mpc, $N^3 = 8640^3$, $mc^2 = 7 \times 10^{-23}$ eV, and FDM ICs. The projection depth is $2R_{200}$ for each halo. Circles with radius R_{200} (comoving) are centered on halos.

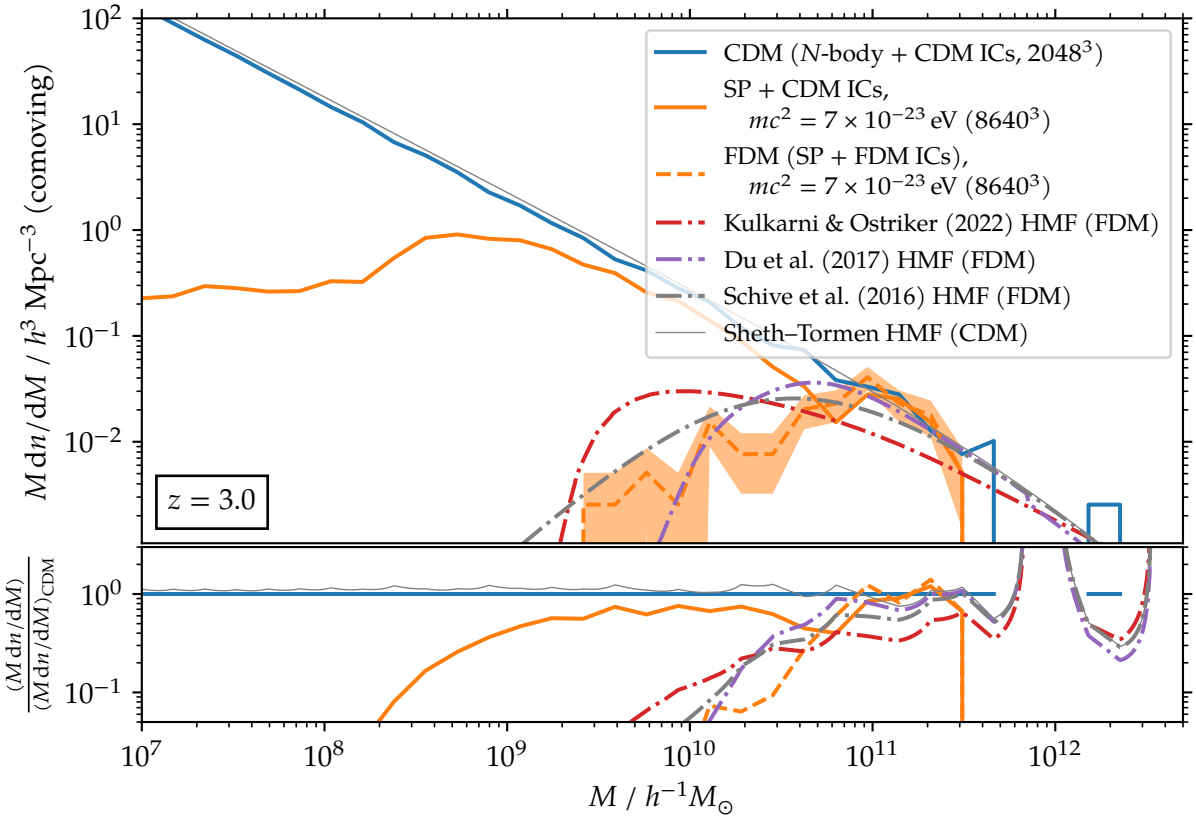

Figure 5.9: Halo mass functions in cosmological FDM and CDM simulations at $z = 3$ with FDM mass $mc^2 = 7 \times 10^{-23}\,\mathrm{eV}$, a box size of $L = 10\,h^{-1}\,\mathrm{Mpc}$ and different ICs. The HMF derived for CDM by Sheth and Tormen (1999) is shown for comparison. Different predictions for the FDM HMF from Schive et al. (2016), Du, Behrens, and Niemeyer (2017), and Kulkarni and Ostriker (2022) are shown using dash-dotted lines. The lower panel shows the ratios of the mass functions to the result of the CDM simulation.

Figure 5.9 shows the resulting HMF for the new FDM simulations compared to the mass functions for CDM ICs. The results for the HMF make it clear that the ICs have a much stronger impact than the choice of SP or N-body dynamics. Due to the cut-off on small scales in the initial FDM power spectrum, the seeds of structure formation are suppressed, and it becomes virtually impossible for halos below a certain mass threshold to form. While the SP dynamics also implies a threshold for the formation of virialized objects, this is drastically raised when introducing FDM ICs. Indeed, this behavior is illustrated quite starkly by the fact that there are only 68 bound halos in the entire simulation volume at $z = 3$, whereas the SP simulation with CDM ICs contains thousands of halos in the same (quite moderate) volume at the same time.

Although the statistics are somewhat poor due to the small number of resulting

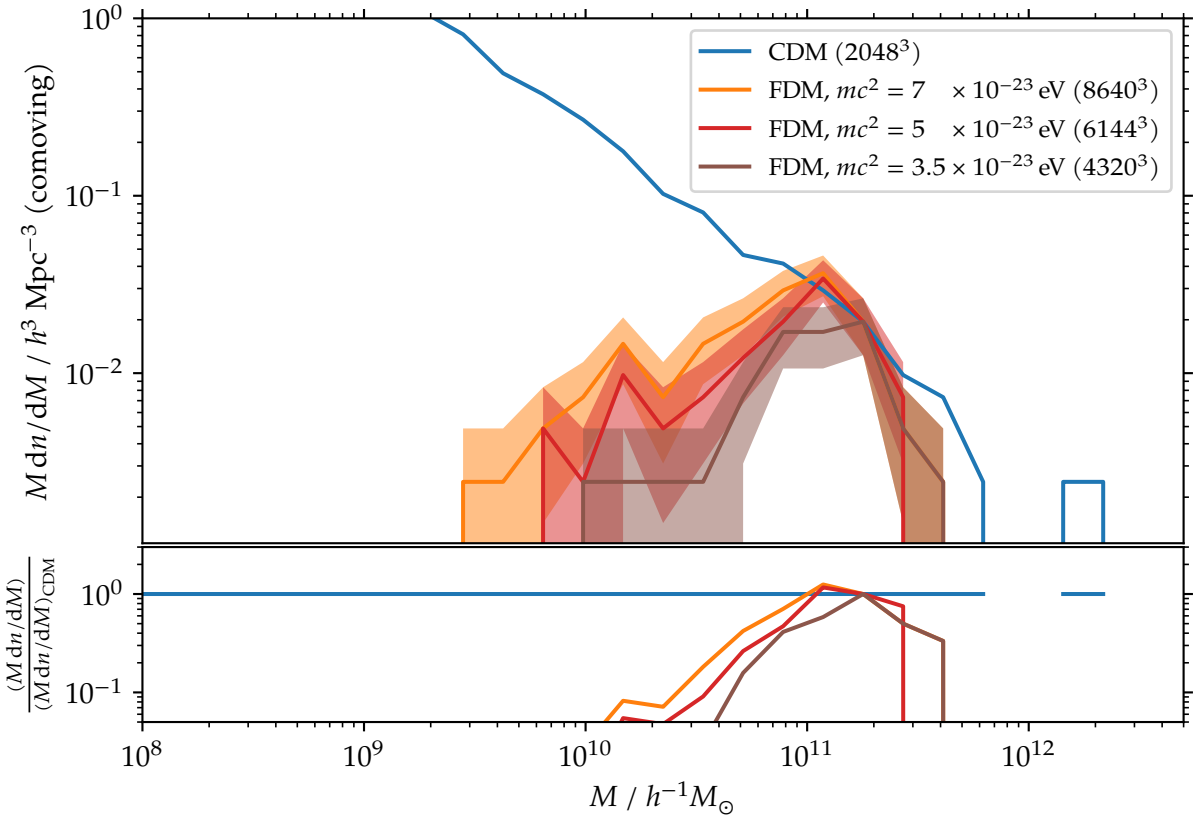

Figure 5.10: Halo mass functions in cosmological FDM simulations at $z = 3$ with different FDM mass values m and a box size of $L = 10\,h^{-1}\,\mathrm{Mpc}$, compared to a corresponding CDM simulation. The lower panel shows the ratios of the mass functions to the CDM result.

simulated halos – despite the efforts to simulate a larger volume than has been standard for SP simulations thus far – the HMF can be measured with reasonable accuracy and arrive at an important result: There appears to be quite good agreement between the HMF measured in the fully self-consistent FDM wave simulation and the fitting function of Schive et al. (2016), which was determined from N-body simulations with an FDM power spectrum. Consistent with my conclusion above, this confirms that the ICs are the primary factor in determining the FDM HMF, and that the use of N-body simulations with filtering techniques is in principle sufficient to obtain at least an approximate estimate of the HMF in an FDM cosmology.

The (semi-)analytic estimates obtained in Du, Behrens, and Niemeyer (2017) and Kulkarni and Ostriker (2022), on the other hand, differ from the result obtained using this technique: Both feature a steeper turnover and drop at low masses than

measured in simulations, and while the former results in a similar cut-off mass, it underestimates the number of halos for high masses and overestimates it for low masses near the cut-off, whereas the latter features a lower cut-off mass, which does not account for the presence of the lower-mass halos found. Thus, compared to methods based on the extended Press–Schechter formalism, the N-body approximation seems to be the most reliable "simplified" approach to the HMF so far, as it yields the best match to my measurements in full FDM simulations.

Figure 5.10 shows the dependence of the HMF on different values of the particle mass m (at $z = 3$). It behaves as expected, although simulations with lower m have even fewer halos and thus suffer more severely from low statistical counts.

Unfortunately, given the small number of halos present even at the final redshift $z = 3$, it is not possible to meaningfully study the evolution of the HMF with redshift in the simulation volume and to explicitly confirm that this conclusion also holds at other redshifts. Likewise, reliably characterizing the quantitative precision of fitting functions such as those of Schive et al. (2016) will have to await FDM simulations with much larger volumes and hence better statistics.

5.6 Fuzzy dark matter filaments

My findings concerning the stark departure of the nature of filaments in FDM compared to CDM in section 5.3.3, and in particular their different visual appearance of the outer parts of halos, compelled me to further investigate FDM filaments. Figure 5.2 provides an image of the cosmic web in FDM, which is quite unlike the more familiar structures expected for CDM, and is thus adding to this motivation. Previous results for WDM and smaller "proof-of-concept" simulations of FDM have already provided hints that filaments represent another, second type of significant object when it comes to baryonic processes such as star formation (e. g. Gao and Theuns 2007; Gao, Theuns, and Springel 2015; Mocz et al. 2020). Contrary to CDM, where filamentary structures fragment into (sub-)halos down to the smallest scales, the filaments in these models stay intact as smooth, overdense, large-scale dark matter structures.

I used the DisPerSE code (Sousbie 2011; Sousbie, Pichon, and Kawahara 2011) in order to attempt to identify the filaments in my FDM simulations. Unfortunately, this approach was met with technical difficulty due to computational limitations of the DisPerSE software. In particular, the code can only operate on shared-memory architectures, with no capabilities for today's distributed-memory parallel computing

infrastructure using the MPI standard. This means that it is inherently limited to the amount of memory present on a single node on a computing cluster. Gravely exacerbating the problem is the fact that the DisPerSE algorithms are rather memory-intensive. Even with the largest amount of memory available on the MPCDF's Raven system, which amounts to two terabytes, the simulated density fields were forced to be severely "down-sampled" in order to allow them to be processed by DisPerSE. As a note, previous work applying DisPerSE to find filaments in large cosmological simulations, such as Galárraga-Espinosa et al. (2020), was able to avoid this problem by applying the code on a tracer field only, such as the distribution of galaxies. However, this option is not available in this case, because there are no objects (such as halos) present in the dark matter filaments which could serve as a tracer field with appropriate sampling rate. Rather, filaments need to be identified directly in the smooth dark matter density field.

In order to ensure comparability between SP (FDM) and N-body (CDM) simulations, I constructed a CDM density field by binning the N-body particles to a Cartesian grid using CIC mass assignment, making the input to DisPerSE a uniform Cartesian grid of the density field in both cases. For FDM, the original 8640^3 grid was smoothed and scaled down by computing the means of neighboring cells. A cut value of $60\rho_\mathrm{m}$ was used, which is the minimum threshold difference between critical points in order for DisPerSE to keep them. This ensures that the identified structural complex is not swamped by tiny small-scale fluctuations in the density field.

In the end, the largest grids which that could be processed with DisPerSE were of size 288^3 for the largest FDM simulation, 260^3 for the N-body simulations (including its CDM counterpart), and only 216^3 for the largest SP simulation with CDM ICs. This of course drastically reduces the spatial resolution (to $34.7\,h^{-1}$ kpc, $38.5\,h^{-1}$ kpc, and $46.3\,h^{-1}$ kpc, respectively), and thus cannot be expected to yield either the full population of simulated filaments or their precise locations. However, it nevertheless results in a reasonable filament catalog which traces a good fraction of the cosmic web, as shown in fig. 5.11.

While it is of course possible to investigate quantities like the radial filament density profiles even with this degraded resolution, the results are not very insightful since the identified filament lines will not generally trace the true, fully-resolved filament centers, making the computed profiles in a significant fraction of the total radial extent unreliable due to the varying offsets from the center. In order to improve the measurement, the filaments output by DisPerSE were post-processed by evaluating each filament segment and gradually shrinking a cylindrical region around the

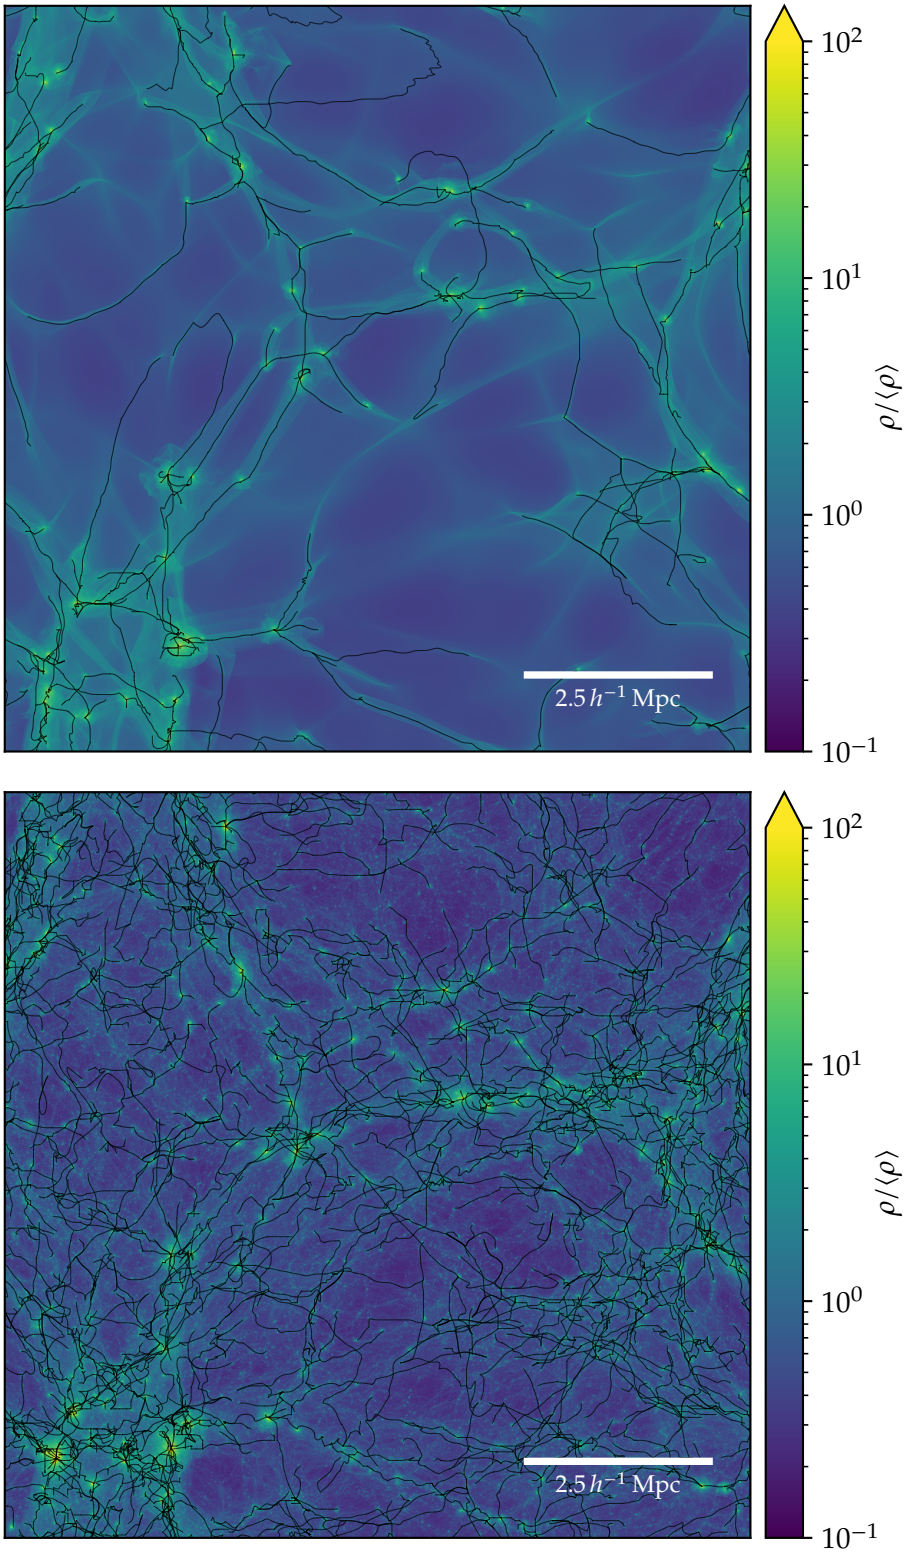

Figure 5.11: Projected dark matter density showing the filaments identified by DisPerSE (marked using black lines) in a cosmological box simulation of FDM at $z = 3$ with box size $L = 10\,h^{-1}$ Mpc, grid size $N^3 = 8640^3$, FDM mass $mc^2 = 7 \times 10^{-23}$ eV (top), and in a similar CDM simulation with 2048^3 particles (bottom). In order to enable DisPerSE to process the simulated density field, it had to be scaled down to a 288^3 (FDM) or 260^3 (CDM) grid.

segment towards the center of mass contained within the region at each step, similar to the "shrinking spheres" technique often used to determine the densest (central) point of a halo. In this case, the segments were only allowed to be offset in a direction perpendicular to themselves, in order to avoid several segments clustering around a single dense point and preserving the property that the segments trace the full length of the filament. While they do not remain continuously connected with this approach, the resulting error in the density profiles will be rather small as long as the originally identified filament line roughly traces the actual filament, since there will only be a slight inaccuracy in the measured distances of matter from the line center due to the tilt of the segments compared to the "true" continuous filament line.

Selected FDM and CDM filaments are shown in fig. 5.12, demonstrating the outcome of this procedure. Especially for FDM, it can clearly be seen that the original line which runs offset from the densest inner region of the filament is correctly shifted towards the center. For CDM, on the other hand, the definition of a filament center is more problematic, since the filaments are peppered with numerous small halos, which are locally the densest points. While the procedure also correctly shifts the segments to these dense points, the set of segments can become quite discontinuous in this case, and it is visually not clear that the halos trace what one would intuitively identify as the center of the filament as a whole. It may be the case that another tracer, such as a suitably smoothed halo distribution, is more appropriate to describe the shape of CDM filaments.

The final (cylindrically) radial density profiles, both as "stacked" profiles and for a sample of individual ones, can be found in fig. 5.13. It is evident that, for both FDM and CDM, filaments can reach rather high central densities. The main difference between the two is the slope and extent of the density profiles: The FDM filament density remains relatively flat and falls more slowly towards the outer regions, reaching much further outwards before dropping to the background matter density, while CDM filament profiles are relatively steep. Whereas CDM filaments on average decline around a radius of $500\,h^{-1}$ kpc, FDM ones often extend beyond $1\,h^{-1}$ Mpc. This difference in the shape of the density profiles is mostly determined by the choice of ICs, independently of the simulation method.

In the inner regions, CDM filaments seem to reach higher central densities, although this statement should be interpreted with caution. Firstly, as mentioned before, it should not be expected that the identified filaments represent the full population present in the simulation due to the coarse-graining which had to be performed

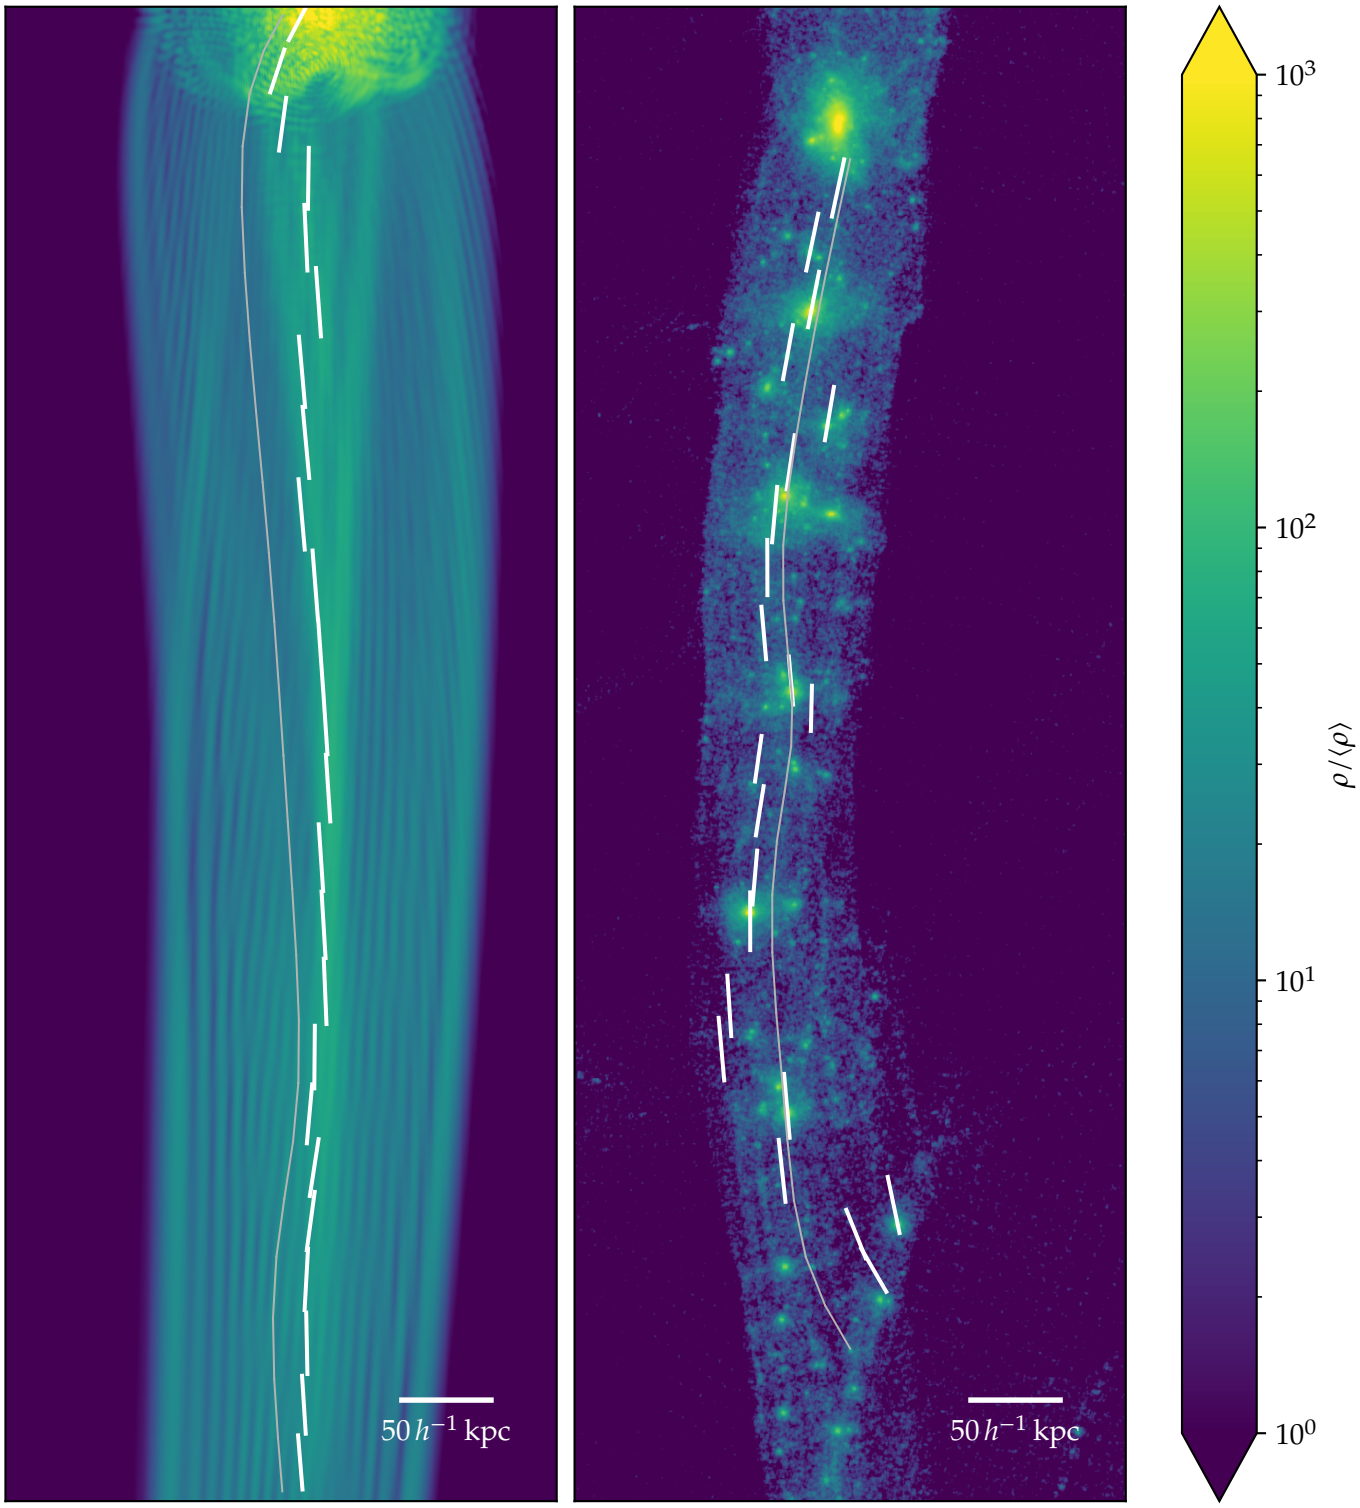

Figure 5.12: Projected dark matter density for an example FDM ($mc^2 = 7 \times 10^{-23}$ eV; left) and CDM (right) filament at $z = 3$. A thin gray line indicates the location as identified by DisPerSE on the coarse grid, while the thicker white line segments show how this initial DisPerSE output was shifted to more accurately capture the filament center for the purposes of computing radial density profiles.

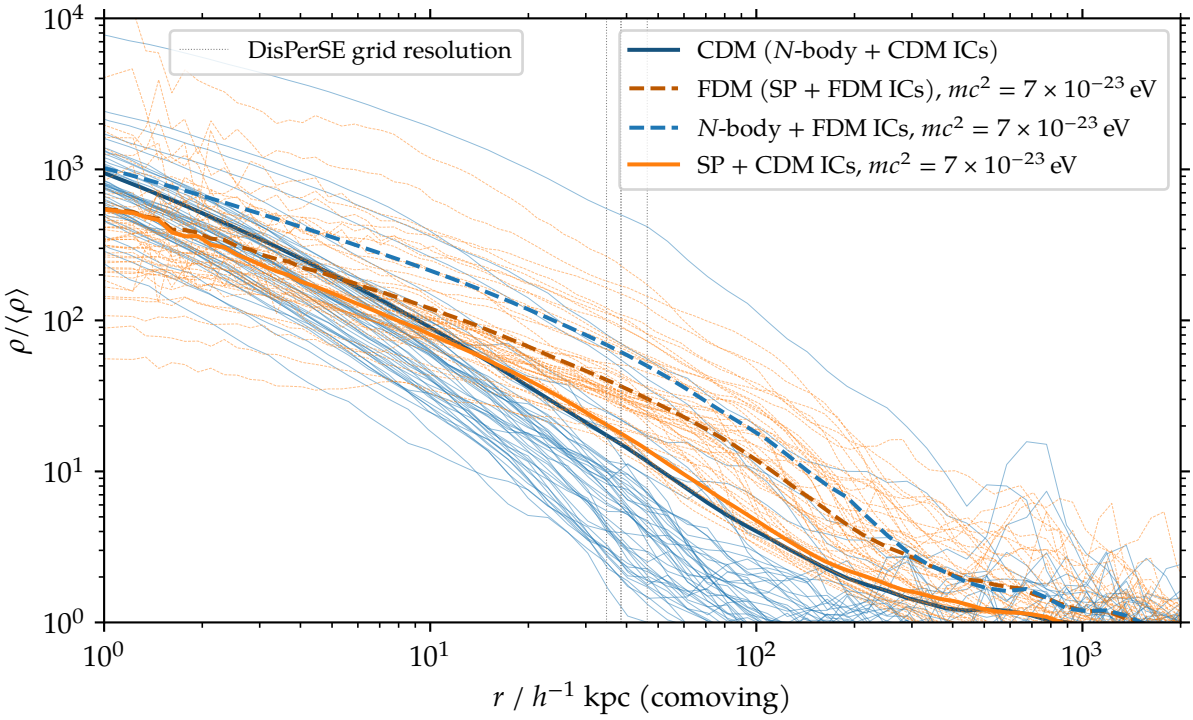

Figure 5.13: Filament profiles at $z = 3$ for the filaments identified by DisPerSE in $L = 10\,h^{-1}$ Mpc cosmological box simulations of CDM (2048^3 particles) and FDM ($N^3 = 8640^3$, $mc^2 = 7 \times 10^{-23}$ eV), scaled down to 260^3 and 288^3 grids, respectively, for processing with DisPerSE. The corresponding grid resolutions used with DisPerSE are marked by dashed lines. Filaments identified by DisPerSE were post-processed using the full simulation data, allowing to resolve the density profiles at much smaller distances than the resolution imposed by the DisPerSE grids. The thin lines show a random sample of individual filament profiles, while the thick lines are the mean ("stacked") density profiles of all filaments with a total length $> 400\,h^{-1}$ kpc. Stacked profiles for corresponding SP and N-body simulations with the other set of ICs are shown for comparison (the SP + CDM ICs case had to be scaled down to 216^3 for use with DisPerSE).

on the density fields. Secondly, the grid resolution for the FDM simulation fundamentally limits the scales on which the density profile can be measured, whereas N-body particles can gather on much smaller scales, so the innermost regions may be subject to resolution effects. This is emphasized by the fact that both N-body and SP simulations, respectively, exhibit the same central density in the stacked profiles independently of the ICs. Thirdly, the filament segments tend to center on halos present in the CDM filaments, which of course tend to have very high inner densities, meaning that the density profile measurement is in part "contaminated" by these halo profiles. This once again raises the question of how to define the filament center.

Finally, the scatter in FDM filament densities is much greater than for CDM, in particular featuring a population of filaments with inner densities of around $100\rho_{\mathrm{m}}$ or even below. Further investigation into different populations of filaments may reveal why these only seem to occur with FDM. Another interesting observation is that the inner FDM filament profiles are much noisier than the CDM ones. While this can be partially attributed to resolution, which leads to poor sampling rates in the small inner cylindrical shells, it could also be a hint of the wave patterns and fluctuations present with FDM.

Generally, it is clear that the procedure employed here is not the ideal way to measure dark matter filaments. However, in the absence of superior methods and tools, fig. 5.13 demonstrates that it can still provide adequate results even far below the nominal grid resolution. The main drawback of the approach is that the filament population, as identified by DisPerSE on the coarse density grids, is likely incomplete, leading to potential biases in quantities like the mean density profiles. Further, the sample is contaminated by pathological cases, such as spurious filaments which do not directly correspond to any structure in the original, fine density field. Unfortunately, the need for "correcting" the output in post-processing makes it difficult to distinguish legitimate cases which can be fixed from pathological ones. Enforcing a minimal filament length somewhat alleviates this, but is not a complete solution.

5.7 Fuzzy dark matter halo profiles from a self-consistent cosmology

The halo catalogs from my simulations in chapter 4 are also very useful to study a large sample of FDM halos, e. g. in the context of the core–halo mass relation (Chan, Ferreira, May, et al. 2022). Due to the use of a CDM initial matter power spectrum, the HMF is strongly enhanced compared to a "pure" self-consistent FDM model, even for low halo masses, as shown in section 5.5. This significantly boosts the statistics by increasing the number of simulated halos by several orders of magnitude, which can alleviate the issue of limited simulation volumes available in FDM simulations to some extent.

However, there is the lingering question of whether the cosmological context, which differs from that of a "real" FDM cosmology due to the use of a CDM initial matter power spectrum, has an impact on the halo properties, and thus whether this approach is ultimately valid for studying halo profiles. If this were the case,

then the halo population may not be representative of a true sample of FDM halos with self-consistent ICs, making the approach less suitable for studying FDM halo properties. With the addition of FDM IC simulations of sufficient size, it is now possible to explicitly check for differences in halo properties for the different ICs.

Figure 5.14 shows the radial density profiles of all halos in the largest FDM simulation, highlighting a small number of individual profiles for visual clarity. The profiles present the familiar solitonic core (fit e. g. by the analytic approximation from Schive, Chiueh, and Broadhurst 2014),

$$\rho_{\mathrm{soliton}}(r) = \frac{1.9 \times 10^9 \, a^{-1} \left(\frac{mc^2}{10^{-23}\,\mathrm{eV}}\right)^{-2} \left(\frac{r_\mathrm{s}}{\mathrm{kpc}}\right)^{-4}}{\left(1 + 0.091 \left(\frac{r}{r_\mathrm{s}}\right)^2\right)^8} M_\odot \,\mathrm{kpc}^{-3}, \tag{5.13}$$

in the center and NFW density profile (Navarro, Frenk, and White 1996; Navarro, Frenk, and White 1997),

$$\rho_{\mathrm{NFW}}(r) = \frac{\rho_0}{\frac{r}{R_\mathrm{s}}\left(1 + \frac{r}{R_\mathrm{s}}\right)^2}, \tag{5.14}$$

in the outskirts.[7] In order to compare to the simulation with CDM ICs, which contains many more halos, an equal number of halos whose virial mass most closely match those in the FDM IC simulation was selected. Figure 5.14 also shows "stacked" profiles[8] in different mass bins defined by these halos in a separate panel, with both sets of stacked profiles matching excellently, although there seems to be a slight trend towards higher central densities for the CDM ICs case. There is one exception in the lowest-mass bin, where the central soliton differs quite drastically. However, this bin only contains a single halo, so this only highlights that halos with similar masses can have different properties. In this case in particular, it seems that the halo in the FDM simulation has its mass somewhat inflated by excess overdensity in its outskirts due to its presence in a filament (see last halo in fig. 5.8). Thus, while there might certainly be a statistical impact in the form that a larger sample of halos is more likely to contain outliers, e. g. "unusual" halos with a "rare" merger history, both cases agree very well at the level of the averaged density profiles, and the results do not show evidence for large systematic differences. There are slight hints for higher central densities in the "true" FDM simulation, which could be investigated in future work.

[7] All given quantities are comoving.

[8] The procedure of "stacking" refers to taking the mean density across all halos for each radial bin.

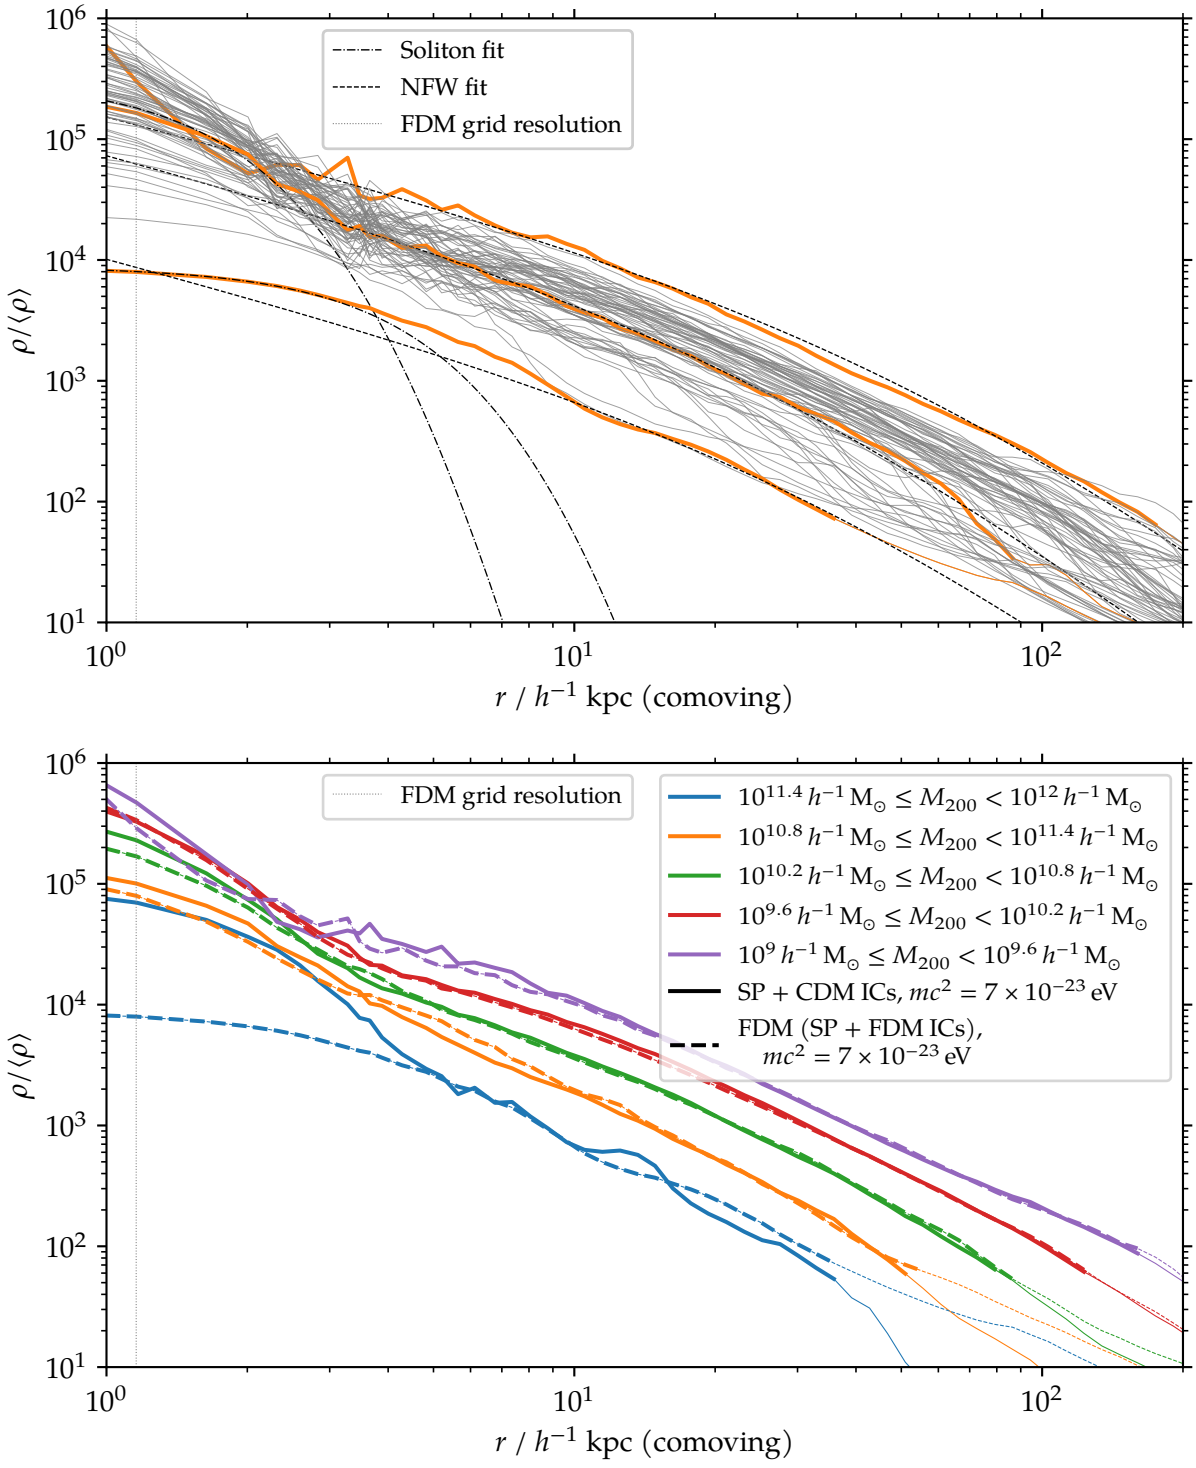

Figure 5.14: *Top:* Halo profiles for all 68 identified halos (gray lines) in an $L = 10\,h^{-1}$ Mpc cosmological box simulation of FDM ($N^3 = 8640^3$, $mc^2 = 7 \times 10^{-23}$ eV, FDM ICs), with the minimum-, maximum-, and an intermediate-mass halo highlighted. *Bottom:* "Stacked" profiles in several mass bins of all halos for the FDM ICs case and the subset of halos which match these most closely in mass for the CDM ICs case. Thin dot-dashed and dashed lines show soliton (5.14) and NFW fits (5.14) to the regions of the stacked profile within the soliton radius and virial radius R_{200}, which are indicated with thick lines.

5.8 Summary and Conclusions

The simulations presented in this chapter constitute the largest numerical simulations of cosmic structure formation with FDM that account for its full wave nature via the SP equations, using a pseudo-spectral method. By approaching simulation volumes more representative of cosmic large-scale structure, I was able to follow the combined linear large-scale, and non-linear small-scale evolution of wave FDM in a fully self-consistent manner. The pseudo-spectral method can be considered the most accurate approach to solve the SP system of equations numerically (although usually, the issue is not about the accuracy of the numerical approach to SP equations, but rather whether an attempt is made to solve these equations at all, or whether the cost is shirked by choosing one of the approximate options). It encompasses all aspects of the wave amplitude's temporal evolution, and in particular accounts for the oscillatory, order-unity fluctuations of the local density due to the wave effects in the axion-like dynamics.

I stress again that, unfortunately, the numerical resolution requirements to faithfully follow the FDM dynamics are much more stringent than for the familiar N-body techniques applicable in the CDM case. Even large-scale modes require a very fine mesh due to the velocity criterion, eq. (5.4), since the spatial oscillations of the phase factors are otherwise not resolved, resulting in a halted evolution and eventually just numerical noise. Nevertheless, the time step criterion, eq. (5.2), with its quadratic (rather than linear) dependence on the spatial resolution, enforces very small time steps when the mesh is made fine enough to resolve the de Broglie wavelength λ_{dB}, even when the coarsest possible mesh resolution is selected. These requirements continue to severely limit the regimes available to full numerical simulations of FDM, and in particular preclude large cosmological volumes to be computed at low resolution in a way similar to the standard practice in CDM.

The large, fully self-consistent FDM simulations performed here have allowed me to gain new insights into the evolution of the non-linear power spectrum in these cosmologies, especially in comparison to CDM and the simplified approach of "N-body with FDM ICs". I could also make the first direct, self-consistent measurements of the halo mass function in such models. In the process, I discovered that the filamentary structure of the cosmic web behaves more similarly to WDM than to CDM, and characterized the properties of these filaments in comparison to CDM.

My main findings can be summarized as follows:

- In this chapter, I could determine the fully non-linear matter power spectrum of

self-consistently evolved FDM cosmologies over an exceptionally broad range of scales, and I have also computed matching CDM N-body simulations with the same initial conditions. On large scales, I confirmed the consistency between all simulations – a manifestation of the Schrödinger–Poisson–Vlasov–Poisson correspondence –, verifying the expectation that FDM behaves like CDM on these scales. On small scales, the "quantum pressure" in FDM suppresses structure formation compared to CDM. Self-consistently accounting for FDM yields initial conditions with substantially reduced power on small scales, however, producing yet bigger suppression effects for low-mass halos, in a way strongly reminiscent of WDM models.

- I was able to draw a four-way comparison of the evolution of the matter power spectrum for the different choices of simulated dynamics and ICs. The findings demonstrate that the differences between each option are time-dependent, with large variations in their relative behaviors. Especially at early times, the ICs are the determining factor, dominating the shape of the power spectrum at small scales due to the cut-off from at the FDM Jeans scale. By $z = 15$, bottom-up structure formation is already in full swing for CDM, with non-linear excess power accumulating on small scales. SP with CDM ICs starts to reach the linear theory power spectrum again at this point, whereas simulations with FDM ICs are still significantly lagging behind on scales smaller than the initial cut-off. Finally, for smaller redshifts, the shapes of all power spectra start to become rather similar, and by $z = 3$, the impact of going from CDM to FDM ICs and from N-body to SP dynamics is roughly similar in size. Of course, only SP simulations exhibit a small-scale excess "bump" (exceeding even the CDM power) due to wave interference.

- I have developed a technique which allows me to reliably identify halos in self-consistent FDM SP simulations, including the proper "WDM-like" ICs with suppressed small-scale power. This proved non-trivial because, similar to WDM, standard halo finding algorithms like FoF are not able to handle on their own the presence of extended, smoothly overdense filaments. By employing an increased overdensity threshold combined with a binding criterion and a filter on the identified objects' virial ratios, I was able to identify the subset of gravitationally collapsed, virialized objects.

- I could, for the first time, measure the self-consistent halo mass function directly from pseudo-spectral FDM simulations. By comparing SP simulations with

FDM and CDM ICs, I could confirm that the FDM ICs have a larger impact on the HMF than the late-time non-linear SP dynamics, although the suppression caused by the latter is still significant. For the "true" FDM HMF, I found broad agreement with the approximate scheme of obtaining the HMF from *N*-body simulations with an FDM initial power spectrum (as performed in this chapter and previously in Schive et al. 2016). In contrast, predictions based on the extended Press–Schechter formalism (Du, Behrens, and Niemeyer 2017; Kulkarni and Ostriker 2022) do not fit the HMF results so well.

- I have investigated the differences between cosmic filaments for FDM and CDM. Like WDM, FDM filaments do not fragment into (sub-)halos like those of CDM, and instead feature the characteristic wave fluctuation patterns. I found significant differences in the shapes of filamentary density profiles, with FDM filaments being much more extended and featuring a more slowly-declining density than CDM ones. These differences were mostly determined by the choice of the initial power spectrum. The scatter in identified filaments was also larger for FDM, which features a population with very high densities (even comparable to lower-mass halos). Such objects have been found to be the sites of first star formation with FDM (Mocz et al. 2020).

- One advantage of the "SP with CDM ICs" simulations which I have performed previously is the much larger population of simulated halos, which can be used to infer FDM halo properties using a large sample. However, it has not been a priori clear whether all the halos obtained in such a simulation are representative of a fully self-consistent FDM cosmology that is based on the correct primordial power spectrum, yielding only a much suppressed halo abundance. In this chapter, I was able to show that the properties of the simulated halos match quite closely, validating the use of the larger halo catalogs for the purposes of studying FDM halos. This is reminiscent of *N*-body simulations of CDM and WDM cosmologies, which show very similar halo density profiles apart from a slightly reduced concentration of WDM halos, consistent with their later formation time (Avila-Reese et al. 2001).

Cosmological simulations of FDM are numerically much more challenging than CDM, but my work shows that at least small volumes can be studied with decent spatial resolution. Including baryons explicitly would clearly be a very worthwhile and interesting next step to arrive at a more reliable assessment of whether these cosmologies are still viable, and how tight some of the constraints placed on the

particle mass really are. In this regard, the high computational cost of FDM simulations can actually be viewed as an encouragement. Unlike for "cheap" CDM simulations, adding hydrodynamics and a modeling of galaxy formation physics should be comparatively easy to do – at least it is not expected to dominate the computational expense, and in this sense appears quite feasible. This will be attempted in forthcoming work.

Fuzzy dark matter simulations with baryons and the Lyman-α forest

6

6.1 The Lyman-α forest

The so-called Lyman-α forest is an important observational probe of cosmological large-scale structure and the intergalactic medium (IGM) (Weinberg et al. 2003). Figure 6.1 illustrates how it arises: A distant object (typically a quasar, since quasars are strong emitters in the relevant part of the electromagnetic spectrum, i. e. ultraviolet light) emits radiation with a relatively smooth spectrum in the ultraviolet range below the wavelength of the Lyman-α line at around 1216 Å (i. e. it emits more energetic radiation). The Lyman-α line is the atomic transition line of hydrogen between the ground state ($n = 1$) and the first excited state ($n = 2$). As the radiation emitted by the source (e. g. quasar) propagates, it experiences redshift due to the cosmic expansion. On the way, it can encounter clouds of neutral hydrogen, which are present in the IGM. When ultraviolet light with the wavelength of the Lyman-α line encounters neutral hydrogen, it will excite the hydrogen atoms and be absorbed (the excited hydrogen will eventually fall back down to the ground state and re-emit the Lyman-α photon; however, it will do so in a random direction, and thus almost certainly not reach the same observer as the rest of the radiation which was not absorbed). Thus, whichever part of the originally emitted spectrum which has been redshifted to the wavelength of the Lyman-α line when it reaches neutral hydrogen will be absorbed, corresponding to an absorption line in the spectrum measured by an observer of the quasar.

As the emitted light travels, more and more energetic (lower-wavelength) parts of the emitted spectrum are redshifted to the Lyman-α wavelength. Whenever a neutral hydrogen cloud is encountered, the corresponding absorption line is imprinted in a different part of the spectrum, since the light is continuously redshifted as it

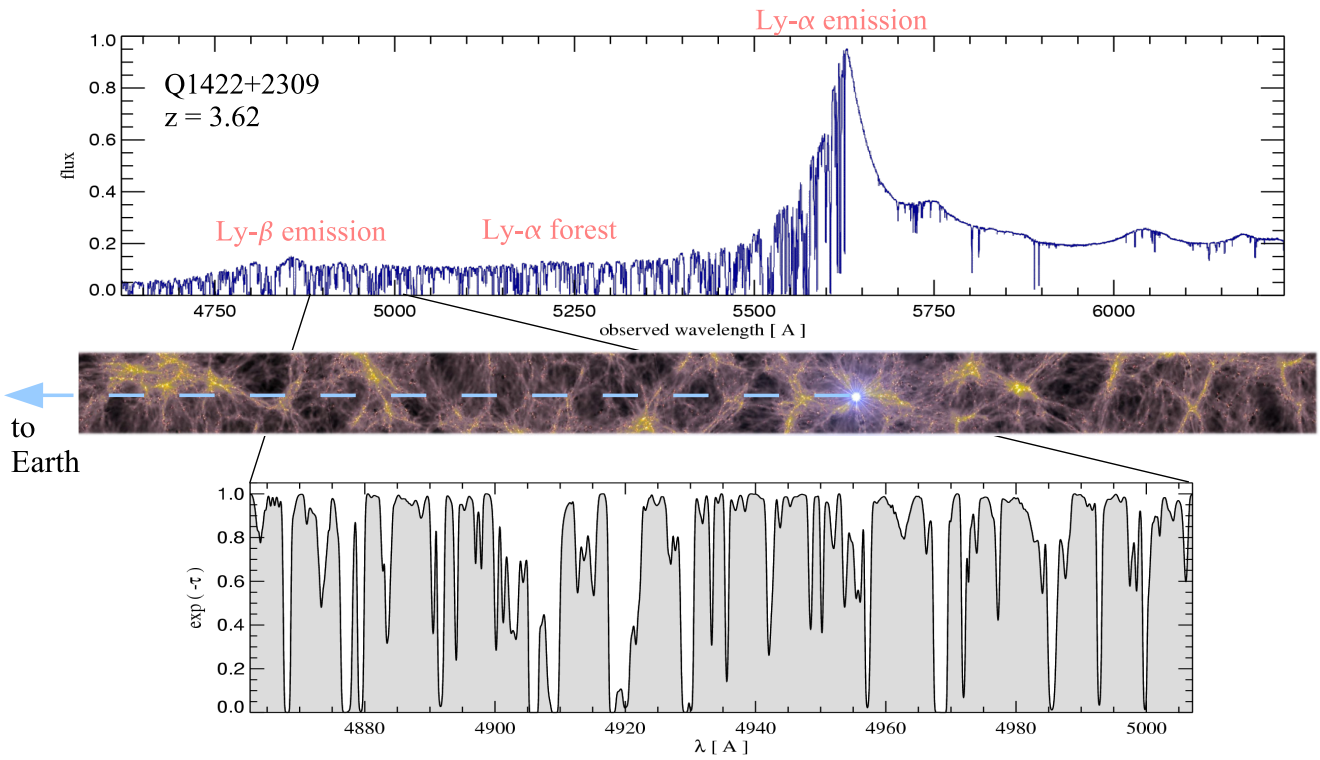

Figure 6.1: Illustration of how the observed Lyman-α forest arises. Taken from Springel, Frenk, and White (2006).

propagates. When it eventually reaches the observer, the result is a "forest" of many absorption lines as shown in fig. 6.1, each corresponding to a different absorber along the line of sight. In this way, the Lyman-α forest represents a one-dimensional map of the neutral hydrogen column density. This is an extremely valuable probe, because it is a direct measure of the distribution of hydrogen gas and thus, its large-scale structure. The statistics of this distribution can be used to check the predictions of cosmological models of structure formation.

6.2 Simulation setup

Building on the work developed in the previous chapters, the simulation volumes and redshifts reached in the large FDM simulations developed here are highly suitable to study the Lyman-α forest. In fact, only the development of these large-scale simulations makes it feasible to attempt this for a full wave treatment of FDM in the first place, since simulation volumes of $\approx 1\,h^{-1}\,\mathrm{Mpc}$ are simply too small to

meaningfully measure it. For these simulations, baryons from AREPO's moving mesh are now gravitationally coupled to AxiREPO's FDM grid. In this configuration, AREPO essentially operates analogously in PM-only mode, since there is no tree gravity for FDM, and the gas and dark matter interact via the computed PM potential and forces. This approach does not reduce the accuracy of the simulation, since the dark matter is limited to the Cartesian grid resolution in any case, and this resolution generally has to be set so high that the PM force for the gas is still resolved on quite small scales. In fact, contrary to the usual experience with simulations of CDM and baryons, the computational cost of the gas evolution and physics is essentially negligible compared to the FDM computations.

The simulations presented here follow Springel et al. (2018) for the cosmological parameters, which correspond to Planck Collaboration: Ade et al. (2016), with $\Omega_\mathrm{m} = 0.2603$, $\Omega_\mathrm{b} = 0.0486$, $\Omega_\Lambda = 0.6911$, $H_0 = 67.74\,\mathrm{km\,s^{-1}\,Mpc^{-1}}$ ($h = 0.6774$), $\sigma_8 = 0.8159$, and $n_\mathrm{s} = 0.9667$. The ICs were generated using an initial power spectrum computed with axionCAMB (both FDM and CDM) in an $L = 10\,h^{-1}\,\mathrm{Mpc}$ comoving box at $z = 63$, with the perturbed density field and particle distribution produced using the Zel'dovich approximation, as explained in chapters 4 and 5, and evolved until $z = 3$. For CDM, the baryons were included via the standard functionality of N-GenIC, which splits the initially created particle load into two sets of dark matter and gas particles, assigning mass values and shifting them according to the respective density parameters in order to maintain a valid density field. For FDM, the baryons are simply produced in the standard fashion as particles by N-GenIC – since dark matter is present in the form of a grid, it is not necessary to shift the particles, since there are no other particles which would cause collisions in the positions. Two simulations were performed: An FDM one with an 8640^3 Cartesian mesh for AxiREPO's SP solver and 256^3 gas cells, and a "standard" CDM one with 256^3 dark matter particles and 256^3 gas cells.

For simplicity, in these first simulations, only AREPO's basic built-in radiative cooling and star formation prescriptions were enabled for the baryonic physics (Weinberger, Springel, and Pakmor 2020). In addition, while AREPO's standard star formation model follows Springel and Hernquist (2003), an additional simplification was made since there were no feedback or galaxy formation physics included in these simulations: Instead of forming stars with a certain rate below the star formation threshold, gas will be converted to star particles instantaneously once it reaches the aforementioned threshold.[1] Since the Lyman-α forest is sensitive to the volume of

[1] This is known as the `QUICK_LYALPHA` configuration option in AREPO.

absorbing gas, whereas star formation occurs in small, collapsed regions, this choice should not significantly affect the measurements.

6.3 First results

Unfortunately, the first simulation of the FDM with baryons using AxiREPO reached $z = 3$ only a few weeks before the submission of this dissertation. Thus, there was not enough time to analyze the simulation data thoroughly, so that only some selected first results can be presented here. In particular, the main question of whether the characteristic wave fluctuations of FDM manifest themselves in the Lyman-α forest in some unique way which is distinguishable from CDM, which would represent a "smoking gun" for the astrophysical detection of FDM, still requires more detailed investigation.

6.3.1 The impact of baryons on matter clustering

To begin, fig. 6.2 displays the projected FDM and gas densities of the simulation box at $z = 3$. It is clear that the gas traces the dark matter distribution; however, it can also be seen that the high-density gas structures are more compact than the dark matter structures they reside in. This demonstrates that the gas can condense and collapse due to baryonic physics processes, whereas the dark matter cannot achieve a similarly compact state as easily. FDM in particular resists collapse below the de Broglie scale, and this contrast to baryons is clearly exemplified by fig. 6.2.

However, even beyond the distinct dynamics considered separately, it also shows the impact of the presence of baryons on the FDM structures: Compared to FDM-only simulations, such as those in chapter 5, where corresponding projections are shown in fig. 5.2,[2] the dark matter structures in fig. 6.2 are also slightly more compact than their dark matter-only counterparts. Although some small deviations are to be expected because the cosmological parameters are slightly different, this mostly indicates that the additional gravitational potential due to the compact baryonic structures is able to have an impact on the FDM structures as well.

Moving once again to the matter power spectrum, fig. 6.3 compares the evolution of the matter power spectrum in the FDM simulations with and without baryons, with a high-resolution CDM-only simulation shown again for reference. For the

[2] Due to the use of the same random seed in the IC for these simulations, it is easy to draw a direct comparison.

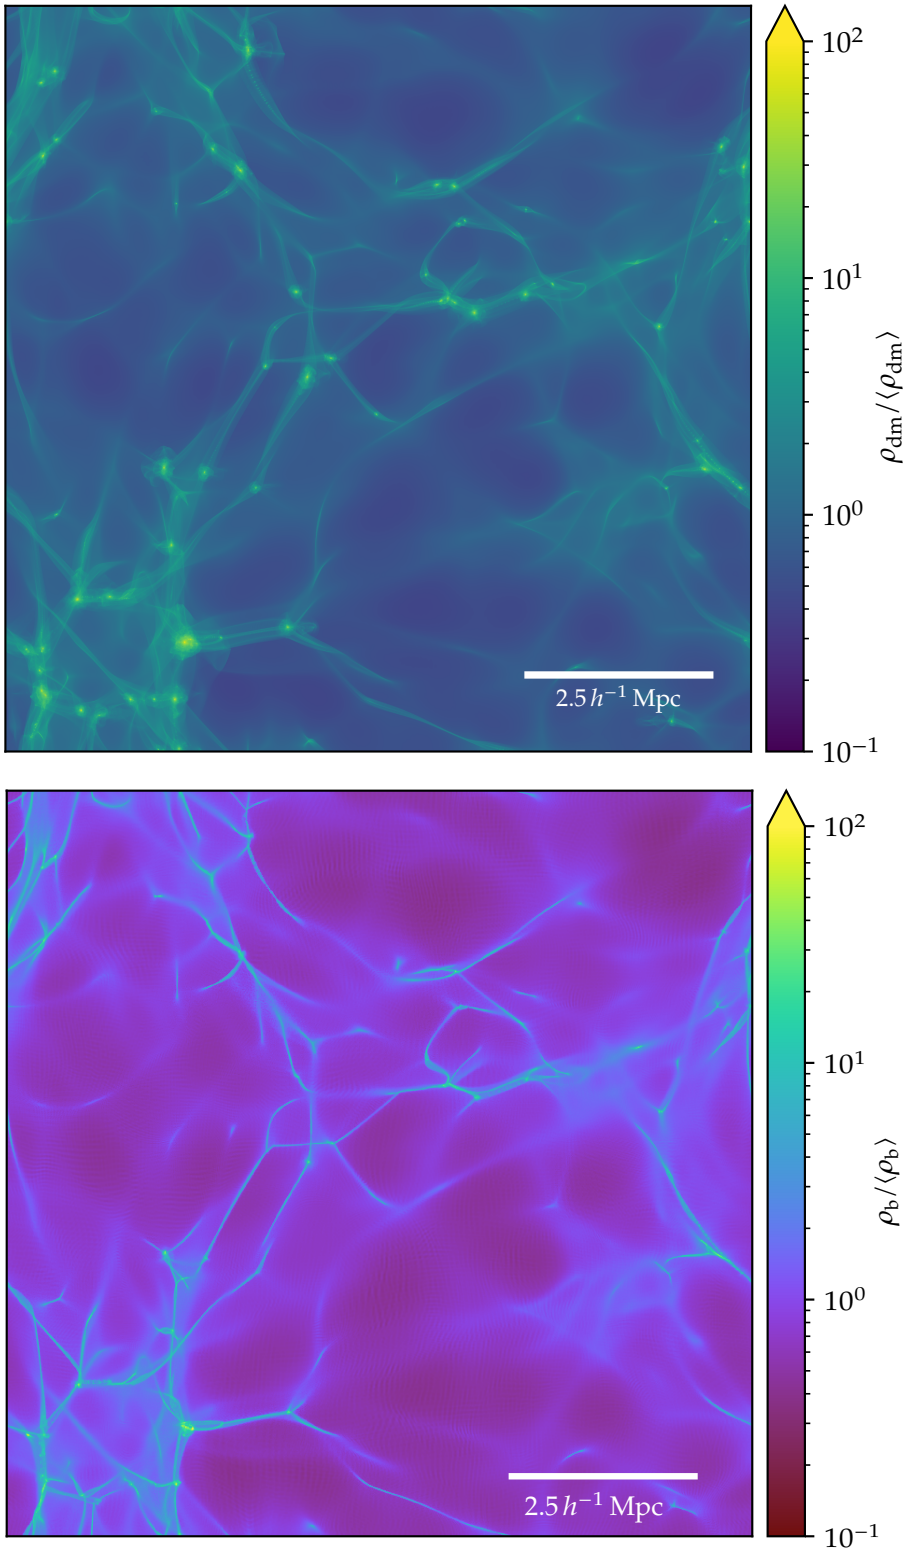

Figure 6.2: Projected dark matter density (left) and baryon gas density (right) in a $10\,h^{-1}$ Mpc box simulation of FDM with $mc^2 = 7 \times 10^{-23}$ eV at $z = 3$. The total numbers of grid and gas cells in the simulation are 8640^3 and 256^3, respectively.

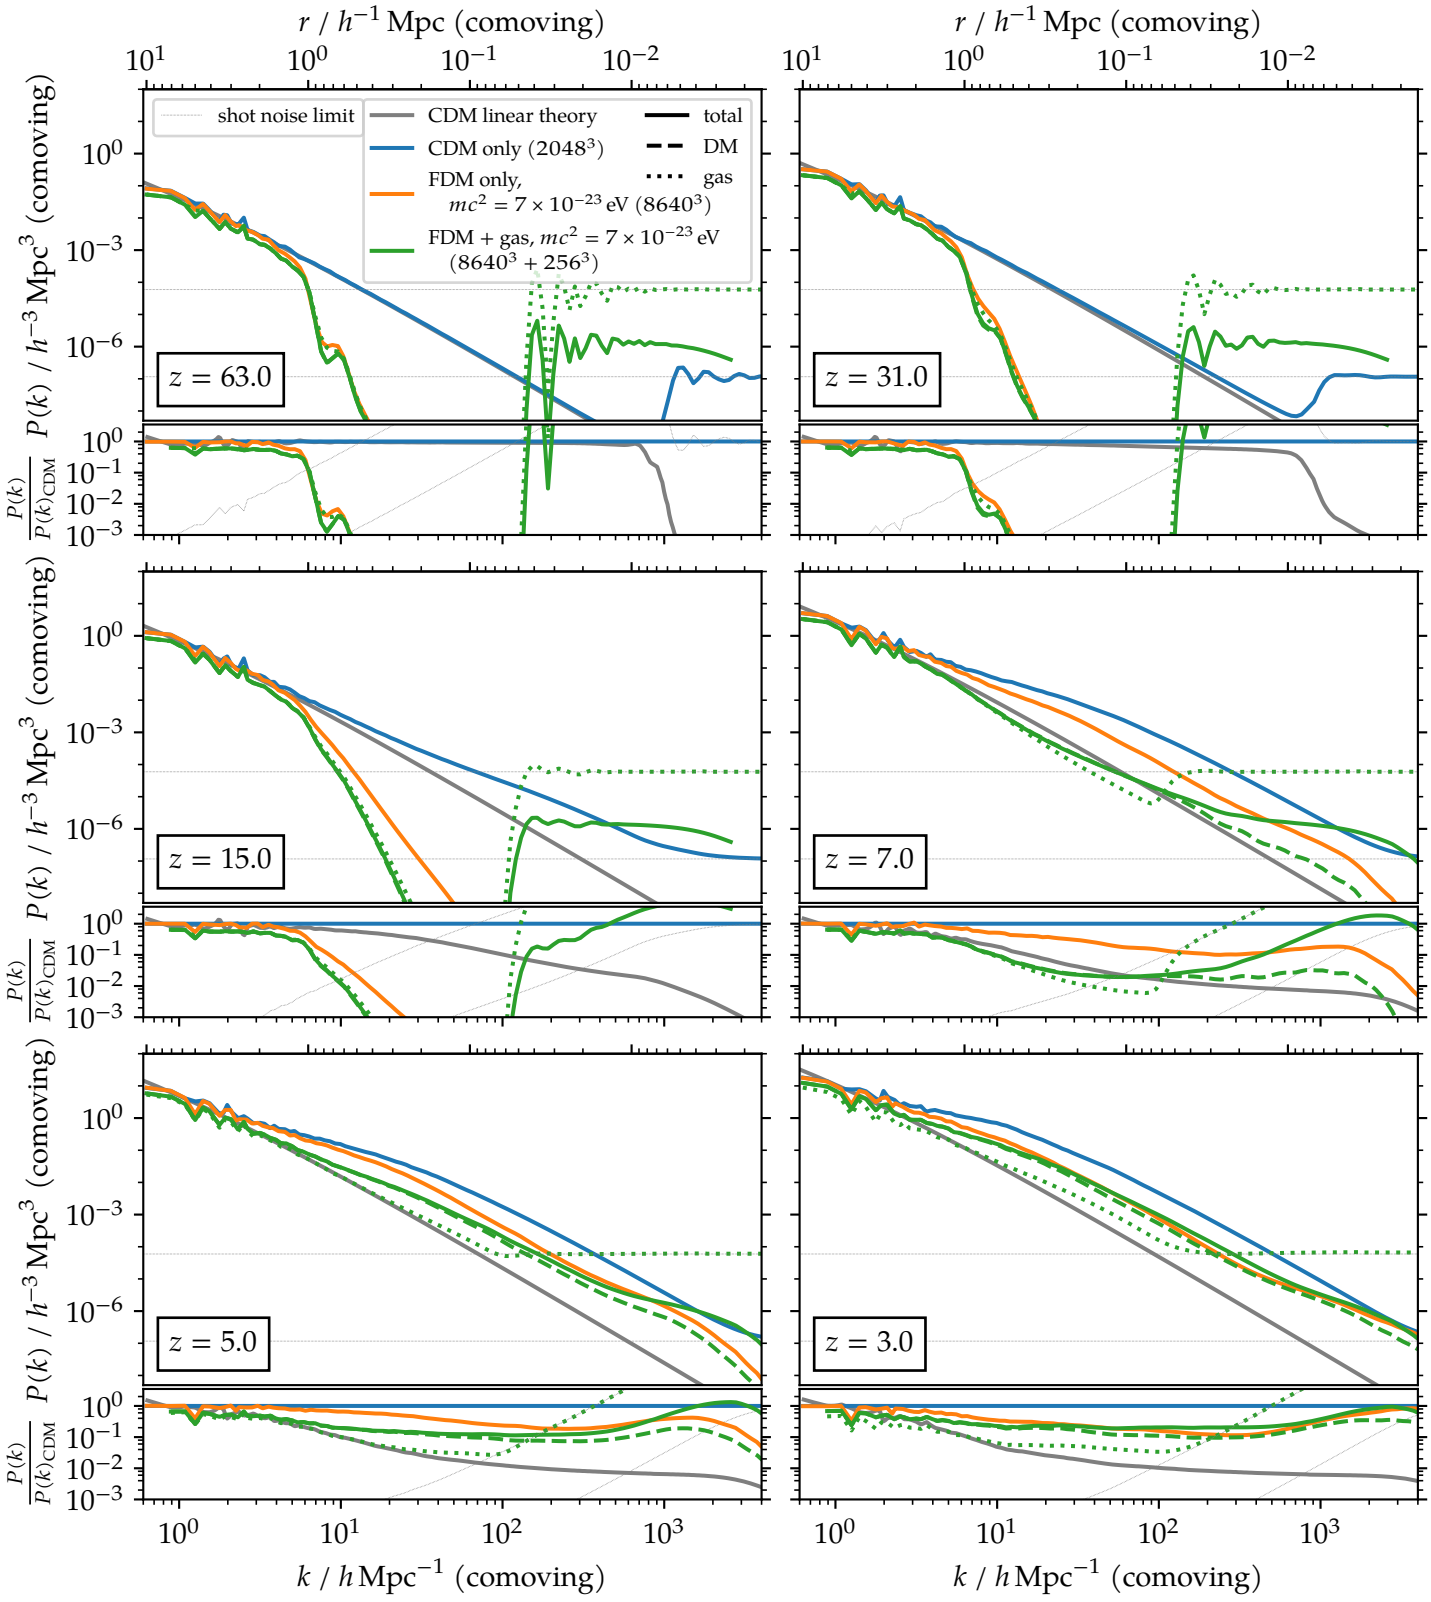

Figure 6.3: Matter power spectrum in a $10\,h^{-1}$ Mpc box simulation of FDM with baryons at $z = 3$ compared to dark matter-only simulations of both FDM and CDM. The power spectrum evolved using linear perturbation theory (scale-independent growth factor) is shown for comparison. The lower panels show the ratio of the power spectra to the CDM result. Faint dotted lines show the shot noise limits of the gas cells (upper) and the N-body CDM simulation (lower); the power spectrum of the corresponding component cannot be measured accurately once it reaches this limit.

FDM simulation with baryons, the separate power spectra of the different matter components (FDM and gas) are displayed in addition to the total matter power spectrum. It should be noted that, although the resolution of the gas cells is relatively high for a simulation of this box size (with 256^3 cells), it is completely dwarfed when compared to the number of resolution elements in the dark matter component (8640^3 for FDM and 2048^3 for CDM). Since the number of particles sets the limit below which so-called shot noise dominates the power spectrum measurement (treating AREPO's Voronoi mesh cells as particles), this limit will be reached for the gas component much sooner than the dark matter, preventing it from being probed down to similarly small scales. Explicitly, the gas power spectrum can only be measured up to around $k \lesssim 10^2\, h\,\mathrm{Mpc}^{-1}$. While the shot noise floor can be subtracted from the power spectrum in an attempt to obtain an estimate on smaller scales, this estimate should not be considered reliable.

Figure 6.3 shows that the inclusion of baryonic gas has an impact on the evolution of the power spectrum compared to the FDM-only case. (Note again that there is a small difference, especially with respect to a global offset in the normalization of the power spectrum, due to the different choice of cosmological parameters.) Interestingly, while the early (≤ 31) and final ($z = 3$) FDM power spectra are very similar, the intermediate evolution shows some discrepancies. Between $z = 15$ and $z = 5$, the simulation with baryons has fallen behind the FDM-only simulation in the growth of power on small scales below the initial FDM Jeans cutoff scale ($k \geq 10^1\, h\,\mathrm{Mpc}^{-1}$). However, by $z = 3$, it has caught up again, such that there is a temporary delay in the evolution of the power spectrum on those scales when including baryons.

Further, there is an interesting development on small scales at $z = 3$. Although it is somewhat difficult to say because the power spectrum quickly approaches the shot noise limit of the gas, it appears that the combined FDM and gas power spectrum exceeds the FDM-only one for $k \geq 20\, h\,\mathrm{Mpc}^{-1}$. Since the former has a lower overall normalization of the power spectrum due to the cosmological parameters, this pulling ahead of the simulation with baryons is indeed significant. It appears to be a quantitative confirmation of the effect observed earlier in fig. 6.2, where baryons are more condensed on small scales than the FDM, and in turn have a compactifying impact on the dark matter.

6.3.2 Fuzzy dark matter and the Lyman-α forest

I will now turn to the main purpose with which the baryonic FDM simulation was designed: The Lyman-α forest. In order to analyze and quantify what impact the

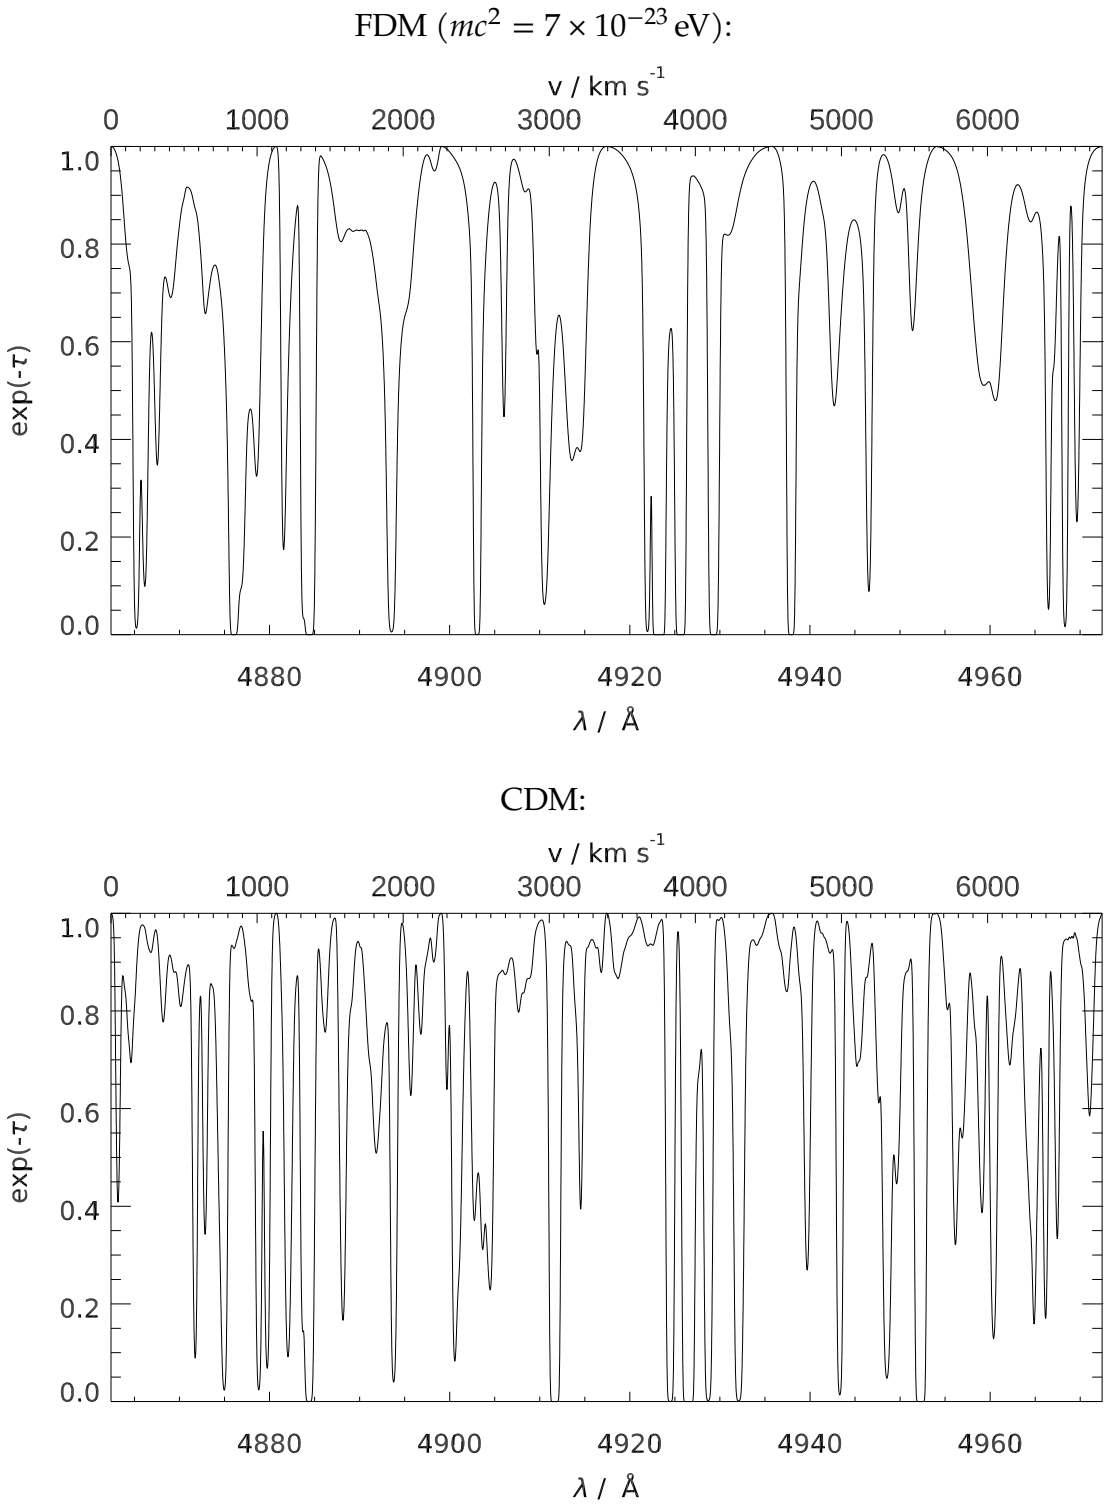

Figure 6.4: Transmission along a line of sight through the gas in a $10\,h^{-1}$ Mpc box simulation of FDM (top) and CDM (bottom) with baryons at $z = 3$.

distribution of gas in the simulation would have on the observed Lyman-α forest, the simulation volume is probed by many randomly chosen lines of sight. Along these lines of sight, synthetic absorption spectra are constructed according to the procedure outlined in Bolton and Haehnelt (2007). This results in the fraction of transmitted flux $T = \exp(-\tau)$, where τ is called the optical depth, from the source to the observer as a function of wavelength λ.

Figure 6.4 shows examples of such absorption spectra from the companion simulations of both FDM and CDM with baryons. It is clear that in the case of FDM with the chosen mass parameter, there are much fewer absorption troughs than for CDM, and those that are there also often appear broader than the CDM ones. This is a consequence of the reduced amount of small-scale structure present with FDM: Since the gas traces the dark matter, a reduction in small-scale structure implies that there are fewer opportunities for neutral gas to condense into (relatively) dense clouds (since there must be an appreciable amount of absorbing gas present along the line of sight to have a noticeable impact on the transmitted flux). A possible hint at the wave fluctuations in the FDM density field could hide, for example, around the first dip on the left in the FDM transmission in fig. 6.4. As the flux rises again towards higher wavelengths, there are what appear to be oscillations along the curve. It will be of interest to study whether the baryons follow the FDM density oscillations in some cases, and whether this can be identified in the Lyman-α forest through patterns like these.

The synthetic absorption spectra can be used to estimate the flux power spectrum, which is shown in fig. 6.5 for the simulations with baryons presented in this chapter with some observational data for context (McDonald et al. 2000). There is very good agreement on large scales, but the FDM starts to drop far below the CDM case very quickly towards smaller scales. However, on the smallest scales, FDM starts to slightly close the gap to again, meaning that the largest discrepancy is on intermediate scales. This effect could be related to the similar phenomenon in the matter power spectrum, where FDM wave fluctuations cause an enhancement of power on small scales.

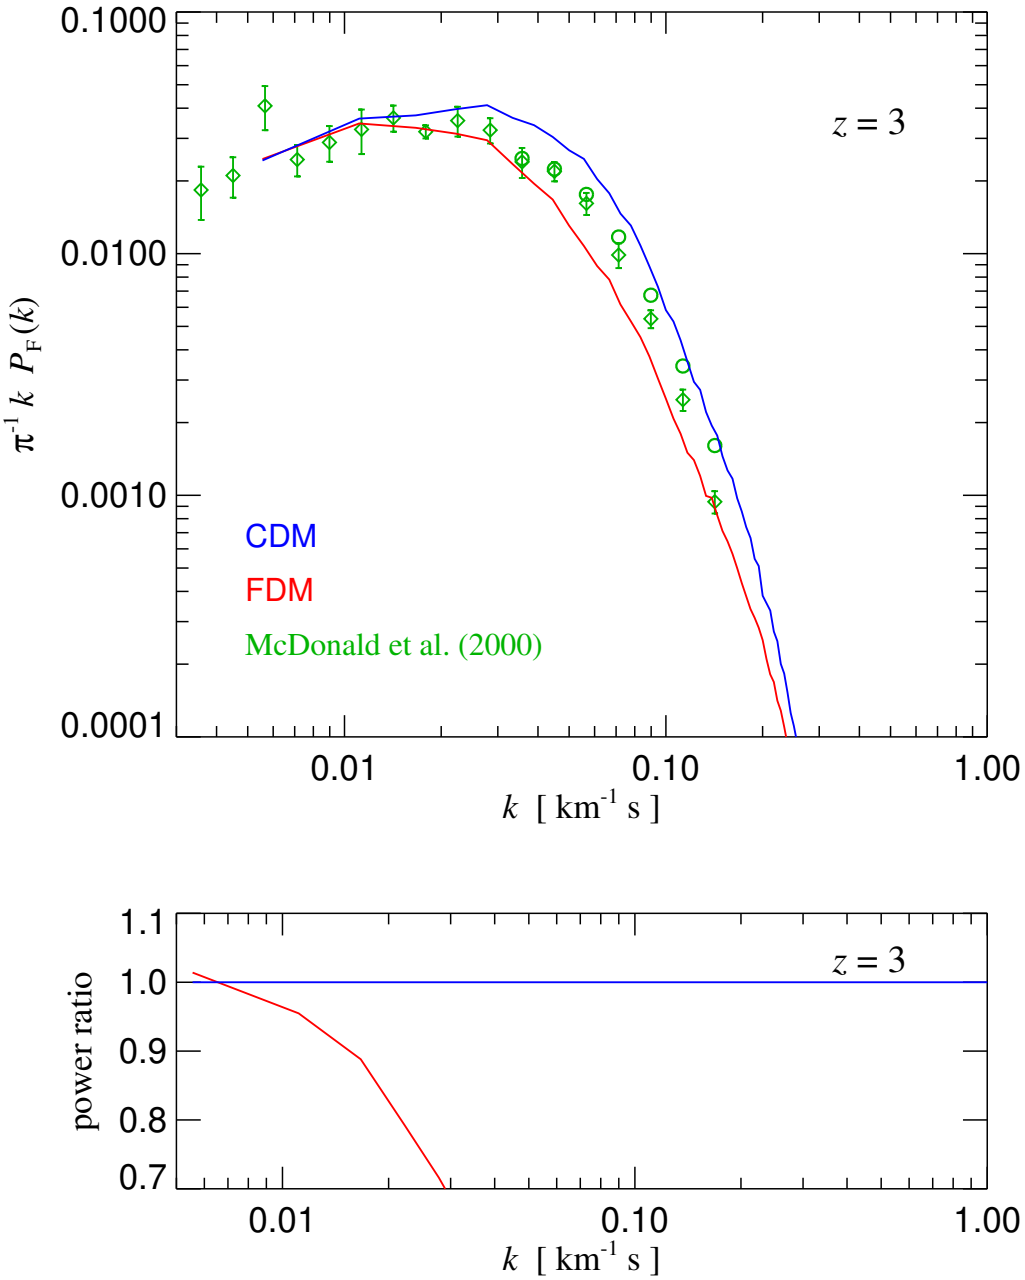

Figure 6.5: The flux power spectrum of FDM and CDM simulations with baryons, obtained by averaging over the synthetic absorption spectra of 1000 randomly selected lines of sight. Observational data from McDonald et al. (2000) is shown for comparison.

7 Conclusion and outlook

In this dissertation, I have presented the development, implementation, and analysis of the largest wave numerical simulations (numerically solving the SP equations) of cosmic structure formation with FDM which have been performed to date. These simulations are pushing the boundary of what is feasible in the size of the simulation volume and computational grid, improving on previous efforts by a large factor. The achieved comoving box size of $10\,h^{-1}$ Mpc has made it possible for the first time to study larger-scale cosmological observables, such as the HMF or the Lyman-α forest, in a direct, fully self-consistent manner without relying on approximations. Not only that, but such extensive simulations can serve as a baseline to verify and calibrate simplified or approximate approaches, providing a robust underpinning.

It is arguably safe to say that the ideal goal of gaining a comprehensive understanding of structure formation with FDM, similar to the current state of ΛCDM, is much more difficult to achieve than originally hoped for at the beginning of this project. Due to the prohibitive computational cost and requirements imposed by the SP equations, it is highly unlikely that there will ever be a fully self-consistent simulation of FDM's wave behavior from large cosmological scales (hundreds of Mpc) down to the wavelengths within halos for the foreseeable future, especially for the larger FDM masses $> 10^{-21}$ eV increasingly being pushed by observational constraints.

However, this is not meant to diminish the achievements made as part of this dissertation! On the contrary, it could be said that a complete framework for wave dark matter simulations (not only limited to FDM!), including many of the "standard" tools which are by now taken for granted in the context of N-body simulations (such as a halo finder), was developed in only a few short years and proven in practice. At this point (fundamental limitations enforced by FDM's computational cost notwithstanding), the work in this dissertation, with the resulting AxiREPO code, in principle makes it possible to run any kind of realistic simulation of cosmic structure formation with FDM, including full models of baryons and galaxy formation! Instead, it is more a question of perspective, especially when one is used to the well-established

sphere of ΛCDM *N*-body simulations: There was simply a lot more groundwork and answering of fundamental questions necessary before setting out to perform "full-physics" simulations than there would be when starting from a mature and complete set of tools as in ΛCDM.

Although any kind of numerical work is bound to encounter technical challenges, it should be pointed out that these were especially severe and numerous in this endeavor of FDM simulations. These challenges were both fundamental and of a more "mundane" nature: Concerning the former, the immense computational requirements posed by FDM simulations have been demonstrated e. g. in chapter 4. However, this is not only a problem for the final "production runs" of a simulation suite, which will obviously be limited in scope by the computational cost. Rather, it already takes a severe toll during the development stage, because it is *simply not possible to run low-resolution FDM simulations*. This is very much unlike the experience with CDM and *N*-body methods, where even a comparatively poor resolution gives reliable results on the largest scales. The fact that similar low-resolution comparisons cannot be done for FDM is a major obstacle for experimentation and debugging. For example, how it is easy to imagine the difficulties involved in developing a reliable measurement for the FDM HMF when the only kind of simulation that features enough halos for any HMF to speak of is an 8640^3 grid, which forces any tests to be run on a massive 10 TiB snapshot. Each iteration in the development process occurs on a time frame measured in *days* when operating on the full snapshot data set.

Similarly, the fact that the analysis of the simulation data constantly involves such large amounts of data makes the process very painful. All tools and scripts created to analyze the simulations eventually had to evolve into MPI-parallelized distributed-memory versions in order to be able to deal with the amount of data on reasonable time scales. Even for the companion CDM simulations, which were done using existing conventions tools, the high resolution used to be able to compare the FDM and CDM power spectra frequently brought analysis software to its limits when working on 2048^3 particle simulations, once again taking inordinate amounts of time for standard operations (if they were possible at all without running out of memory) like visualization or computing density profiles. So even here, the simplest solution was often to start from scratch instead of attempting to make existing tools work.

As a consequence, it must be noted that the simulations presented in this dissertation have *not nearly been analyzed to the fullest extent!* Large simulation projects like the Millennium (Springel et al. 2005), EAGLE (Schaye et al. 2015), or IllustrisTNG (Springel et al. 2018) simulations are are often explored over many years by large

teams of researchers from different (sub-)fields. On the flipside, there are only so many aspects a single person can investigate in a limited amount of time. However, this should also serve as an illustration of the great potential of the FDM simulations presented here: There are still countless facets to be discovered "hidden in the data", such as a more detailed and systematic exploration of the halo catalogs and the properties of halos, e. g. shapes, concentrations, rotation curves, the interaction of cores and granules with the remaining halo, merger events and halo stripping, assembly bias, or the impact of the environment on halo evolution; the investigation of vortices: their properties and prospects for detection; the evolution and history across time, and the question of hierarchical growth; or higher-order statistics.

Concerning the state of the FDM model as a whole, it is undeniable that observational evidence, especially recently, is increasingly pushing the allowed parameter space towards rather large particle masses, in a regime where the "fuzzy" properties of FDM would be relegated to (astrophysically) tiny, practically unobservable scales – at least for the simplest version of the model. While the methods developed in this dissertation have paved the way for a detailed exploration of all aspects of the FDM (within the computational constraints) in realistic scenarios including baryonic physics and galaxy formation, the fact that the observational constraints have all but excluded the parameter space where this is computationally feasible is rather sobering. Indeed, it seems that the biggest challenge for numerical approaches to FDM in the near future will be the issue raised in Dalal and Kravtsov (2022): With their claims of constraining m up to 10^{-19} eV, simulations on any kind of cosmological scales in this mass range are completely unthinkable, as the required resolution increases by three orders of magnitude compared to the simulations shown in this work. However, the numerical method use to obtain this bound is not self-consistent, so if this could be robustly confirmed using a full numerical treatment, it would spell the end for FDM, which becomes almost indistinguishable from CDM in the remaining allowed region.

Supplement to chapter 4

Since there is no analytical solution to the general time-dependent Schrödinger–Poisson equations, reliability of the code in the general case (beyond toy examples and limiting cases) can only be checked through comparison with other implementations. Thus, in Chan, Ferreira, May, et al. (2022), we compared the dark matter density fluctuations at $z = 0$ in a test simulation using two independent implementations. The codes are independently developed, but adopted the same pseudo-spectral splitting method in second order.

We ran a cosmological fuzzy dark matter simulation separately with identical initial conditions generated by MUSIC with box size $L = 10\,h^{-1}\,\mathrm{Mpc}$, particle mass $mc^2 = 2.5 \times 10^{-24}\,\mathrm{eV}$, and number of grid cells $N^3 = 1024^3$. The cosmological simulations are evolved until $z = 0$, and the density fluctuations are measured as the power spectrum shown in fig. A.1.

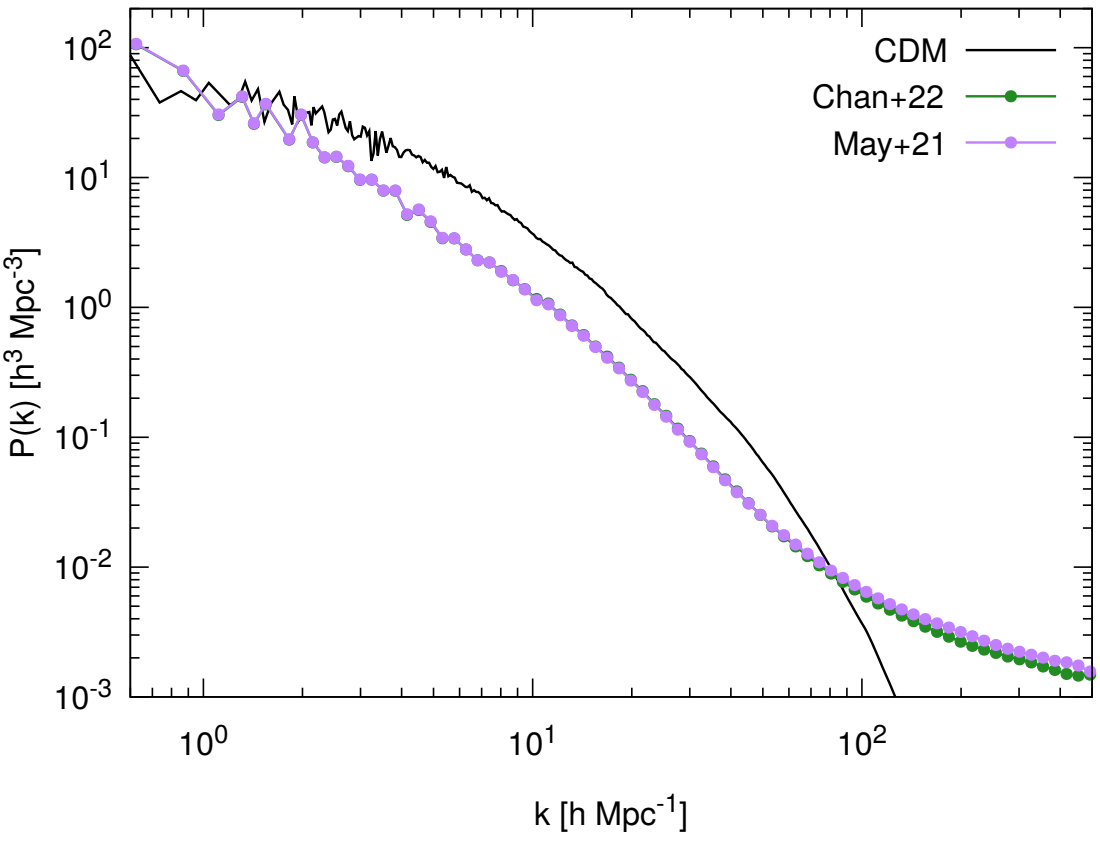

Figure A.1: Comparison of the power spectrum at $z = 0$ between the codes used in Chan, Ferreira, May, et al. (2022) for a cosmological test simulation.

Supplement to chapter 5

B

After performing the simulations with FDM ICs, I discovered that, unfortunately, a small error had occurred with the used input power spectra (generated by AXIONCAMB): Instead of generating the power spectrum at the correct starting redshift of $z = 127$, the power spectrum was generated for $z = 0$ and then scaled back to $z = 127$ using the (scale-independent) *CDM* growth factor. However, since the growth factor for FDM is scale-dependent, this procedure does not result in the correct power spectrum for an earlier redshift.

Figure B.1 shows how the shape of the linear FDM power spectrum compares for different redshifts. In this case, the red line ($z = 0$) was used (incorrectly) instead of the orange line ($z = 127$). Fortunately, the discrepancy is not very serious. Since FDM coincides with CDM on large scales, there is no error incurred there. The difference essentially boils down to a slightly increased suppression on scales smaller than $300\,h^{-1}$ kpc, accompanied by some changes in the details of the oscillations of the power spectrum on these scales. However, the order of magnitudes of the relative suppression involved (compared to CDM) are 10^{-12} and 10^{-9}, or in other words, structure is effectively completely suppressed on these scales in both cases. For all intents and purposes, the difference is negligible, and is erased by the vastly larger amount of power coming in from non-linear power transfer in the later stages of evolution (cf. section 5.4). Nevertheless, I would like to note this slight inconsistency in the ICs here for completeness.

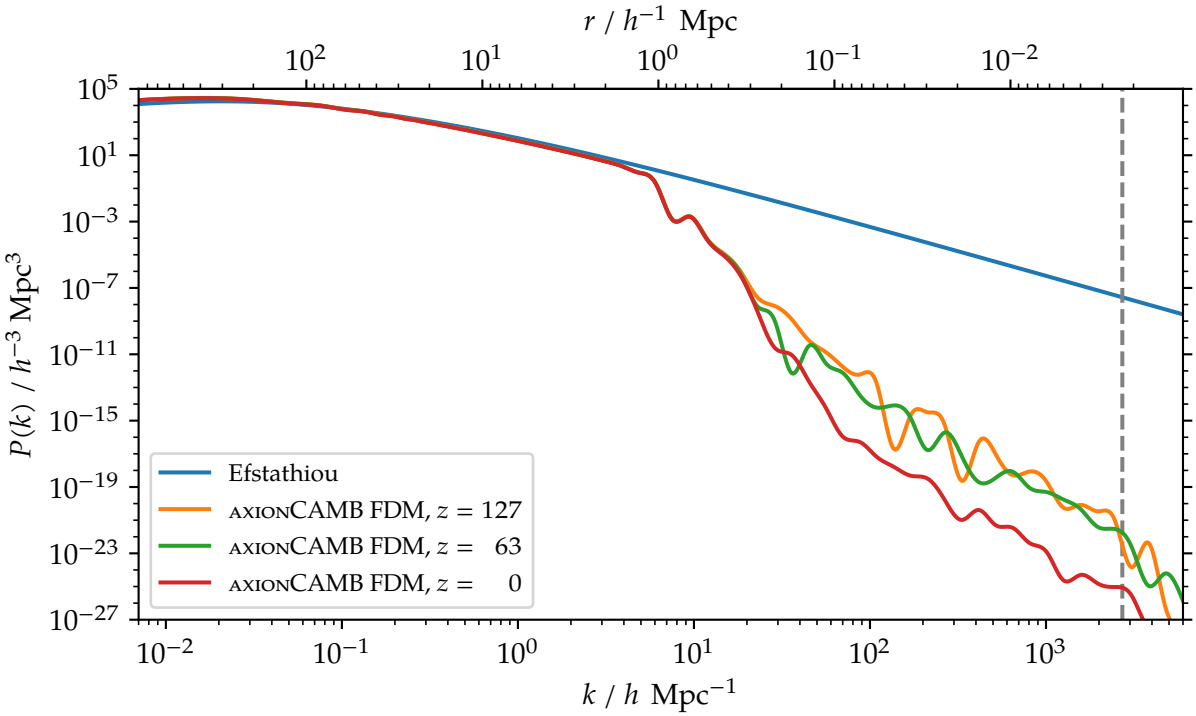

Figure B.1: Linear FDM power spectra ($mc^2 = 7 \times 10^{-23}$ eV) generated by axionCAMB for different redshifts z, with the Efstathiou, Sutherland, and Maddox (1990) and Efstathiou, Bond, and White (1992) power spectrum from eq. (5.8) for reference. For easier comparison, all power spectra were scaled to $z = 0$ using the linear CDM growth factor. The dashed vertical line indicates the grid resolution scale for a simulation in a $10\,h^{-1}$ Mpc box with an 8640^3 grid.

Bibliography

Planck Collaboration: Ade, P. A. R. et al. (2016). "Planck 2015 results. XIII. Cosmological parameters". In: *A&A* 594, A13 (Sept. 2016), A13. DOI: 10.1051/0004-6361/201525830. arXiv: 1502.01589 [astro-ph.CO].

Planck Collaboration: Aghanim, N., Y. Akrami, F. Arroja, et al. (2020). "Planck 2018 results. I. Overview and the cosmological legacy of Planck". In: *A&A* 641, A1 (Sept. 2020), A1. DOI: 10.1051/0004-6361/201833880. arXiv: 1807.06205 [astro-ph.CO].

Planck Collaboration: Aghanim, N., Y. Akrami, M. Ashdown, et al. (2020). "Planck 2018 results. VI. Cosmological parameters". In: *A&A* 641, A6 (Sept. 2020), A6. DOI: 10.1051/0004-6361/201833910. arXiv: 1807.06209 [astro-ph.CO].

BOSS Collaboration: Alam, Shadab et al. (2017). "The clustering of galaxies in the completed SDSS-III Baryon Oscillation Spectroscopic Survey: cosmological analysis of the DR12 galaxy sample". In: *Mon. Not. Roy. Astron. Soc.* 470.3 (2017), pp. 2617–2652. DOI: 10.1093/mnras/stx721. arXiv: 1607.03155 [astro-ph.CO].

Alexander, Stephon, Christian Capanelli, Elisa G. M. Ferreira, and Evan McDonough (2022). "Cosmic filament spin from dark matter vortices". In: *Physics Letters B* 833, 137298 (Oct. 2022), p. 137298. DOI: 10.1016/j.physletb.2022.137298. arXiv: 2111.03061 [astro-ph.CO].

Amin, Mustafa A., Mudit Jain, Rohith Karur, and Philip Mocz (2022). "Small-scale structure in vector dark matter". In: *J. Cosmology Astropart. Phys.* 2022.8, 014 (Aug. 2022), p. 014. DOI: 10.1088/1475-7516/2022/08/014. arXiv: 2203.11935 [astro-ph.CO].

Angulo, Raul E. and Oliver Hahn (2022). "Large-scale dark matter simulations". In: *Living Reviews in Computational Astrophysics* 8.1, 1 (Dec. 2022), p. 1. DOI: 10.1007/s41115-021-00013-z. arXiv: 2112.05165 [astro-ph.CO].

Angulo, Raul E., Oliver Hahn, and Tom Abel (2013). "The warm dark matter halo mass function below the cut-off scale". In: *MNRAS* 434.4 (Oct. 2013), pp. 3337–3347. DOI: 10.1093/mnras/stt1246. arXiv: 1304.2406 [astro-ph.CO].

Avila-Reese, Vladimir et al. (2001). "Formation and Structure of Halos in a Warm Dark Matter Cosmology". In: *ApJ* 559.2 (Oct. 2001), pp. 516–530. DOI: 10.1086/322411. arXiv: astro-ph/0010525 [astro-ph].

Barnes, Joshua E. and Piet Hut (1989). "Error Analysis of a Tree Code". In: *ApJS* 70 (June 1989), p. 389. DOI: 10.1086/191343.

Barrera-Hinojosa, Cristian and Baojiu Li (2020). "GRAMSES: a new route to general relativistic N-body simulations in cosmology. Part I. Methodology and code description". In: *J. Cosmology Astropart. Phys.* 2020.1, 007 (Jan. 2020), p. 007. DOI: 10.1088/1475-7516/2020/01/007. arXiv: 1905.08890 [astro-ph.CO].

Bett, Philip et al. (2007). "The spin and shape of dark matter haloes in the Millennium simulation of a Λ cold dark matter universe". In: *MNRAS* 376.1 (Mar. 2007), pp. 215–232. DOI: 10.1111/j.1365-2966.2007.11432.x. arXiv: astro-ph/0608607 [astro-ph].

Blanton, Michael R. et al. (2017). "Sloan Digital Sky Survey IV: Mapping the Milky Way, Nearby Galaxies, and the Distant Universe". In: *AJ* 154.1, 28 (July 2017), p. 28. DOI: 10.3847/1538-3881/aa7567. arXiv: 1703.00052 [astro-ph.GA].

Bolton, James S. and Martin G. Haehnelt (2007). "The nature and evolution of the highly ionized near-zones in the absorption spectra of z $\sim=$ 6 quasars". In: *MNRAS* 374.2 (Jan. 2007), pp. 493–514. DOI: 10.1111/j.1365-2966.2006.11176.x. arXiv: astro-ph/0607331 [astro-ph].

Bosch, Frank C. van den (2017). "Dissecting the Evolution of Dark Matter Subhaloes in the Bolshoi Simulation". In: *MNRAS* 468.1 (2017), pp. 885–909. DOI: 10.1093/mnras/stx520. arXiv: 1611.02657 [astro-ph.GA].

Boyarsky, A. et al. (2019). "Sterile neutrino Dark Matter". In: *Progress in Particle and Nuclear Physics* 104 (Jan. 2019), pp. 1–45. DOI: 10.1016/j.ppnp.2018.07.004. arXiv: 1807.07938 [hep-ph].

Boyd, J. P. (2001). *Chebyshev and Fourier Spectral Methods: Second Revised Edition*. Dover Books on Mathematics. Dover Publications, 2001. ISBN: 9780486411835.

Boylan-Kolchin, Michael, James S. Bullock, and Manoj Kaplinghat (2011). "Too big to fail? The puzzling darkness of massive Milky Way subhaloes". In: *MNRAS* 415.1 (July 2011), pp. L40–L44. DOI: 10.1111/j.1745-3933.2011.01074.x. arXiv: 1103.0007 [astro-ph.CO].

Broadhurst, Tom et al. (2020). "Ghostly galaxies as solitons of Bose-Einstein dark matter". In: *Phys. Rev. D* 101.8, 083012 (Apr. 2020), p. 083012. DOI: 10.1103/PhysRevD.101.083012. arXiv: 1902.10488 [astro-ph.CO].

Buchert, Thomas and Syksy Räsänen (2012). "Backreaction in Late-Time Cosmology". In: *Annual Review of Nuclear and Particle Science* 62.1 (Nov. 2012), pp. 57–79. DOI: 10.1146/annurev.nucl.012809.104435. arXiv: 1112.5335 [astro-ph.CO].

Bull, Philip et al. (2016). "Beyond Λ CDM: Problems, solutions, and the road ahead". In: *Physics of the Dark Universe* 12 (June 2016), pp. 56–99. DOI: 10.1016/j.dark.2016.02.001. arXiv: 1512.05356 [astro-ph.CO].

Bullock, James S. and Michael Boylan-Kolchin (2017). "Small-Scale Challenges to the ΛCDM Paradigm". In: *ARA&A* 55.1 (Aug. 2017), pp. 343–387. DOI: 10.1146/annurev-astro-091916-055313. arXiv: 1707.04256 [astro-ph.CO].

Burkert, A. (2020). "Fuzzy Dark Matter and Dark Matter Halo Cores". In: *ApJ* 904.2, 161 (Dec. 2020), p. 161. DOI: 10.3847/1538-4357/abb242. arXiv: 2006.11111 [astro-ph.GA].

Calabrese, Erminia and David N. Spergel (2016). "Ultra-Light Dark Matter in Ultra-Faint Dwarf Galaxies". In: *MNRAS* 460.4 (2016), pp. 4397–4402. DOI: 10.1093/mnras/stw1256. arXiv: 1603.07321 [astro-ph.CO].

Chan, Hei Yin Jowett, Elisa G. M. Ferreira, Simon May, et al. (2022). "The diversity of core-halo structure in the fuzzy dark matter model". In: *MNRAS* 511.1 (Mar. 2022), pp. 943–952. DOI: 10.1093/mnras/stac063. arXiv: 2110.11882 [astro-ph.CO].

Chen, Shu-Rong, Hsi-Yu Schive, and Tzihong Chiueh (2017). "Jeans Analysis for Dwarf Spheroidal Galaxies in Wave Dark Matter". In: *Mon. Not. Roy. Astron. Soc.* 468.2 (2017), pp. 1338–1348. DOI: 10.1093/mnras/stx449. arXiv: 1606.09030 [astro-ph.GA].

Clifton, Timothy, Pedro G. Ferreira, Antonio Padilla, and Constantinos Skordis (2012). "Modified gravity and cosmology". In: *Phys. Rep.* 513.1 (Mar. 2012), pp. 1–189. DOI: 10.1016/j.physrep.2012.01.001. arXiv: 1106.2476 [astro-ph.CO].

D'Onghia, Elena and George Lake (2004). "Cold Dark Matter's Small-Scale Crisis Grows Up". In: *ApJ* 612.2 (Sept. 2004), pp. 628–632. DOI: 10.1086/422794. arXiv: astro-ph/0309735 [astro-ph].

Dalal, Neal and Andrey Kravtsov (2022). "Excluding fuzzy dark matter with sizes and stellar kinematics of ultrafaint dwarf galaxies". In: *Phys. Rev. D* 106.6 (2022), p. 063517. DOI: 10.1103/PhysRevD.106.063517. arXiv: 2203.05750 [astro-ph.CO].

Del Popolo, Antonino and Morgan Le Delliou (2017). "Small Scale Problems of the ΛCDM Model: A Short Review". In: *Galaxies* 5.1 (Feb. 2017), p. 17. DOI: 10.3390/galaxies5010017. arXiv: 1606.07790 [astro-ph.CO].

Despali, Giulia et al. (2016). "The universality of the virial halo mass function and models for non-universality of other halo definitions". In: *MNRAS* 456.3 (Mar. 2016), pp. 2486–2504. DOI: 10.1093/mnras/stv2842. arXiv: 1507.05627 [astro-ph.CO].

Di Paolo, Chiara, Paolo Salucci, and Adnan Erkurt (2019). "The universal rotation curve of low surface brightness galaxies – IV. The interrelation between dark and luminous matter". In: *MNRAS* 490.4 (2019), pp. 5451–5477. DOI: 10.1093/mnras/stz2700. arXiv: 1805.07165.

Du, Xiaolong, Christoph Behrens, and Jens C. Niemeyer (2017). "Substructure of fuzzy dark matter haloes". In: *MNRAS* 465.1 (Feb. 2017), pp. 941–951. DOI: 10.1093/mnras/stw2724. arXiv: 1608.02575 [astro-ph.CO].

Du, Xiaolong, Christoph Behrens, Jens C. Niemeyer, and Bodo Schwabe (2017). "Core-halo mass relation of ultralight axion dark matter from merger history". In: *Phys. Rev. D* 95.4 (2017), p. 043519. DOI: 10.1103/PhysRevD.95.043519. arXiv: 1609.09414 [astro-ph.GA].

Edwards, Faber, Emily Kendall, Shaun Hotchkiss, and Richard Easther (2018). "PyUltraLight: a pseudo-spectral solver for ultralight dark matter dynamics". In: *J. Cosmology Astropart. Phys.* 2018.10, 027 (Oct. 2018), p. 027. DOI: 10.1088/1475-7516/2018/10/027. arXiv: 1807.04037 [astro-ph.CO].

Efstathiou, G., J. R. Bond, and S. D. M. White (1992). "COBE background radiation anisotropies and large-scale structure in the universe". In: *MNRAS* 258.1 (Sept. 1992), 1P–6P. DOI: 10.1093/mnras/258.1.1P.

Efstathiou, G., W. J. Sutherland, and S. J. Maddox (1990). "The cosmological constant and cold dark matter". In: *Nature* 348.6303 (Dec. 1990), pp. 705–707. DOI: 10.1038/348705a0.

Ferreira, Elisa G. M. (2021). "Ultra-light dark matter". In: *A&A Rev.* 29.1, 7 (Dec. 2021), p. 7. DOI: 10.1007/s00159-021-00135-6. arXiv: 2005.03254 [astro-ph.CO].

Fiaschi, Juri, Michael Klasen, and Simon May (2019). "Singlet-doublet fermion and triplet scalar dark matter with radiative neutrino masses". In: *Journal of High Energy Physics* 2019.5, 15 (May 2019), p. 15. DOI: 10.1007/JHEP05(2019)015. arXiv: 1812.11133 [hep-ph].

Frenk, Carlos S. and Simon D. M. White (2012). "Dark matter and cosmic structure". In: *Annalen der Physik* 524.9-10 (Oct. 2012), pp. 507–534. DOI: 10.1002/andp.201200212. arXiv: 1210.0544 [astro-ph.CO].

Frigo, Matteo and Steven G. Johnson (2005). "The Design and Implementation of FFTW3". In: *Proceedings of the IEEE* 93.2 (Feb. 2005), pp. 216–231. DOI: 10.1109/jproc.2004.840301.

Galárraga-Espinosa, Daniela et al. (2020). "Populations of filaments from the distribution of galaxies in numerical simulations". In: *A&A* 641, A173 (Sept. 2020), A173. DOI: 10.1051/0004-6361/202037986. arXiv: 2003.09697 [astro-ph.CO].

Gao, Liang and Tom Theuns (2007). "Lighting the Universe with Filaments". In: *Science* 317.5844 (Sept. 2007), p. 1527. DOI: 10.1126/science.1146676. arXiv: 0709.2165 [astro-ph].

Gao, Liang, Tom Theuns, and Volker Springel (2015). "Star-forming filaments in warm dark matter models". In: *MNRAS* 450.1 (June 2015), pp. 45–52. DOI: 10.1093/mnras/stv643. arXiv: 1403.2475 [astro-ph.CO].

Garny, Mathias, Thomas Konstandin, and Henrique Rubira (2020). "The Schrödinger-Poisson method for Large-Scale Structure". In: *J. Cosmology Astropart. Phys.* 2020.4, 003 (Apr. 2020), p. 003. DOI: 10.1088/1475-7516/2020/04/003. arXiv: 1911.04505 [astro-ph.CO].

Geller, Margaret J. and John P. Huchra (1989). "Mapping the Universe". In: *Science* 246.4932 (Nov. 1989), pp. 897–903. DOI: 10.1126/science.246.4932.897.

González-Morales, Alma X., David J. E. Marsh, Jorge Peñarrubia, and Luis A. Ureña-López (2017). "Unbiased constraints on ultralight axion mass from dwarf spheroidal galaxies". In: *Mon. Not. Roy. Astron. Soc.* 472.2 (2017), pp. 1346–1360. DOI: 10.1093/mnras/stx1941. arXiv: 1609.05856 [astro-ph.CO].

Grand, Robert J. J. et al. (2021). "Determining the full satellite population of a Milky Way-mass halo in a highly resolved cosmological hydrodynamic simulation". In: *MNRAS* 507.4 (Nov. 2021), pp. 4953–4967. DOI: 10.1093/mnras/stab2492. arXiv: 2105.04560 [astro-ph.GA].

Grin, Daniel, David J. E. Marsh, and Renee Hlozek (2022). *axionCAMB: Modification of the CAMB Boltzmann code*. Astrophysics Source Code Library, record ascl:2203.026. Mar. 2022.

Gross, E. P. (1961). "Structure of a quantized vortex in boson systems". In: *Il Nuovo Cimento* 20.3 (May 1961), pp. 454–477. DOI: 10.1007/BF02731494.

Hayashi, Kohei, Elisa G. M. Ferreira, and Hei Yin Jowett Chan (2021). "Narrowing the Mass Range of Fuzzy Dark Matter with Ultrafaint Dwarfs". In: *ApJ* 912.1, L3 (May 2021), p. L3. DOI: 10.3847/2041-8213/abf501. arXiv: 2102.05300 [astro-ph.CO].

Hellwing, Wojciech A. et al. (2016). "The effect of baryons on redshift space distortions and cosmic density and velocity fields in the EAGLE simulation". In: *MNRAS* 461.1 (Sept. 2016), pp. L11–L15. DOI: 10.1093/mnrasl/slw081. arXiv: 1603.03328 [astro-ph.CO].

Hinshaw, G. et al. (2013). "Nine-year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Cosmological Parameter Results". In: *ApJS* 208.2, 19 (Oct. 2013), p. 19. DOI: 10.1088/0067-0049/208/2/19. arXiv: 1212.5226 [astro-ph.CO].

Hlozek, Renée, Daniel Grin, David J. E. Marsh, and Pedro G. Ferreira (2015). "A search for ultralight axions using precision cosmological data". In: *Phys. Rev. D* 91.10, 103512 (May 2015), p. 103512. DOI: 10.1103/PhysRevD.91.103512. arXiv: 1410.2896 [astro-ph.CO].

Hložek, Renée et al. (2017). "Future CMB tests of dark matter: Ultralight axions and massive neutrinos". In: *Phys. Rev. D* 95.12, 123511 (June 2017), p. 123511. DOI: 10.1103/PhysRevD.95.123511. arXiv: 1607.08208 [astro-ph.CO].

Hopkins, Philip F. (2017). "A New Public Release of the GIZMO Code". In: *arXiv e-prints*, arXiv:1712.01294 (Dec. 2017), arXiv:1712.01294. arXiv: 1712.01294 [astro-ph.IM].

Hu, Wayne, Rennan Barkana, and Andrei Gruzinov (2000). "Fuzzy Cold Dark Matter: The Wave Properties of Ultralight Particles". In: *Phys. Rev. Lett.* 85.6 (Aug. 2000), pp. 1158–1161. DOI: 10.1103/PhysRevLett.85.1158. arXiv: astro-ph/0003365 [astro-ph].

Hui, Lam (2021). "Wave Dark Matter". In: *ARA&A* 59 (Sept. 2021), pp. 247–289. DOI: 10.1146/annurev-astro-120920-010024. arXiv: 2101.11735 [astro-ph.CO].

Hui, Lam, Austin Joyce, Michael J. Landry, and Xinyu Li (2021). "Vortices and waves in light dark matter". In: *J. Cosmology Astropart. Phys.* 2021.1, 011 (Jan. 2021), p. 011. DOI: 10.1088/1475-7516/2021/01/011. arXiv: 2004.01188 [astro-ph.CO].

Hui, Lam, Jeremiah P. Ostriker, Scott Tremaine, and Edward Witten (2017). "Ultralight scalars as cosmological dark matter". In: *Phys. Rev. D* 95.4, 043541 (Feb. 2017), p. 043541. DOI: 10.1103/PhysRevD.95.043541. arXiv: 1610.08297 [astro-ph.CO].

Jenkins, A. et al. (1998). "Evolution of Structure in Cold Dark Matter Universes". In: *ApJ* 499.1 (May 1998), pp. 20–40. DOI: 10.1086/305615. arXiv: astro-ph/9709010 [astro-ph].

Jenkins, A. et al. (2001). "The mass function of dark matter haloes". In: *MNRAS* 321.2 (Feb. 2001), pp. 372–384. DOI: 10.1046/j.1365-8711.2001.04029.x. arXiv: astro-ph/0005260 [astro-ph].

Jungman, Gerard, Marc Kamionkowski, and Kim Griest (1996). "Supersymmetric dark matter". In: *Phys. Rept.* 267 (1996), pp. 195–373. DOI: 10.1016/0370-1573(95)00058-5. arXiv: hep-ph/9506380 [hep-ph].

Kapteyn, J. C. (1922). "First Attempt at a Theory of the Arrangement and Motion of the Sidereal System". In: *ApJ* 55 (May 1922), p. 302. DOI: 10.1086/142670.

Kulkarni, Mihir and Jeremiah P. Ostriker (2022). "What is the halo mass function in a fuzzy dark matter cosmology?" In: *MNRAS* 510.1 (Feb. 2022), pp. 1425–1430. DOI: 10.1093/mnras/stab3520. arXiv: 2011.02116 [astro-ph.CO].

Laguë, Alex et al. (2021). "Evolving ultralight scalars into non-linearity with Lagrangian perturbation theory". In: *MNRAS* 504.2 (June 2021), pp. 2391–2404. DOI: 10.1093/mnras/stab601. arXiv: 2004.08482 [astro-ph.CO].

Li, Xinyu, Lam Hui, and Greg L. Bryan (2019). "Numerical and perturbative computations of the fuzzy dark matter model". In: *Phys. Rev. D* 99.6, 063509 (Mar. 2019), p. 063509. DOI: 10.1103/PhysRevD.99.063509. arXiv: 1810.01915 [astro-ph.CO].

Li, Xinyu, Lam Hui, and Tomer D. Yavetz (2021). "Oscillations and random walk of the soliton core in a fuzzy dark matter halo". In: *Phys. Rev. D* 103.2, 023508 (Jan. 2021), p. 023508. DOI: 10.1103/PhysRevD.103.023508. arXiv: 2011.11416.

Lovell, Mark R. et al. (2014). "The properties of warm dark matter haloes". In: *MNRAS* 439.1 (Mar. 2014), pp. 300–317. DOI: 10.1093/mnras/stt2431. arXiv: 1308.1399 [astro-ph.CO].

Lundmark, Knut (1930). "Über die Bestimmung der Entfernungen, Dimensionen, Massen und Dichtigkeit fur die nächstgelegenen anagalacktischen Sternsysteme." In: *Meddelanden fran Lunds Astronomiska Observatorium Serie I* 125 (Jan. 1930), pp. 1–13.

Madelung, Erwin (1927). "Quantentheorie in hydrodynamischer Form". In: *Zeitschrift für Physik* 40.3-4 (Mar. 1927), pp. 322–326. DOI: 10.1007/bf01400372.

Mannheim, Philip D. (2006). "Alternatives to dark matter and dark energy". In: *Progress in Particle and Nuclear Physics* 56.2 (Apr. 2006), pp. 340–445. DOI: 10.1016/j.ppnp.2005.08.001. arXiv: astro-ph/0505266 [astro-ph].

Marsh, David J. E. (2016). "Axion cosmology". In: *Phys. Rep.* 643 (July 2016), pp. 1–79. DOI: 10.1016/j.physrep.2016.06.005. arXiv: 1510.07633 [astro-ph.CO].

Marsh, David J. E. and Joseph Silk (2014). "A model for halo formation with axion mixed dark matter". In: *MNRAS* 437.3 (Jan. 2014), pp. 2652–2663. DOI: 10.1093/mnras/stt2079. arXiv: 1307.1705 [astro-ph.CO].

Martin, G. et al. (2019). "The formation and evolution of low-surface-brightness galaxies". In: *MNRAS* 485.1 (May 2019), pp. 796–818. DOI: 10.1093/mnras/stz356. arXiv: 1902.04580.

May, Simon and Volker Springel (2021). "Structure formation in large-volume cosmological simulations of fuzzy dark matter: impact of the non-linear dynamics". In: *MNRAS* 506.2 (Sept. 2021), pp. 2603–2618. DOI: 10.1093/mnras/stab1764. arXiv: 2101.01828 [astro-ph.CO].

– (2023). "The halo mass function and filaments in full cosmological simulations with fuzzy dark matter". In: *MNRAS* 524.3 (Sept. 2023), pp. 4256–4274. DOI: 10.1093/mnras/stad2031. arXiv: 2209.14886 [astro-ph.CO].

McDonald, Patrick et al. (2000). "The Observed Probability Distribution Function, Power Spectrum, and Correlation Function of the Transmitted Flux in the Lyα Forest". In: *ApJ* 543.1 (Nov. 2000), pp. 1–23. DOI: 10.1086/317079. arXiv: astro-ph/9911196 [astro-ph].

Mina, Mattia, David F. Mota, and Hans A. Winther (2020). "SCALAR: an AMR code to simulate axion-like dark matter models". In: *A&A* 641, A107 (Sept. 2020), A107. DOI: 10.1051/0004-6361/201936272. arXiv: 1906.12160 [physics.comp-ph].

– (2022). "Solitons in the dark: First approach to non-linear structure formation with fuzzy dark matter". In: *A&A* 662, A29 (June 2022), A29. DOI: 10.1051/0004-6361/202038876.

Mocz, Philip et al. (2017). "Galaxy formation with BECDM - I. Turbulence and relaxation of idealized haloes". In: *MNRAS* 471.4 (Nov. 2017), pp. 4559–4570. DOI: 10.1093/mnras/stx1887. arXiv: 1705.05845 [astro-ph.CO].

Mocz, Philip et al. (2018). "Schrödinger-Poisson-Vlasov-Poisson correspondence". In: *Phys. Rev. D* 97.8, 083519 (Apr. 2018), p. 083519. DOI: 10.1103/PhysRevD.97.083519. arXiv: 1801.03507 [astro-ph.CO].

Mocz, Philip et al. (2019). "First Star-Forming Structures in Fuzzy Cosmic Filaments". In: *Phys. Rev. Lett.* 123.14, 141301 (Oct. 2019), p. 141301. DOI: 10.1103/PhysRevLett.123.141301. arXiv: 1910.01653 [astro-ph.GA].

Mocz, Philip et al. (2020). "Galaxy formation with BECDM - II. Cosmic filaments and first galaxies". In: *MNRAS* 494.2 (May 2020), pp. 2027–2044. DOI: 10.1093/mnras/staa738. arXiv: 1911.05746 [astro-ph.CO].

Nakama, Tomohiro, Jens Chluba, and Marc Kamionkowski (2017). "Shedding light on the small-scale crisis with CMB spectral distortions". In: *Phys. Rev. D* 95.12, 121302 (June 2017), p. 121302. DOI: 10.1103/PhysRevD.95.121302. arXiv: 1703.10559 [astro-ph.CO].

Navarro, Julio F., Carlos S. Frenk, and Simon D. M. White (1996). "The Structure of Cold Dark Matter Halos". In: *ApJ* 462 (May 1996), p. 563. DOI: 10.1086/177173. arXiv: astro-ph/9508025 [astro-ph].

– (1997). "A Universal Density Profile from Hierarchical Clustering". In: *ApJ* 490.2 (Dec. 1997), pp. 493–508. DOI: 10.1086/304888. arXiv: astro-ph/9611107 [astro-ph].

Neto, Angelo F. et al. (2007). "The statistics of Λ CDM halo concentrations". In: *MNRAS* 381.4 (Nov. 2007), pp. 1450–1462. DOI: 10.1111/j.1365-2966.2007.12381.x. arXiv: 0706.2919 [astro-ph].

Niemeyer, Jens C. (2020). "Small-scale structure of fuzzy and axion-like dark matter". In: *Progress in Particle and Nuclear Physics* 113, 103787 (July 2020), p. 103787. DOI: 10.1016/j.ppnp.2020.103787. arXiv: 1912.07064 [astro-ph.CO].

Nori, Matteo and Marco Baldi (2018). "AX-GADGET: a new code for cosmological simulations of Fuzzy Dark Matter and Axion models". In: *MNRAS* 478.3 (Aug. 2018), pp. 3935–3951. DOI: 10.1093/mnras/sty1224. arXiv: 1801.08144 [astro-ph.CO].

– (2021). "Scaling relations of fuzzy dark matter haloes - I. Individual systems in their cosmological environment". In: *MNRAS* 501.1 (Feb. 2021), pp. 1539–1556. DOI: 10.1093/mnras/staa3772. arXiv: 2007.01316 [astro-ph.CO].

Nori, Matteo et al. (2019). "Lyman α forest and non-linear structure characterization in Fuzzy Dark Matter cosmologies". In: *MNRAS* 482.3 (Jan. 2019), pp. 3227–3243. DOI: 10.1093/mnras/sty2888. arXiv: 1809.09619 [astro-ph.CO].

Peacock, John A. et al. (2001). "A measurement of the cosmological mass density from clustering in the 2dF Galaxy Redshift Survey". In: *Nature* 410.6825 (Mar. 2001), pp. 169–173. arXiv: astro-ph/0103143 [astro-ph].

Peccei, R. D. and Helen R. Quinn (1977). "CP conservation in the presence of pseudoparticles". In: *Phys. Rev. Lett.* 38 (June 1977), pp. 1440–1443. DOI: 10.1103/PhysRevLett.38.1440.

Pillepich, Annalisa et al. (2018). "Simulating galaxy formation with the IllustrisTNG model". In: *MNRAS* 473.3 (Jan. 2018), pp. 4077–4106. DOI: 10.1093/mnras/stx2656. arXiv: 1703.02970 [astro-ph.GA].

Pitaevskii, L. P. (1961). "Vortex Lines in an Imperfect Bose Gas". In: *Journal of Experimental and Theoretical Physics* 13.2 (1961), p. 451.

Press, William H. and Paul Schechter (1974). "Formation of Galaxies and Clusters of Galaxies by Self-Similar Gravitational Condensation". In: *ApJ* 187 (Feb. 1974), pp. 425–438. DOI: 10.1086/152650.

Restrepo, Diego, Oscar Zapata, and Carlos E. Yaguna (2013). "Models with radiative neutrino masses and viable dark matter candidates". In: *Journal of High Energy Physics* 2013, 11 (Nov. 2013), p. 11. DOI: 10.1007/JHEP11(2013)011. arXiv: 1308.3655 [hep-ph].

Rogers, Keir K. and Hiranya V. Peiris (2021). "Strong Bound on Canonical Ultralight Axion Dark Matter from the Lyman-Alpha Forest". In: *Phys. Rev. Lett.* 126.7 (2021), p. 071302. DOI: 10.1103/PhysRevLett.126.071302. arXiv: 2007.12705 [astro-ph.CO].

Rubin, Vera C. and Jr. Ford W. Kent (1970). "Rotation of the Andromeda Nebula from a Spectroscopic Survey of Emission Regions". In: *ApJ* 159 (Feb. 1970), p. 379. DOI: 10.1086/150317.

Safarzadeh, Mohammadtaher and David N. Spergel (2020). "Ultra-light Dark Matter Is Incompatible with the Milky Way's Dwarf Satellites". In: *ApJ* 893.1, 21 (Apr. 2020), p. 21. DOI: 10.3847/1538-4357/ab7db2. arXiv: 1906.11848 [astro-ph.CO].

Sales, Laura V., Andrew Wetzel, and Azadeh Fattahi (2022). "Baryonic solutions and challenges for cosmological models of dwarf galaxies". In: *Nature Astronomy* 6 (June 2022), pp. 897–910. DOI: 10.1038/s41550-022-01689-w. arXiv: 2206.05295 [astro-ph.GA].

Salucci, Paolo et al. (2007). "The Universal Rotation Curve of Spiral Galaxies. 2. The Dark Matter Distribution out to the Virial Radius". In: *MNRAS* 378 (2007), pp. 41–47. DOI: 10.1111/j.1365-2966.2007.11696.x. arXiv: astro-ph/0703115.

Santos-Santos, Isabel M. E. et al. (2020). "Baryonic clues to the puzzling diversity of dwarf galaxy rotation curves". In: *MNRAS* 495.1 (June 2020), pp. 58–77. DOI: 10.1093/mnras/staa1072. arXiv: 1911.09116 [astro-ph.GA].

Schaller, Matthieu, Pedro Gonnet, Aidan B. G. Chalk, and Peter W. Draper (2016). "SWIFT: Using Task-Based Parallelism, Fully Asynchronous Communication, and Graph Partition-Based Domain Decomposition for Strong Scaling on more than 100,000 Cores". In: *Proceedings of the Platform for Advanced Scientific Computing Conference*. June 2016, 2, p. 2. DOI: 10.1145/2929908.2929916. arXiv: 1606.02738 [cs.DC].

Schaye, Joop et al. (2015). "The EAGLE project: simulating the evolution and assembly of galaxies and their environments". In: *MNRAS* 446.1 (Jan. 2015), pp. 521–554. DOI: 10.1093/mnras/stu2058. arXiv: 1407.7040 [astro-ph.GA].

Schive, Hsi-Yu, Tzihong Chiueh, and Tom Broadhurst (2014). "Cosmic structure as the quantum interference of a coherent dark wave". In: *Nature Physics* 10.7 (July 2014), pp. 496–499. DOI: 10.1038/nphys2996. arXiv: 1406.6586 [astro-ph.GA].

Schive, Hsi-Yu, Tzihong Chiueh, Tom Broadhurst, and Kuan-Wei Huang (2016). "Contrasting Galaxy Formation from Quantum Wave Dark Matter, ψDM, with ΛCDM, using Planck and Hubble Data". In: *ApJ* 818.1, 89 (Feb. 2016), p. 89. DOI: 10.3847/0004-637X/818/1/89. arXiv: 1508.04621 [astro-ph.GA].

Schive, Hsi-Yu et al. (2014). "Understanding the Core-Halo Relation of Quantum Wave Dark Matter from 3D Simulations". In: *Phys. Rev. Lett.* 113.26, 261302 (Dec. 2014), p. 261302. DOI: 10.1103/PhysRevLett.113.261302. arXiv: 1407.7762 [astro-ph.GA].

Schneider, Peter (2015). *Extragalactic Astronomy and Cosmology*. 2nd ed. Springer, 2015. ISBN: 978-3-642-54082-0. DOI: 10.1007/978-3-642-54083-7.

Schwabe, Bodo and Jens C. Niemeyer (2022). "Deep Zoom-In Simulation of a Fuzzy Dark Matter Galactic Halo". In: *Phys. Rev. Lett.* 128.18, 181301 (May 2022), p. 181301. DOI: 10.1103/PhysRevLett.128.181301. arXiv: 2110.09145 [astro-ph.CO].

Schwabe, Bodo, Jens C. Niemeyer, and Jan F. Engels (2016). "Simulations of solitonic core mergers in ultralight axion dark matter cosmologies". In: *Phys. Rev. D* 94.4 (2016), p. 043513. DOI: 10.1103/PhysRevD.94.043513. arXiv: 1606.05151 [astro-ph.CO].

Schwabe, Bodo et al. (2020). "Simulating mixed fuzzy and cold dark matter". In: *Phys. Rev. D* 102.8, 083518 (Oct. 2020), p. 083518. DOI: 10.1103/PhysRevD.102.083518. arXiv: 2007.08256 [astro-ph.CO].

Sheth, Ravi K. and Giuseppe Tormen (1999). "Large-scale bias and the peak background split". In: *MNRAS* 308.1 (Sept. 1999), pp. 119–126. DOI: 10.1046/j.1365-8711.1999.02692.x. arXiv: astro-ph/9901122 [astro-ph].

Sousbie, T. (2011). "The persistent cosmic web and its filamentary structure - I. Theory and implementation". In: *MNRAS* 414.1 (June 2011), pp. 350–383. DOI: 10.1111/j.1365-2966.2011.18394.x. arXiv: 1009.4015 [astro-ph.CO].

Sousbie, T., C. Pichon, and H. Kawahara (2011). "The persistent cosmic web and its filamentary structure - II. Illustrations". In: *MNRAS* 414.1 (June 2011), pp. 384–403.

DOI: 10.1111/j.1365-2966.2011.18395.x. arXiv: 1009.4014 [astro-ph.CO].

Springel, Volker (2010). "E pur si muove: Galilean-invariant cosmological hydrodynamical simulations on a moving mesh". In: *MNRAS* 401.2 (Jan. 2010), pp. 791–851. DOI: 10.1111/j.1365-2966.2009.15715.x. arXiv: 0901.4107 [astro-ph.CO].

– (2015). *N-GenIC: Cosmological structure initial conditions*. Astrophysics Source Code Library, record ascl:1502.003. Feb. 2015.

Springel, Volker, Carlos S. Frenk, and Simon D. M. White (2006). "The large-scale structure of the Universe". In: *Nature* 440.7088 (Apr. 2006), pp. 1137–1144. DOI: 10.1038/nature04805. arXiv: astro-ph/0604561 [astro-ph].

Springel, Volker and Lars Hernquist (2003). "Cosmological smoothed particle hydrodynamics simulations: a hybrid multiphase model for star formation". In: *MNRAS* 339.2 (Feb. 2003), pp. 289–311. DOI: 10.1046/j.1365-8711.2003.06206.x. arXiv: astro-ph/0206393 [astro-ph].

Springel, Volker, Rüdiger Pakmor, Oliver Zier, and Martin Reinecke (2021). "Simulating cosmic structure formation with the GADGET-4 code". In: *MNRAS* 506.2 (Sept. 2021), pp. 2871–2949. DOI: 10.1093/mnras/stab1855. arXiv: 2010.03567 [astro-ph.IM].

Springel, Volker et al. (2005). "Simulations of the formation, evolution and clustering of galaxies and quasars". In: *Nature* 435.7042 (June 2005), pp. 629–636. DOI: 10.1038/nature03597. arXiv: astro-ph/0504097 [astro-ph].

Springel, Volker et al. (2018). "First results from the IllustrisTNG simulations: matter and galaxy clustering". In: *MNRAS* 475.1 (Mar. 2018), pp. 676–698. DOI: 10.1093/mnras/stx3304. arXiv: 1707.03397 [astro-ph.GA].

Stücker, Jens, Raul E. Angulo, and Philipp Busch (2021). "The boosted potential". In: *MNRAS* 508.4 (Dec. 2021), pp. 5196–5216. DOI: 10.1093/mnras/stab2913. arXiv: 2107.13008 [astro-ph.CO].

Tinker, Jeremy et al. (2008). "Toward a Halo Mass Function for Precision Cosmology: The Limits of Universality". In: *ApJ* 688.2 (Dec. 2008), pp. 709–728. DOI: 10.1086/591439. arXiv: 0803.2706 [astro-ph].

Tulin, Sean and Hai-Bo Yu (2018). "Dark matter self-interactions and small scale structure". In: *Phys. Rep.* 730 (Feb. 2018), pp. 1–57. DOI: 10.1016/j.physrep.2017.11.004. arXiv: 1705.02358 [hep-ph].

Veltmaat, Jan and Jens C. Niemeyer (2016). "Cosmological particle-in-cell simulations with ultralight axion dark matter". In: *Phys. Rev. D* 94.12, 123523 (Dec. 2016),

p. 123523. DOI: 10.1103/PhysRevD.94.123523. arXiv: 1608.00802 [astro-ph.CO].

Veltmaat, Jan, Jens C. Niemeyer, and Bodo Schwabe (2018). "Formation and structure of ultralight bosonic dark matter halos". In: *Phys. Rev. D* 98.4, 043509 (Aug. 2018), p. 043509. DOI: 10.1103/PhysRevD.98.043509. arXiv: 1804.09647 [astro-ph.CO].

Veltmaat, Jan, Bodo Schwabe, and Jens C. Niemeyer (2020). "Baryon-driven growth of solitonic cores in fuzzy dark matter halos". In: *Phys. Rev. D* 101.8, 083518 (Apr. 2020), p. 083518. DOI: 10.1103/PhysRevD.101.083518. arXiv: 1911.09614 [astro-ph.CO].

Viel, Matteo, George D. Becker, James S. Bolton, and Martin G. Haehnelt (2013). "Warm dark matter as a solution to the small scale crisis: New constraints from high redshift Lyman-α forest data". In: *Phys. Rev. D* 88.4, 043502 (Aug. 2013), p. 043502. DOI: 10.1103/PhysRevD.88.043502. arXiv: 1306.2314 [astro-ph.CO].

Vogelsberger, Mark, Federico Marinacci, Paul Torrey, and Ewald Puchwein (2020). "Cosmological simulations of galaxy formation". In: *Nature Reviews Physics* 2.1 (Jan. 2020), pp. 42–66. DOI: 10.1038/s42254-019-0127-2. arXiv: 1909.07976 [astro-ph.GA].

Wang, Jie and Simon D. M. White (2007). "Discreteness effects in simulations of hot/warm dark matter". In: *MNRAS* 380.1 (Sept. 2007), pp. 93–103. DOI: 10.1111/j.1365-2966.2007.12053.x. arXiv: astro-ph/0702575 [astro-ph].

Weinberg, David H., Romeel Davé, Neal Katz, and Juna A. Kollmeier (2003). "The Lyman-α Forest as a Cosmological Tool". In: *The Emergence of Cosmic Structure*. Ed. by Stephen H. Holt and Christopher S. Reynolds. Vol. 666. American Institute of Physics Conference Series. May 2003, pp. 157–169. DOI: 10.1063/1.1581786. arXiv: astro-ph/0301186 [astro-ph].

Weinberg, David H. et al. (2015). "Cold dark matter: Controversies on small scales". In: *Proceedings of the National Academy of Science* 112.40 (Oct. 2015), pp. 12249–12255. DOI: 10.1073/pnas.1308716112. arXiv: 1306.0913 [astro-ph.CO].

Weinberger, Rainer, Volker Springel, and Rüdiger Pakmor (2020). "The AREPO Public Code Release". In: *ApJS* 248.2, 32 (June 2020), p. 32. DOI: 10.3847/1538-4365/ab908c. arXiv: 1909.04667 [astro-ph.IM].

Weinberger, Rainer et al. (2017). "Simulating galaxy formation with black hole driven thermal and kinetic feedback". In: *MNRAS* 465.3 (Mar. 2017), pp. 3291–3308. DOI: 10.1093/mnras/stw2944. arXiv: 1607.03486 [astro-ph.GA].

Widrow, Lawrence M. and Nick Kaiser (1993). "Using the Schrödinger Equation to Simulate Collisionless Matter". In: *ApJ* 416 (Oct. 1993), p. L71. DOI: 10.1086/187073.

Woo, Tak-Pong and Tzihong Chiueh (2009). "High-Resolution Simulation on Structure Formation with Extremely Light Bosonic Dark Matter". In: *ApJ* 697.1 (May 2009), pp. 850–861. DOI: 10.1088/0004-637X/697/1/850. arXiv: 0806.0232 [astro-ph].

X-ray: NASA/CXC/CfA/M.Markevitch et al.; Optical: NASA/STScI; Magellan/U.Arizona/D.Clowe et al.; Lensing Map: NASA/STScI; ESO WFI; Magellan/U.Arizona/D.Clowe et al. (2016). *Lensing map of the bullet cluster*. Aug. 12, 2016. URL: http://chandra.harvard.edu/photo/2006/1e0657/.

Yavetz, Tomer D., Xinyu Li, and Lam Hui (2021). *Construction of Wave Dark Matter Halos: Numerical Algorithm and Analytical Constraints*. Sept. 2021. arXiv: 2109.06125 [astro-ph.CO].

Zhang, Jiajun, Hantao Liu, and Ming-Chung Chu (2018). "Cosmological Simulation for Fuzzy Dark Matter Model". In: *Frontiers in Astronomy and Space Sciences* 5, 48 (Dec. 2018), p. 48. DOI: 10.3389/fspas.2018.00048. arXiv: 1809.09848 [astro-ph.CO].

Zhang, Jiajun et al. (2018). "The Importance of Quantum Pressure of Fuzzy Dark Matter on Lyα Forest". In: *ApJ* 863.1, 73 (Aug. 2018), p. 73. DOI: 10.3847/1538-4357/aacf3f. arXiv: 1708.04389 [astro-ph.CO].

Zoutendijk, Sebastiaan L. et al. (2021). "The MUSE-Faint survey. II. The dark-matter density profile of the ultra-faint dwarf galaxy Eridanus 2". In: *A&A* 651, A80 (July 2021), A80. DOI: 10.1051/0004-6361/202040239. arXiv: 2101.00253 [astro-ph.GA].

Zwicky, Fritz (1933). "Die Rotverschiebung von extragalaktischen Nebeln". In: *Helvetica Physica Acta* 6 (Jan. 1933), pp. 110–127.

Particle Data Group: Zyla, P. A. et al. (2020). "Review of Particle Physics". In: *Progress of Theoretical and Experimental Physics* 2020.8, 083C01 (Aug. 2020), p. 083C01. DOI: 10.1093/ptep/ptaa104.

Acknowledgements

It is probably safe to say that I had a very different kind of adventure than most others who embarked on their PhD journey at MPA. Not only was more than half of my time there mired in a global pandemic, but I also became a father during this time. Taken together with the fact that I was, for all intents and purposes, the only person working on the subject of fuzzy dark matter simulations at MPA, the result was that, although I thoroughly enjoyed my time there, I wish that I could have connected and integrated with everyone else in the institute more.

First and foremost, I would like to thank Elisa Ferreira for her incredible support and invaluable advice, while she was there at MPA and beyond. Further, my thanks extend to my advisor, Volker Springel, who, although chronically busy and short of time, has nevertheless always been friendly and insightful. And of course, to all the people who made social events like the WAWBeers (which, fortunately, survived the pandemic after a long break) and MPA in general a fun place to be!

Finally, my greatest thanks are due to my family: my wife Alex and our daughter Stella. Although I cannot say that they have always made it easy for me to work on my scientific projects and this dissertation, especially during the final weeks, they are, after all, the ones who ultimately make it all worth it!